



Kea Administrator Reference Manual Documentation

Release 3.0.2

Internet Systems Consortium

Oct 17, 2025

CONTENTS

1	Introduction	3
1.1	Supported Platforms	3
1.1.1	Regularly Tested Platforms	3
1.1.2	Best-Effort	3
1.1.3	Community-Maintained	4
1.1.4	Unsupported Platforms	4
1.2	Required Software at Runtime	4
1.3	Kea Software	5
2	Quick Start	7
2.1	Quick Start Guide Using tarball	7
2.2	Quick Start Guide Using Native Packages	7
2.3	Quick Start Guide Using Docker Containers	10
2.4	Quick Start Guide for DHCPv4 and DHCPv6 Services	10
2.5	Running the Kea Servers Directly	11
3	Installation	13
3.1	Packages	13
3.1.1	Installation From Cloudsmith Packages	13
3.1.2	Caveats When Upgrading Kea Packages	14
3.2	Installation Hierarchy	14
3.3	Build Requirements	15
3.4	Installation From Source	16
3.4.1	Download Tar File	16
3.4.2	Verify The Tar File Signature	16
3.4.3	Retrieve From Git	16
3.4.4	Set Up the Build	17
3.4.5	Build	17
3.4.6	Install	17
3.4.7	Cross-Building	18
3.5	DHCP Database Installation and Configuration	18
3.5.1	Building with MySQL Support	18
3.5.2	Building with PostgreSQL Support	18
3.6	Hammer Building Tool	19
3.7	Running Kea From a Non-root Account on Linux	21
3.8	Deprecated Features	22
3.8.1	Sysrepo 0.x or 1.x	22
4	Kea Database Administration	23
4.1	Databases and Schema Versions	23

4.2	The kea-admin Tool	23
4.3	Supported Backends	24
4.3.1	Memfile	24
4.3.1.1	Upgrading Memfile Lease Files From an Earlier Version of Kea	24
4.3.2	MySQL	24
4.3.2.1	MySQL 5.7 vs MySQL 8 vs MariaDB 10 and 11	25
4.3.2.2	First-Time Creation of the MySQL Database	25
4.3.2.3	Upgrading a MySQL Database From an Earlier Version of Kea	27
4.3.2.4	Improved Performance With MySQL	28
4.3.3	PostgreSQL	28
4.3.3.1	First-Time Creation of the PostgreSQL Database	28
4.3.3.2	Initialize the PostgreSQL Database Using kea-admin	30
4.3.3.3	Upgrading a PostgreSQL Engine From an Earlier Version	30
4.3.3.4	Upgrading a PostgreSQL Database From an Earlier Version of Kea	31
4.3.3.5	Improved Performance With PostgreSQL	31
4.3.4	Using Read-Only Databases With Host Reservations	31
4.3.5	Limitations Related to the Use of SQL Databases	31
4.3.5.1	Year 2038 Issue	31
5	Kea Security	33
5.1	TLS/HTTPS Support	33
5.1.1	Building Kea with TLS/HTTPS Support	34
5.1.2	TLS/HTTPS Configuration	34
5.1.3	OpenSSL Tuning	35
5.2	Secure Kea Control Agent	36
5.3	Securing a Kea Deployment	36
5.3.1	Component-Based Design	36
5.3.2	Limiting Application Permissions	37
5.3.3	Securing Kea Administrative Access	37
5.3.4	Securing Database Connections	38
5.3.5	Information Leakage Through Logging	38
5.3.6	Summary of Path Restrictions	39
5.3.7	Cryptography Components	39
5.3.8	TSIG Signatures	39
5.3.9	Raw Socket Support	40
5.3.10	Remote Administrative Access	40
5.3.11	Authentication for Kea's RESTful API	40
5.4	Kea Security Processes	40
5.4.1	Vulnerability Handling	41
5.4.2	Code Quality and Testing	41
5.4.3	Fuzz Testing	41
5.4.4	Release Integrity	42
5.4.5	Bus Factor	42
6	Kea Configuration	43
6.1	JSON Configuration	43
6.1.1	JSON Syntax	43
6.1.2	Comments and User Context	44
6.1.3	Simplified Notation	47
6.1.4	Configuration Files Inclusion	47
6.2	Kea Configuration Backend	48
6.2.1	Applicability	48
6.2.1.1	Example Scenario	48
6.2.2	Limitations and Warnings	49

6.2.2.1	Availability	49
6.2.2.2	Preparation is Required	49
6.2.2.3	Database Management	49
6.2.2.4	Config Conflicts	49
6.2.2.5	Incompatible Software	49
6.2.2.6	Limitations	50
6.2.2.7	Custom Options	50
6.2.3	Components	50
6.2.4	Installation	51
6.2.5	Configuration Sharing and Server Tags	51
7	Managing Kea with keactrl	53
7.1	Overview	53
7.2	Command Line Options	53
7.3	The keactrl Configuration File	53
7.4	Commands	55
7.5	Overriding the Server Selection	57
7.6	Native Packages and systemd	57
8	The Kea Control Agent	59
8.1	Overview of the Kea Control Agent	59
8.2	Configuration	59
8.3	Secure Connections	62
8.4	Starting and Stopping the Control Agent	62
8.5	Connecting to the Control Agent	63
9	The DHCPv4 Server	65
9.1	Starting and Stopping the DHCPv4 Server	65
9.2	DHCPv4 Server Configuration	66
9.2.1	Introduction	66
9.2.2	Lease Storage	69
9.2.2.1	Memfile - Basic Storage for Leases	69
9.2.2.2	Why Is Lease File Cleanup Necessary?	70
9.2.2.3	Lease Database Configuration	71
9.2.2.4	Tuning Database Timeouts	73
9.2.3	Hosts Storage	74
9.2.3.1	DHCPv4 Hosts Database Configuration	75
9.2.3.2	Using Read-Only Databases for Host Reservations With DHCPv4	76
9.2.3.3	Tuning Database Timeouts for Hosts Storage	77
9.2.4	Interface Configuration	77
9.2.5	Issues With Unicast Responses to DHCPINFORM	80
9.2.6	IPv4 Subnet Identifier	81
9.2.7	IPv4 Subnet Prefix	81
9.2.8	Configuration of IPv4 Address Pools	82
9.2.9	Sending T1 (Option 58) and T2 (Option 59)	84
9.2.10	Standard DHCPv4 Options	85
9.2.11	CableLabs Client Conf Suboptions	95
9.2.12	Custom DHCPv4 Options	96
9.2.13	DHCPv4 Private Options	98
9.2.14	DHCPv4 Vendor-Specific Options	101
9.2.15	Nested DHCPv4 Options (Custom Option Spaces)	105
9.2.16	Unspecified Parameters for DHCPv4 Option Configuration	107
9.2.17	Support for Long Options	107
9.2.18	Support for IPv6-Only Preferred Option	109

9.2.19	Stateless Configuration of DHCPv4 Clients	109
9.2.20	Client Classification in DHCPv4	109
9.2.20.1	Setting Fixed Fields in Classification	111
9.2.20.2	Using Vendor Class Information in Classification	111
9.2.20.3	Defining and Using Custom Classes	112
9.2.20.4	Additional Classification	112
9.2.21	DDNS for DHCPv4	114
9.2.21.1	DHCP-DDNS Server Connectivity	118
9.2.21.2	When Does the <code>kea-dhcp4</code> Server Generate a DDNS Request?	119
9.2.21.3	<code>kea-dhcp4</code> Name Generation for DDNS Update Requests	120
9.2.21.4	Sanitizing Client Host Name and FQDN Names	121
9.2.22	Next Server (<code>siaddr</code>)	123
9.2.23	Echoing Client-ID (RFC 6842)	123
9.2.24	Using Client Identifier and Hardware Address	123
9.2.25	Authoritative DHCPv4 Server Behavior	125
9.2.26	DHCPv4-over-DHCPv6: DHCPv4 Side	126
9.2.27	Sanity Checks in DHCPv4	127
9.2.28	Storing Extended Lease Information	128
9.2.29	Stash Agent Options	129
9.2.30	Multi-Threading Settings	129
9.2.31	Multi-Threading Settings With Different Database Backends	129
9.2.32	IPv6-Only Preferred Networks	130
9.2.33	Lease Caching	131
9.2.34	Temporary Allocation on DHCPDISCOVER	132
9.2.35	DNR (Discovery of Network-designated Resolvers) Options for DHCPv4	132
9.3	Host Reservations in DHCPv4	133
9.3.1	Address Reservation Types	135
9.3.2	Conflicts in DHCPv4 Reservations	135
9.3.3	Reserving a Hostname	136
9.3.4	Including Specific DHCPv4 Options in Reservations	137
9.3.5	Reserving Next Server, Server Hostname, and Boot File Name	139
9.3.6	Reserving Client Classes in DHCPv4	139
9.3.7	Storing Host Reservations in MySQL or PostgreSQL	141
9.3.8	Fine-Tuning DHCPv4 Host Reservation	141
9.3.9	Global Reservations in DHCPv4	147
9.3.10	Pool Selection with Client Class Reservations	148
9.3.11	Subnet Selection with Client Class Reservations	149
9.3.12	Multiple Reservations for the Same IP	151
9.3.13	Host Reservations as Basic Access Control	153
9.4	Shared Networks in DHCPv4	154
9.4.1	Local and Relayed Traffic in Shared Networks	157
9.4.2	Client Classification in Shared Networks	160
9.4.3	Host Reservations in Shared Networks	162
9.5	Server Identifier in DHCPv4	163
9.6	How the DHCPv4 Server Selects a Subnet for the Client	164
9.6.1	Using a Specific Relay Agent for a Subnet	165
9.6.2	Segregating IPv4 Clients in a Cable Network	165
9.7	Duplicate Addresses (DHCPDECLINE Support)	166
9.8	Statistics in the DHCPv4 Server	167
9.9	Management API for the DHCPv4 Server	172
9.9.1	UNIX Control Socket	172
9.9.2	HTTP/HTTPS Control Socket	174
9.10	User Contexts in IPv4	175
9.11	Supported DHCP Standards	176

9.11.1	Known RFC Violations	179
9.12	DHCPv4 Server Limitations	179
9.13	Kea DHCPv4 Server Examples	179
9.14	Configuration Backend in DHCPv4	179
9.14.1	Supported Parameters	179
9.14.2	Enabling the Configuration Backend	181
9.15	Kea DHCPv4 Compatibility Configuration Parameters	183
9.15.1	Lenient Option Parsing	184
9.15.2	Ignore DHCP Server Identifier	184
9.15.3	Ignore RAI Link Selection	184
9.15.4	Exclude First Last Addresses in /24 Subnets or Larger	185
9.16	Address Allocation Strategies in DHCPv4	185
9.16.1	Allocators Comparison	185
9.16.2	Iterative Allocator	186
9.16.3	Random Allocator	186
9.16.4	Free Lease Queue Allocator	186
10	The DHCPv6 Server	189
10.1	Starting and Stopping the DHCPv6 Server	189
10.2	DHCPv6 Server Configuration	190
10.2.1	Introduction	190
10.2.2	Lease Storage	193
10.2.2.1	Memfile - Basic Storage for Leases	193
10.2.2.2	Why Is Lease File Cleanup Necessary?	194
10.2.2.3	Lease Database Configuration	194
10.2.2.4	Tuning Database Timeouts	196
10.2.3	Hosts Storage	197
10.2.3.1	DHCPv6 Hosts Database Configuration	198
10.2.3.2	Using Read-Only Databases for Host Reservations with DHCPv6	200
10.2.3.3	Tuning Database Timeouts for Hosts Storage	200
10.2.4	Interface Configuration	201
10.2.5	IPv6 Subnet Identifier	202
10.2.6	IPv6 Subnet Prefix	203
10.2.7	Unicast Traffic Support	203
10.2.8	Configuration of IPv6 Address Pools	204
10.2.9	Subnet and Prefix Delegation Pools	205
10.2.10	Prefix Exclude Option	206
10.2.11	Standard DHCPv6 Options	207
10.2.12	Common Software46 Options	215
10.2.12.1	Software46 Container Options	215
10.2.12.2	S46 Rule Option	216
10.2.12.3	S46 BR Option	216
10.2.12.4	S46 DMR Option	217
10.2.12.5	S46 IPv4/IPv6 Address Binding Option	217
10.2.12.6	S46 Port Parameters	217
10.2.13	DNR (Discovery of Network-designated Resolvers) Options for DHCPv6	217
10.2.14	NTP Server Suboptions	220
10.2.15	Custom DHCPv6 Options	220
10.2.16	DHCPv6 Vendor-Specific Options	222
10.2.17	Nested DHCPv6 Options (Custom Option Spaces)	225
10.2.18	Unspecified Parameters for DHCPv6 Option Configuration	226
10.2.19	Controlling the Values Sent for T1 and T2 Times	227
10.2.20	IPv6 Subnet Selection	228
10.2.21	Rapid Commit	229

10.2.22	DHCPv6 Relays	229
10.2.23	Relay-Supplied Options	230
10.2.24	Client Classification in DHCPv6	230
10.2.24.1	Defining and Using Custom Classes	231
10.2.24.2	Additional Classification	233
10.2.25	DDNS for DHCPv6	234
10.2.25.1	DHCP-DDNS Server Connectivity	237
10.2.25.2	When Does the kea-dhcp6 Server Generate a DDNS Request?	238
10.2.25.3	kea-dhcp6 Name Generation for DDNS Update Requests	240
10.2.25.4	Sanitizing Client FQDN Names	241
10.2.26	DHCPv4-over-DHCPv6: DHCPv6 Side	242
10.2.27	Sanity Checks in DHCPv6	244
10.2.28	Storing Extended Lease Information	244
10.2.29	Multi-Threading Settings	246
10.2.30	Multi-Threading Settings With Different Database Backends	246
10.2.31	Lease Caching	247
10.3	Host Reservations in DHCPv6	248
10.3.1	Address/Prefix Reservation Types	249
10.3.2	Conflicts in DHCPv6 Reservations	250
10.3.3	Reserving a Hostname	251
10.3.4	Including Specific DHCPv6 Options in Reservations	252
10.3.5	Reserving Client Classes in DHCPv6	253
10.3.6	Storing Host Reservations in MySQL or PostgreSQL	255
10.3.7	Fine-Tuning DHCPv6 Host Reservation	255
10.3.8	Global Reservations in DHCPv6	261
10.3.9	Pool Selection with Client Class Reservations	262
10.3.10	Subnet Selection with Client Class Reservations	264
10.3.11	Multiple Reservations for the Same IP	265
10.3.12	Host Reservations as Basic Access Control	267
10.4	Shared Networks in DHCPv6	268
10.4.1	Local and Relayed Traffic in Shared Networks	272
10.4.2	Client Classification in Shared Networks	275
10.4.3	Host Reservations in Shared Networks	277
10.5	Server Identifier in DHCPv6	278
10.6	DHCPv6 Data Directory	280
10.7	Stateless DHCPv6 (INFORMATION-REQUEST Message)	281
10.8	Support for RFC 7550 (now part of RFC 8415)	281
10.9	Using a Specific Relay Agent for a Subnet	282
10.10	Segregating IPv6 Clients in a Cable Network	283
10.11	MAC/Hardware Addresses in DHCPv6	284
10.12	Duplicate Addresses (DHCPDECLINE Support)	285
10.13	Address Registration (RFC 9686 Support)	286
10.14	Statistics in the DHCPv6 Server	287
10.15	Management API for the DHCPv6 Server	293
10.15.1	UNIX Control Socket	294
10.15.2	HTTP/HTTPS Control Socket	295
10.16	User Contexts in IPv6	297
10.17	Supported DHCPv6 Standards	298
10.18	DHCPv6 Server Limitations	299
10.19	Kea DHCPv6 Server Examples	300
10.20	Configuration Backend in DHCPv6	300
10.20.1	Supported Parameters	300
10.20.2	Enabling the Configuration Backend	302
10.21	Kea DHCPv6 Compatibility Configuration Parameters	304

10.21.1	Lenient Option Parsing	304
10.22	Allocation Strategies in DHCPv6	304
10.22.1	Allocators Comparison	305
10.22.2	Iterative Allocator	305
10.22.3	Random Allocator	306
10.22.4	Free Lease Queue Allocator (Prefix Delegation Only)	306
11	Database Connectivity	309
12	Lease Expiration	311
12.1	Lease Reclamation	311
12.2	Lease Reclamation Configuration Parameters	312
12.3	Configuring Lease Reclamation	312
12.4	Configuring Lease Affinity	314
12.5	Reclaiming Expired Leases via Command	315
13	Congestion Handling	317
13.1	What is Congestion?	317
13.2	Configuring Congestion Handling	317
14	The DHCP-DDNS Server	319
14.1	Overview	319
14.1.1	DNS Server Selection	319
14.1.2	Conflict Resolution	319
14.1.3	Dual-Stack Environments	320
14.2	Starting and Stopping the DHCP-DDNS Server	320
14.3	Configuring the DHCP-DDNS Server	321
14.3.1	Global Server Parameters	322
14.3.2	Management API for the D2 Server	322
14.3.2.1	UNIX Control Socket	323
14.3.2.2	HTTP/HTTPS Control Socket	324
14.3.3	TSIG Key List	325
14.3.4	Forward DDNS	326
14.3.4.1	Adding Forward DDNS Domains	327
14.3.4.1.1	Adding Forward DNS Servers	327
14.3.5	Reverse DDNS	328
14.3.5.1	Adding Reverse DDNS Domains	328
14.3.5.1.1	Adding Reverse DNS Servers	329
14.3.5.2	Per-DNS-Server TSIG Keys	330
14.3.6	User Contexts in DDNS	331
14.3.7	Example DHCP-DDNS Server Configuration	332
14.4	DHCP-DDNS Server Statistics	334
14.4.1	NCR Statistics	334
14.4.2	DNS Update Statistics	334
14.4.3	Per-TSIG-Key DNS Update Statistics	335
14.5	DHCP-DDNS Server Limitations	335
14.6	Supported Standards	335
15	The LFC Process	337
15.1	Overview	337
15.2	Command-Line Options	337
16	Client Classification	339
16.1	Client Classification Overview	339
16.1.1	Classification Steps	340

16.2	Built-in Client Classes	341
16.3	Using Expressions in Classification	342
16.3.1	Predicates	345
16.3.2	Logical Operators	346
16.3.3	Substring	346
16.3.4	Concat	346
16.3.5	Split	347
16.3.6	Ifelse	347
16.3.7	Hexstring	347
16.4	Configuring Classes	347
16.4.1	Template Classes	350
16.5	Using Static Host Reservations in Classification	352
16.6	Configuring Subnets With Class Information	352
16.7	Configuring Pools With Class Information	353
16.8	Class Priority	355
16.9	Option Class-Tagging	356
16.10	Classes and Hooks	358
16.11	Debugging Expressions	358
17	Hook Libraries	361
17.1	Introduction	361
17.2	Installing Hook Packages	361
17.3	Configuring Hook Libraries	363
17.3.1	Order of Configuration:	365
17.3.2	User Contexts in Hooks	365
17.3.3	Parked-Packet Limit	367
17.4	Available Hook Libraries	368
17.5	libdhcp_bootp.so: Support for BOOTP Clients	371
17.5.1	BOOTP Hooks Limitations	372
17.6	libdhcp_cb_cmds.so: Configuration Backend Commands	372
17.6.1	Command Structure	373
17.6.2	Control Commands for DHCP Servers	374
17.6.3	Metadata	374
17.6.4	The remote-server4-del, remote-server6-del Commands	375
17.6.5	The remote-server4-get, remote-server6-get Commands	375
17.6.6	The remote-server4-get-all, remote-server6-get-all Commands	376
17.6.7	The remote-server4-set, remote-server6-set Commands	377
17.6.8	The remote-global-parameter4-del, remote-global-parameter6-del Commands	378
17.6.9	The remote-global-parameter4-get, remote-global-parameter6-get Commands	378
17.6.10	The remote-global-parameter4-get-all, remote-global-parameter6-get-all Commands	380
17.6.11	The remote-global-parameter4-set, remote-global-parameter6-set Commands	381
17.6.12	The remote-network4-del, remote-network6-del Commands	381
17.6.13	The remote-network4-get, remote-network6-get Commands	382
17.6.14	The remote-network4-list, remote-network6-list Commands	382
17.6.15	The remote-network4-set, remote-network6-set Commands	384
17.6.16	The remote-option-def4-del, remote-option-def6-del Commands	385
17.6.17	The remote-option-def4-get, remote-option-def6-get Commands	386
17.6.18	The remote-option-def4-get-all, remote-option-def6-get-all Commands	386
17.6.19	The remote-option-def4-set, remote-option-def6-set Commands	387
17.6.20	The remote-option4-global-del, remote-option6-global-del Commands	388
17.6.21	The remote-option4-global-get, remote-option6-global-get Commands	388
17.6.22	The remote-option4-global-get-all, remote-option6-global-get-all Commands	389
17.6.23	The remote-option4-global-set, remote-option6-global-set Commands	389

17.6.24	The remote-option4-network-del, remote-option6-network-del Commands . . .	390
17.6.25	The remote-option4-network-set, remote-option6-network-set Commands . . .	391
17.6.26	The remote-option6-pd-pool-del Command	391
17.6.27	The remote-option6-pd-pool-set Command	392
17.6.28	The remote-option4-pool-del, remote-option6-pool-del Commands	393
17.6.29	The remote-option4-pool-set, remote-option6-pool-set Commands	393
17.6.30	The remote-option4-subnet-del, remote-option6-subnet-del Commands	394
17.6.31	The remote-option4-subnet-set, remote-option6-subnet-set Commands	395
17.6.32	The remote-subnet4-del-by-id, remote-subnet6-del-by-id Commands	395
17.6.33	The remote-subnet4-del-by-prefix, remote-subnet6-del-by-prefix Commands .	396
17.6.34	The remote-subnet4-get-by-id, remote-subnet6-get-by-id Commands	396
17.6.35	The remote-subnet4-get-by-prefix, remote-subnet6-get-by-prefix Commands .	396
17.6.36	The remote-subnet4-list, remote-subnet6-list Commands	397
17.6.37	The remote-subnet4-set, remote-subnet6-set Commands	398
17.6.38	The remote-class4-del, remote-class6-del Commands	400
17.6.39	The remote-class4-get, remote-class6-get Commands	400
17.6.40	The remote-class4-get-all, remote-class6-get-all Commands	401
17.6.41	The remote-class4-set, remote-class6-set Commands	402
17.7	libdhcp_class_cmds.so: Class Commands	403
17.7.1	The class-add Command	404
17.7.2	The class-update Command	404
17.7.3	The class-del Command	405
17.7.4	The class-list Command	406
17.7.5	The class-get Command	406
17.8	libdhcp_ddns_tuning.so: DDNS Tuning	407
17.8.1	Procedural Host-Name Generation	408
17.8.1.1	DHCPv4 Host-Name Generation	409
17.8.1.2	DHCPv6 Host-Name Generation	409
17.8.2	Skipping DDNS Updates	410
17.9	libdhcp_flex_id.so: Flexible Identifier for Host Reservations	411
17.9.1	The replace-client-id Flag	413
17.9.2	The ignore-iaid Flag	414
17.10	libdhcp_flex_option.so: Flexible Option Actions for Option Value Settings .	415
17.11	libddns_gss_tsig.so: Sign DNS Updates With GSS-TSIG	418
17.12	libdhcp_ha.so: High Availability Outage Resilience for Kea Servers	418
17.12.1	Supported Configurations	418
17.12.2	Clocks on Active Servers	420
17.12.3	HTTPS Support	420
17.12.4	Server States	423
17.12.5	Scope Transition in a Partner-Down Case	426
17.12.6	Load-Balancing Configuration	427
17.12.7	Load Balancing With Advanced Classification	431
17.12.8	Hot-Standby Configuration	433
17.12.9	Passive-Backup Configuration	434
17.12.10	Lease Information Sharing	436
17.12.11	Controlling Lease-Page Size Limit	437
17.12.12	Timeouts	437
17.12.13	Pausing the HA State Machine	438
17.12.14	Control Agent Configuration	440
17.12.15	Multi-Threaded Configuration (HA+MT)	441
17.12.16	Parked-Packet Limit	443
17.12.17	Controlled Shutdown and Maintenance of DHCP Servers	443
17.12.18	Control Commands for High Availability	444
17.12.18.1	The ha-sync Command	444

17.12.18.2	The ha-scopes Command	445
17.12.18.3	The ha-continue Command	446
17.12.18.4	The ha-heartbeat Command	446
17.12.18.5	The status-get Command	447
17.12.18.6	The ha-maintenance-start Command	450
17.12.18.7	The ha-maintenance-cancel Command	450
17.12.18.8	The ha-maintenance-notify Command	450
17.12.18.9	The ha-reset Command	451
17.12.18.10	The ha-sync-complete-notify Command	451
17.12.19	Hub and Spoke Configuration	452
17.13	libdhcp_host_cache.so: Host Cache Reservations for Improved Performance	456
17.13.1	The cache-flush Command	457
17.13.2	The cache-clear Command	457
17.13.3	The cache-size Command	457
17.13.4	The cache-write Command	457
17.13.5	The cache-load Command	458
17.13.6	The cache-get Command	458
17.13.7	The cache-get-by-id Command	458
17.13.8	The cache-insert Command	459
17.13.9	The cache-remove Command	460
17.14	libdhcp_host_cmds.so: Host Commands	460
17.14.1	The subnet-id Parameter	462
17.14.2	The operation-target Parameter	462
17.14.3	The reservation-add Command	462
17.14.4	The reservation-get Command	464
17.14.5	The reservation-get-all Command	466
17.14.6	The reservation-get-page Command	467
17.14.7	The reservation-get-by-address Command	469
17.14.8	The reservation-get-by-hostname Command	473
17.14.9	The reservation-get-by-id Command	474
17.14.10	The reservation-del Command	475
17.14.11	The reservation-update Command	477
17.14.12	General Mentions	479
17.15	libdhcp_lease_cmds.so: Lease Commands for Easier Lease Management	479
17.15.1	The lease4-add, lease6-add Commands	480
17.15.2	The lease6-bulk-apply Command	482
17.15.3	The lease4-get, lease6-get Commands	484
17.15.4	The lease4-get-all, lease6-get-all Commands	486
17.15.5	The lease4-get-page, lease6-get-page Commands	487
17.15.6	The lease4-get-by-*, lease6-get-by-* Commands	489
17.15.7	The lease4-del, lease6-del Commands	490
17.15.8	The lease4-update, lease6-update Commands	491
17.15.9	The lease4-wipe, lease6-wipe Commands	492
17.15.10	The lease4-resend-ddns, lease6-resend-ddns Commands	492
17.15.11	The lease4-write, lease6-write Commands	493
17.15.12	Binding Variables	493
17.16	libdhcp_lease_query.so: Leasequery Support	495
17.16.1	DHCPv4 Leasequery	495
17.16.2	DHCPv4 Leasequery Configuration	496
17.16.3	DHCPv6 Leasequery	497
17.16.4	DHCPv6 Leasequery Configuration	499
17.16.5	DHCPv4 Bulk Leasequery	499
17.16.6	DHCPv6 Bulk Leasequery	500
17.16.7	Bulk Leasequery Configuration	501

17.16.8	Updating Existing Leases in SQL Lease Backends	503
17.17	<code>libdhcp_legal_log.so</code> : Forensic Logging	504
17.17.1	Log File Naming	504
17.17.2	Configuring the Forensic Logging Hooks	505
17.17.3	DHCPv4 Log Entries	508
17.17.4	DHCPv6 Log Entries	512
17.17.5	Database Backend	518
17.17.6	Syslog Backend	519
17.18	<code>libdhcp_limits.so</code> : Limits to Manage Lease Allocation and Packet Processing	519
17.18.1	Configuration	519
17.18.2	Lease Limiting	521
17.18.3	Rate Limiting	521
17.19	<code>libdhcp_mysql.so</code> : Database Backend for MySQL	522
17.19.1	MySQL Configuration Backend	522
17.20	<code>libdhcp_perfmom.so</code> : Performance Monitoring	523
17.20.1	Overview	523
17.20.2	Passive Event Logging	523
17.20.3	Duration Monitoring	524
17.20.4	Statistics Reporting	525
17.20.5	Alarms	526
17.20.6	API Commands	526
17.20.6.1	The <code>perfmom-control</code> Command	526
17.20.6.2	The <code>perfmom-get-all-durations</code> Command	526
17.20.7	Configuration	528
17.21	<code>libdhcp_ping_check.so</code> : Ping Check	530
17.21.1	Overview	530
17.21.2	Configuration	531
17.22	<code>libdhcp_pgsq1.so</code> : Database Backend for PostgreSQL	532
17.22.1	PostgreSQL Configuration Backend	532
17.23	<code>libdhcp_radius.so</code> : RADIUS Server Support	533
17.24	<code>libdhcp_rbac.so</code> : Role-Based Access Control	533
17.24.1	Role-Based Access Control (RBAC) Overview	533
17.24.2	Role-Based Access Control Configuration	534
17.24.2.1	Role Assignment	534
17.24.2.2	Role Configuration	534
17.24.2.3	API Commands	534
17.24.2.4	Access Control Lists	535
17.24.2.5	Response Filters	535
17.24.2.6	Global Parameters	535
17.24.3	Sample Configuration	536
17.24.4	Accept/Reject Algorithm	538
17.24.5	Custom Hook Commands and Command Redefinition	538
17.24.6	Extensive Example	539
17.25	<code>libdhcp_run_script.so</code> : Run Script Support for External Hook Scripts	542
17.26	<code>libdhcp_stat_cmds.so</code> : Statistics Commands for Supplemental Lease Statistics	555
17.26.1	The <code>stat-lease4-get</code> , <code>stat-lease6-get</code> Commands	556
17.27	<code>libdhcp_subnet_cmds.so</code> : Subnet Commands to Manage Subnets and Shared Networks	558
17.27.1	The <code>subnet4-list</code> Command	559
17.27.2	The <code>subnet6-list</code> Command	560
17.27.3	The <code>subnet4-get</code> Command	560
17.27.4	The <code>subnet6-get</code> Command	561
17.27.5	The <code>subnet4-add</code> Command	562
17.27.6	The <code>subnet6-add</code> Command	563
17.27.7	The <code>subnet4-update</code> Command	564

17.27.8	The subnet6-update Command	565
17.27.9	The subnet4-del Command	565
17.27.10	The subnet6-del Command	566
17.27.11	The subnet4-delta-add Command	567
17.27.12	The subnet6-delta-add Command	568
17.27.13	The subnet4-delta-del Command	570
17.27.14	The subnet6-delta-del Command	571
17.27.15	The network4-list, network6-list Commands	572
17.27.16	The network4-get, network6-get Commands	573
17.27.17	The network4-add, network6-add Commands	574
17.27.18	The network4-del, network6-del Commands	575
17.27.19	The network4-subnet-add, network6-subnet-add Commands	576
17.27.20	The network4-subnet-del, network6-subnet-del Commands	577
17.28	libdhcp_user_chk.so: User Check	577
18	Statistics	579
18.1	Statistics Overview	579
18.2	Statistics Lifecycle	579
18.3	Commands for Manipulating Statistics	580
18.3.1	The statistic-get Command	580
18.3.2	The statistic-reset Command	581
18.3.3	The statistic-remove Command	581
18.3.4	The statistic-get-all Command	581
18.3.5	The statistic-reset-all Command	594
18.3.6	The statistic-remove-all Command	595
18.3.7	The statistic-sample-age-set Command	595
18.3.8	The statistic-sample-age-set-all Command	595
18.3.9	The statistic-sample-count-set Command	596
18.3.10	The statistic-sample-count-set-all Command	596
18.4	Time Series	596
19	Management API	597
19.1	Data Syntax	597
19.2	Control Agent Command Response Format	600
19.3	Using the Control Channel	601
19.4	Commands Supported by Both the DHCPv4 and DHCPv6 Servers	602
19.4.1	The build-report Command	602
19.4.2	The config-get Command	602
19.4.3	The config-hash-get Command	602
19.4.4	The config-reload Command	603
19.4.5	The config-test Command	603
19.4.6	The config-write Command	604
19.4.7	The leases-reclaim Command	604
19.4.8	The list-commands Command	604
19.4.9	The config-set Command	605
19.4.10	The shutdown Command	605
19.4.11	The dhcp-disable Command	606
19.4.12	The dhcp-enable Command	607
19.4.13	The status-get Command	607
19.4.14	The server-tag-get Command:	608
19.4.15	The config-backend-pull Command:	608
19.4.16	The version-get Command	608
19.5	Commands Supported by the DHCPv4 Server	608
19.5.1	The subnet4-select-test Command	608

19.5.2	The subnet4o6-select-test Command	609
19.6	Commands Supported by the DHCPv6 Server	609
19.6.1	The subnet6-select-test Command	609
19.7	Commands Supported by the D2 Server	610
19.8	Commands Supported by the Control Agent	610
19.9	Migration from the Control Agent	611
20	Logging	613
20.1	Logging Configuration	613
20.1.1	Loggers	613
20.1.1.1	The name (string) Logger	613
20.1.1.2	The severity (string) Logger	619
20.1.1.3	The debuglevel (integer) Logger	619
20.1.1.4	The output-options (list) Logger	619
20.1.1.4.1	The output (string) Option	620
20.1.1.4.2	The flush (boolean) Option	620
20.1.1.4.3	The maxsize (integer) Option	620
20.1.1.4.4	The maxver (integer) Option	620
20.1.1.4.5	The pattern (string) Option	620
20.1.2	Logging Message Format	621
20.1.2.1	Example Logger Configurations	623
20.1.3	Logging During Kea Startup	624
20.2	Logging Levels	625
21	The Kea Shell	627
21.1	Overview of the Kea Shell	627
21.2	Shell Usage	627
21.3	TLS Support	629
22	Integration With External Systems	631
22.1	YANG/NETCONF	631
22.1.1	Overview	631
22.1.2	Installing NETCONF	631
22.1.2.1	Installing libyang From Sources	632
22.1.2.2	Installing sysrepo From Sources	632
22.1.2.3	Installing libyang-cpp From Sources	632
22.1.2.4	Installing sysrepo-cpp From Sources	632
22.1.3	Compiling With NETCONF	632
22.1.4	Quick Sysrepo Overview	633
22.1.5	Supported YANG Models	636
22.1.6	Using the NETCONF Agent	636
22.1.7	Configuration	637
22.1.8	A kea-netconf Configuration Example	638
22.1.9	Starting and Stopping the NETCONF Agent	641
22.1.10	A Step-by-Step NETCONF Agent Operation Example	641
22.1.10.1	Setup of NETCONF Agent Operation Example	642
22.1.10.2	Example of Error Handling in NETCONF Operation	644
22.1.10.3	NETCONF Operation Example with Two Pools	645
22.1.10.4	NETCONF Operation Example with Two Subnets	646
22.1.10.5	NETCONF Operation Example With Logging	646
22.1.10.6	Migrating YANG Data From a Prior Sysrepo Version	648
22.2	GSS-TSIG	648
22.2.1	GSS-TSIG Overview	648
22.2.2	GSS-TSIG Compilation	650

22.2.3	GSS-TSIG Deployment	650
22.2.3.1	Kerberos 5 Setup	650
22.2.3.2	BIND 9 with GSS-TSIG Configuration	653
22.2.3.3	Windows Active Directory Configuration	654
22.2.3.4	GSS-TSIG Troubleshooting	655
22.2.4	Using GSS-TSIG	655
22.2.4.1	GSS-TSIG Automatic Key Removal	662
22.2.4.2	GSS-TSIG Configuration for Deployment	662
22.2.5	GSS-TSIG Statistics	663
22.2.6	GSS-TSIG Commands	663
22.2.6.1	The <code>gss-tsig-get-all</code> Command	663
22.2.6.2	The <code>gss-tsig-get</code> Command	664
22.2.6.3	The <code>gss-tsig-list</code> Command	665
22.2.6.4	The <code>gss-tsig-key-get</code> Command	665
22.2.6.5	The <code>gss-tsig-key-expire</code> Command	666
22.2.6.6	The <code>gss-tsig-key-del</code> Command	666
22.2.6.7	The <code>gss-tsig-purge-all</code> Command	667
22.2.6.8	The <code>gss-tsig-purge</code> Command	667
22.2.6.9	The <code>gss-tsig-rekey-all</code> Command	667
22.2.6.10	The <code>gss-tsig-rekey</code> Command	668
22.3	RADIUS	668
22.3.1	RADIUS Overview	668
22.3.2	RADIUS Hook Library Configuration	669
22.3.3	RADIUS Server Setup Example	674
22.3.4	RADIUS Workflows for Lease Allocation	675
22.3.4.1	Parked-Packet Limit	675
22.3.5	Differences Between RADIUS Hook Libraries Prior To Kea 2.4.0 and As Of 2.6.0	675
23	Monitoring Kea With Stork	679
23.1	Kea Statistics in Grafana	679
24	API Reference	681
24.1	<code>build-report</code>	683
24.2	<code>cache-clear</code>	683
24.3	<code>cache-flush</code>	684
24.4	<code>cache-get</code>	684
24.5	<code>cache-get-by-id</code>	685
24.6	<code>cache-insert</code>	686
24.7	<code>cache-load</code>	687
24.8	<code>cache-remove</code>	688
24.9	<code>cache-size</code>	689
24.10	<code>cache-write</code>	689
24.11	<code>class-add</code>	690
24.12	<code>class-del</code>	691
24.13	<code>class-get</code>	691
24.14	<code>class-list</code>	692
24.15	<code>class-update</code>	693
24.16	<code>config-backend-pull</code>	693
24.17	<code>config-get</code>	694
24.18	<code>config-hash-get</code>	695
24.19	<code>config-reload</code>	695
24.20	<code>config-set</code>	696
24.21	<code>config-test</code>	697
24.22	<code>config-write</code>	697

24.23	dhcp-disable	698
24.24	dhcp-enable	699
24.25	extended-info4-upgrade	700
24.26	extended-info6-upgrade	700
24.27	gss-tsig-get	701
24.28	gss-tsig-get-all	702
24.29	gss-tsig-key-del	703
24.30	gss-tsig-key-expire	704
24.31	gss-tsig-key-get	704
24.32	gss-tsig-list	705
24.33	gss-tsig-purge	706
24.34	gss-tsig-purge-all	707
24.35	gss-tsig-rekey	707
24.36	gss-tsig-rekey-all	708
24.37	ha-continue	709
24.38	ha-heartbeat	709
24.39	ha-maintenance-cancel	710
24.40	ha-maintenance-notify	711
24.41	ha-maintenance-start	711
24.42	ha-reset	712
24.43	ha-scopes	713
24.44	ha-sync	713
24.45	ha-sync-complete-notify	714
24.46	lease4-add	715
24.47	lease4-del	716
24.48	lease4-get	716
24.49	lease4-get-all	717
24.50	lease4-get-by-client-id	718
24.51	lease4-get-by-hostname	719
24.52	lease4-get-by-hw-address	720
24.53	lease4-get-page	721
24.54	lease4-resend-ddns	722
24.55	lease4-update	723
24.56	lease4-wipe	723
24.57	lease4-write	724
24.58	lease6-add	725
24.59	lease6-bulk-apply	725
24.60	lease6-del	727
24.61	lease6-get	728
24.62	lease6-get-all	729
24.63	lease6-get-by-duid	730
24.64	lease6-get-by-hostname	731
24.65	lease6-get-page	732
24.66	lease6-resend-ddns	733
24.67	lease6-update	733
24.68	lease6-wipe	734
24.69	lease6-write	735
24.70	leases-reclaim	735
24.71	list-commands	736
24.72	network4-add	737
24.73	network4-del	738
24.74	network4-get	739
24.75	network4-list	740
24.76	network4-subnet-add	741

24.77	network4-subnet-del	741
24.78	network6-add	742
24.79	network6-del	743
24.80	network6-get	744
24.81	network6-list	745
24.82	network6-subnet-add	746
24.83	network6-subnet-del	747
24.84	perfmon-control	747
24.85	perfmon-get-all-durations	748
24.86	remote-class4-del	750
24.87	remote-class4-get	751
24.88	remote-class4-get-all	752
24.89	remote-class4-set	753
24.90	remote-class6-del	754
24.91	remote-class6-get	755
24.92	remote-class6-get-all	756
24.93	remote-class6-set	757
24.94	remote-global-parameter4-del	758
24.95	remote-global-parameter4-get	759
24.96	remote-global-parameter4-get-all	760
24.97	remote-global-parameter4-set	761
24.98	remote-global-parameter6-del	762
24.99	remote-global-parameter6-get	763
24.100	remote-global-parameter6-get-all	764
24.101	remote-global-parameter6-set	765
24.102	remote-network4-del	766
24.103	remote-network4-get	767
24.104	remote-network4-list	768
24.105	remote-network4-set	769
24.106	remote-network6-del	770
24.107	remote-network6-get	771
24.108	remote-network6-list	772
24.109	remote-network6-set	773
24.110	remote-option-def4-del	774
24.111	remote-option-def4-get	775
24.112	remote-option-def4-get-all	776
24.113	remote-option-def4-set	777
24.114	remote-option-def6-del	778
24.115	remote-option-def6-get	779
24.116	remote-option-def6-get-all	780
24.117	remote-option-def6-set	781
24.118	remote-option4-global-del	782
24.119	remote-option4-global-get	783
24.120	remote-option4-global-get-all	784
24.121	remote-option4-global-set	785
24.122	remote-option4-network-del	786
24.123	remote-option4-network-set	787
24.124	remote-option4-pool-del	788
24.125	remote-option4-pool-set	789
24.126	remote-option4-subnet-del	790
24.127	remote-option4-subnet-set	791
24.128	remote-option6-global-del	792
24.129	remote-option6-global-get	793
24.130	remote-option6-global-get-all	794

24.131remote-option6-global-set	795
24.132remote-option6-network-del	796
24.133remote-option6-network-set	797
24.134remote-option6-pd-pool-del	798
24.135remote-option6-pd-pool-set	799
24.136remote-option6-pool-del	800
24.137remote-option6-pool-set	801
24.138remote-option6-subnet-del	803
24.139remote-option6-subnet-set	804
24.140remote-server4-del	805
24.141remote-server4-get	806
24.142remote-server4-get-all	806
24.143remote-server4-set	807
24.144remote-server6-del	808
24.145remote-server6-get	809
24.146remote-server6-get-all	810
24.147remote-server6-set	811
24.148remote-subnet4-del-by-id	812
24.149remote-subnet4-del-by-prefix	813
24.150remote-subnet4-get-by-id	814
24.151remote-subnet4-get-by-prefix	815
24.152remote-subnet4-list	816
24.153remote-subnet4-set	817
24.154remote-subnet6-del-by-id	818
24.155remote-subnet6-del-by-prefix	819
24.156remote-subnet6-get-by-id	820
24.157remote-subnet6-get-by-prefix	821
24.158remote-subnet6-list	822
24.159remote-subnet6-set	823
24.160reservation-add	824
24.161reservation-del	825
24.162reservation-get	826
24.163reservation-get-all	827
24.164reservation-get-by-address	828
24.165reservation-get-by-hostname	828
24.166reservation-get-by-id	829
24.167reservation-get-page	830
24.168reservation-update	830
24.169server-tag-get	831
24.170shutdown	832
24.171stat-lease4-get	833
24.172stat-lease6-get	834
24.173statistic-get	835
24.174statistic-get-all	835
24.175statistic-remove	848
24.176statistic-remove-all	849
24.177statistic-reset	850
24.178statistic-reset-all	851
24.179statistic-sample-age-set	851
24.180statistic-sample-age-set-all	852
24.181statistic-sample-count-set	853
24.182statistic-sample-count-set-all	854
24.183status-get	854
24.184subnet4-add	856

24.185	subnet4-del	857
24.186	subnet4-delta-add	858
24.187	subnet4-delta-del	859
24.188	subnet4-get	860
24.189	subnet4-list	861
24.190	subnet4-select-test	861
24.191	subnet4-update	862
24.192	subnet4o6-select-test	863
24.193	subnet6-add	864
24.194	subnet6-del	865
24.195	subnet6-delta-add	865
24.196	subnet6-delta-del	866
24.197	subnet6-get	867
24.198	subnet6-list	868
24.199	subnet6-select-test	869
24.200	subnet6-update	870
24.201	version-get	871
25	Manual Pages	873
25.1	kea-dhcp4 - DHCPv4 server in Kea	873
25.1.1	Synopsis	873
25.1.2	Description	873
25.1.3	Arguments	873
25.1.4	Documentation	874
25.1.5	Mailing Lists and Support	874
25.1.6	History	874
25.1.7	See Also	874
25.2	kea-dhcp6 - DHCPv6 server in Kea	874
25.2.1	Synopsis	874
25.2.2	Description	874
25.2.3	Arguments	874
25.2.4	Documentation	875
25.2.5	Mailing Lists and Support	875
25.2.6	History	875
25.2.7	See Also	875
25.3	kea-ctrl-agent - Control Agent process in Kea	876
25.3.1	Synopsis	876
25.3.2	Description	876
25.3.3	Arguments	876
25.3.4	Documentation	876
25.3.5	Mailing Lists and Support	876
25.3.6	History	877
25.3.7	See Also	877
25.4	keactrl - Shell script for managing Kea	877
25.4.1	Synopsis	877
25.4.2	Description	877
25.4.3	Configuration File	877
25.4.4	Options	877
25.4.5	Documentation	878
25.4.6	Mailing Lists and Support	878
25.4.7	See Also	878
25.5	kea-admin - Shell script for managing Kea databases	878
25.5.1	Synopsis	878
25.5.2	Description	879

25.5.3	Arguments	879
25.5.4	Documentation	880
25.5.5	Mailing Lists and Support	880
25.5.6	See Also	880
25.6	kea-dhcp-ddns - DHCP-DDNS process in Kea	880
25.6.1	Synopsis	880
25.6.2	Description	881
25.6.3	Arguments	881
25.6.4	Documentation	881
25.6.5	Mailing Lists and Support	881
25.6.6	History	881
25.6.7	See Also	882
25.7	kea-lfc - Lease File Cleanup process in Kea	882
25.7.1	Synopsis	882
25.7.2	Description	882
25.7.3	Arguments	882
25.7.4	Documentation	883
25.7.5	Mailing Lists and Support	883
25.7.6	History	883
25.7.7	See Also	883
25.8	kea-shell - Text client for kea-dhcp4, kea-dhcp6 and kea-dhcp-ddns servers	883
25.8.1	Synopsis	883
25.8.2	Description	883
25.8.3	Arguments	884
25.8.4	Documentation	885
25.8.5	Mailing Lists and Support	885
25.8.6	History	885
25.8.7	See Also	885
25.9	kea-netconf - NETCONF agent for configuring Kea	885
25.9.1	Synopsis	885
25.9.2	Description	885
25.9.3	Arguments	885
25.9.4	Documentation	886
25.9.5	Mailing Lists and Support	886
25.9.6	History	886
25.9.7	See Also	886
25.10	perfdhcp - DHCP benchmarking tool	886
25.10.1	Synopsis	886
25.10.2	Description	886
25.10.3	Templates	887
25.10.4	Options	888
25.10.5	DHCPv4-Only Options	890
25.10.6	DHCPv6-Only Options	891
25.10.7	Template-Related Options	891
25.10.8	Options Controlling a Test	891
25.10.9	Arguments	892
25.10.10	Errors	892
25.10.11	Exit Status	892
25.10.12	Usage Examples	892
25.10.13	Documentation	893
25.10.14	Mailing Lists and Support	893
25.10.15	History	893
25.10.16	See Also	893

26 Kea Messages Manual	895
26.1 ALLOC	895
26.1.1 ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS	895
26.1.2 ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS6	895
26.1.3 ALLOC_ENGINE_LEASE_RECLAIMED	895
26.1.4 ALLOC_ENGINE_V4_ALLOC_ERROR	895
26.1.5 ALLOC_ENGINE_V4_ALLOC_FAIL	896
26.1.6 ALLOC_ENGINE_V4_ALLOC_FAIL_CLASSES	896
26.1.7 ALLOC_ENGINE_V4_ALLOC_FAIL_NO_POOLS	896
26.1.8 ALLOC_ENGINE_V4_ALLOC_FAIL_SHARED_NETWORK	896
26.1.9 ALLOC_ENGINE_V4_ALLOC_FAIL_SUBNET	897
26.1.10 ALLOC_ENGINE_V4_DECLINED_RECOVERED	897
26.1.11 ALLOC_ENGINE_V4_DISCOVER_ADDRESS_CONFLICT	897
26.1.12 ALLOC_ENGINE_V4_DISCOVER_HR	897
26.1.13 ALLOC_ENGINE_V4_LEASES_RECLAMATION_COMPLETE	897
26.1.14 ALLOC_ENGINE_V4_LEASES_RECLAMATION_FAILED	897
26.1.15 ALLOC_ENGINE_V4_LEASES_RECLAMATION_SLOW	898
26.1.16 ALLOC_ENGINE_V4_LEASES_RECLAMATION_START	898
26.1.17 ALLOC_ENGINE_V4_LEASES_RECLAMATION_TIMEOUT	898
26.1.18 ALLOC_ENGINE_V4_LEASE_RECLAIM	898
26.1.19 ALLOC_ENGINE_V4_LEASE_RECLAMATION_FAILED	898
26.1.20 ALLOC_ENGINE_V4_NO_MORE_EXPIRED_LEASES	898
26.1.21 ALLOC_ENGINE_V4_OFFER_EXISTING_LEASE	899
26.1.22 ALLOC_ENGINE_V4_OFFER_NEW_LEASE	899
26.1.23 ALLOC_ENGINE_V4_OFFER_REQUESTED_LEASE	899
26.1.24 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE	899
26.1.25 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE_COMPLETE	899
26.1.26 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE_FAILED	899
26.1.27 ALLOC_ENGINE_V4_REQUEST_ADDRESS_RESERVED	899
26.1.28 ALLOC_ENGINE_V4_REQUEST_ALLOC_REQUESTED	900
26.1.29 ALLOC_ENGINE_V4_REQUEST_EXTEND_LEASE	900
26.1.30 ALLOC_ENGINE_V4_REQUEST_INVALID	900
26.1.31 ALLOC_ENGINE_V4_REQUEST_IN_USE	900
26.1.32 ALLOC_ENGINE_V4_REQUEST_OUT_OF_POOL	900
26.1.33 ALLOC_ENGINE_V4_REQUEST_PICK_ADDRESS	900
26.1.34 ALLOC_ENGINE_V4_REQUEST_REMOVE_LEASE	900
26.1.35 ALLOC_ENGINE_V4_REQUEST_USE_HR	901
26.1.36 ALLOC_ENGINE_V4_REUSE_EXPIRED_LEASE_DATA	901
26.1.37 ALLOC_ENGINE_V6_ALLOC_ERROR	901
26.1.38 ALLOC_ENGINE_V6_ALLOC_FAIL	901
26.1.39 ALLOC_ENGINE_V6_ALLOC_FAIL_CLASSES	901
26.1.40 ALLOC_ENGINE_V6_ALLOC_FAIL_NO_POOLS	901
26.1.41 ALLOC_ENGINE_V6_ALLOC_FAIL_SHARED_NETWORK	902
26.1.42 ALLOC_ENGINE_V6_ALLOC_FAIL_SUBNET	902
26.1.43 ALLOC_ENGINE_V6_ALLOC_HR_LEASE_EXISTS	902
26.1.44 ALLOC_ENGINE_V6_ALLOC_LEASES_HR	902
26.1.45 ALLOC_ENGINE_V6_ALLOC_LEASES_NO_HR	903
26.1.46 ALLOC_ENGINE_V6_ALLOC_NO_LEASES_HR	903
26.1.47 ALLOC_ENGINE_V6_ALLOC_NO_V6_HR	903
26.1.48 ALLOC_ENGINE_V6_ALLOC_UNRESERVED	903
26.1.49 ALLOC_ENGINE_V6_CALCULATED_PREFERRED_LIFETIME	903
26.1.50 ALLOC_ENGINE_V6_DECLINED_RECOVERED	903
26.1.51 ALLOC_ENGINE_V6_EXPIRED_HINT_RESERVED	904
26.1.52 ALLOC_ENGINE_V6_EXTEND_ALLOC_UNRESERVED	904

26.1.53	ALLOC_ENGINE_V6_EXTEND_ERROR	904
26.1.54	ALLOC_ENGINE_V6_EXTEND_LEASE	904
26.1.55	ALLOC_ENGINE_V6_EXTEND_LEASE_DATA	904
26.1.56	ALLOC_ENGINE_V6_EXTEND_NEW_LEASE_DATA	904
26.1.57	ALLOC_ENGINE_V6_HINT_RESERVED	904
26.1.58	ALLOC_ENGINE_V6_HR_ADDR_GRANTED	905
26.1.59	ALLOC_ENGINE_V6_HR_PREFIX_GRANTED	905
26.1.60	ALLOC_ENGINE_V6_LEASES_RECLAMATION_COMPLETE	905
26.1.61	ALLOC_ENGINE_V6_LEASES_RECLAMATION_FAILED	905
26.1.62	ALLOC_ENGINE_V6_LEASES_RECLAMATION_SLOW	905
26.1.63	ALLOC_ENGINE_V6_LEASES_RECLAMATION_START	905
26.1.64	ALLOC_ENGINE_V6_LEASES_RECLAMATION_TIMEOUT	906
26.1.65	ALLOC_ENGINE_V6_LEASE_RECLAIM	906
26.1.66	ALLOC_ENGINE_V6_LEASE_RECLAMATION_FAILED	906
26.1.67	ALLOC_ENGINE_V6_NO_MORE_EXPIRED_LEASES	906
26.1.68	ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE	906
26.1.69	ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE_COMPLETE	906
26.1.70	ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE_FAILED	907
26.1.71	ALLOC_ENGINE_V6_RENEW_HR	907
26.1.72	ALLOC_ENGINE_V6_RENEW_REMOVE_RESERVED	907
26.1.73	ALLOC_ENGINE_V6_REUSE_EXPIRED_LEASE_DATA	907
26.1.74	ALLOC_ENGINE_V6_REVOKED_ADDR_LEASE	907
26.1.75	ALLOC_ENGINE_V6_REVOKED_PREFIX_LEASE	907
26.1.76	ALLOC_ENGINE_V6_REVOKED_SHARED_ADDR_LEASE	908
26.1.77	ALLOC_ENGINE_V6_REVOKED_SHARED_PREFIX_LEASE	908
26.2	ASIODNS	908
26.2.1	ASIODNS_FETCH_COMPLETED	908
26.2.2	ASIODNS_FETCH_STOPPED	908
26.2.3	ASIODNS_OPEN_SOCKET	908
26.2.4	ASIODNS_READ_DATA	908
26.2.5	ASIODNS_READ_TIMEOUT	909
26.2.6	ASIODNS_SEND_DATA	909
26.2.7	ASIODNS_UNKNOWN_ORIGIN	909
26.2.8	ASIODNS_UNKNOWN_RESULT	909
26.3	BAD	909
26.3.1	BAD_CLIENT_CREDENTIALS	909
26.4	BOOTP	909
26.4.1	BOOTP_BOOTP_QUERY	909
26.4.2	BOOTP_LOAD	910
26.4.3	BOOTP_PACKET_OPTIONS_SKIPPED	910
26.4.4	BOOTP_PACKET_PACK	910
26.4.5	BOOTP_PACKET_PACK_FAIL	910
26.4.6	BOOTP_PACKET_UNPACK_FAILED	910
26.4.7	BOOTP_UNLOAD	910
26.5	BULK	910
26.5.1	BULK_LEASE_QUERY4_UNSUPPORTED_MSG_TYPE	910
26.5.2	BULK_LEASE_QUERY6_UNSUPPORTED_MSG_TYPE	911
26.5.3	BULK_LEASE_QUERY_AT_MAX_CONCURRENT_QUERIES	911
26.5.4	BULK_LEASE_QUERY_DEQUEUED	911
26.5.5	BULK_LEASE_QUERY_DUPLICATE_XID	911
26.5.6	BULK_LEASE_QUERY_EMPTY_REQUEST	911
26.5.7	BULK_LEASE_QUERY_INVALID_REQUEST	911
26.5.8	BULK_LEASE_QUERY_LISTENER_START_FAILED	911
26.5.9	BULK_LEASE_QUERY_PAUSE_CHECK_PERMISSIONS_FAILED	912

26.5.10	BULK_LEASE_QUERY_PAUSE_LISTENER_FAILED	912
26.5.11	BULK_LEASE_QUERY_PAUSE_LISTENER_ILLEGAL	912
26.5.12	BULK_LEASE_QUERY_PROCESSING_UNEXPECTED_FAILURE	912
26.5.13	BULK_LEASE_QUERY_QUERY_RECEIVED	912
26.5.14	BULK_LEASE_QUERY_REJECTED_CONNECTION	912
26.5.15	BULK_LEASE_QUERY_RESPONSE_SEND_ERROR	912
26.5.16	BULK_LEASE_QUERY_RESPONSE_SENT	913
26.5.17	BULK_LEASE_QUERY_RESUME_LISTENER_FAILED	913
26.5.18	BULK_LEASE_QUERY_UNPACK_ERROR	913
26.6	CB	913
26.6.1	CB_CMDS_CLASS4_DEL_HANDLER_FAILED	913
26.6.2	CB_CMDS_CLASS4_GET_ALL_HANDLER_FAILED	913
26.6.3	CB_CMDS_CLASS4_GET_HANDLER_FAILED	913
26.6.4	CB_CMDS_CLASS4_SET_HANDLER_FAILED	913
26.6.5	CB_CMDS_CLASS6_DEL_HANDLER_FAILED	914
26.6.6	CB_CMDS_CLASS6_GET_ALL_HANDLER_FAILED	914
26.6.7	CB_CMDS_CLASS6_GET_HANDLER_FAILED	914
26.6.8	CB_CMDS_CLASS6_SET_HANDLER_FAILED	914
26.6.9	CB_CMDS_DEINIT_OK	914
26.6.10	CB_CMDS_GLOBAL_PARAMETER4_DEL_HANDLER_FAILED	914
26.6.11	CB_CMDS_GLOBAL_PARAMETER4_GET_ALL_HANDLER_FAILED	914
26.6.12	CB_CMDS_GLOBAL_PARAMETER4_GET_HANDLER_FAILED	914
26.6.13	CB_CMDS_GLOBAL_PARAMETER4_SET_HANDLER_FAILED	915
26.6.14	CB_CMDS_GLOBAL_PARAMETER6_DEL_HANDLER_FAILED	915
26.6.15	CB_CMDS_GLOBAL_PARAMETER6_GET_ALL_HANDLER_FAILED	915
26.6.16	CB_CMDS_GLOBAL_PARAMETER6_GET_HANDLER_FAILED	915
26.6.17	CB_CMDS_GLOBAL_PARAMETER6_SET_HANDLER_FAILED	915
26.6.18	CB_CMDS_INIT_OK	915
26.6.19	CB_CMDS_NETWORK4_DEL_HANDLER_FAILED	915
26.6.20	CB_CMDS_NETWORK4_GET_HANDLER_FAILED	915
26.6.21	CB_CMDS_NETWORK4_LIST_HANDLER_FAILED	916
26.6.22	CB_CMDS_NETWORK4_SET_HANDLER_FAILED	916
26.6.23	CB_CMDS_NETWORK6_DEL_HANDLER_FAILED	916
26.6.24	CB_CMDS_NETWORK6_GET_HANDLER_FAILED	916
26.6.25	CB_CMDS_NETWORK6_LIST_HANDLER_FAILED	916
26.6.26	CB_CMDS_NETWORK6_SET_HANDLER_FAILED	916
26.6.27	CB_CMDS_OPTION4_GLOBAL_DEL_HANDLER_FAILED	916
26.6.28	CB_CMDS_OPTION4_GLOBAL_GET_ALL_HANDLER_FAILED	917
26.6.29	CB_CMDS_OPTION4_GLOBAL_GET_HANDLER_FAILED	917
26.6.30	CB_CMDS_OPTION4_GLOBAL_SET_HANDLER_FAILED	917
26.6.31	CB_CMDS_OPTION4_NETWORK_DEL_HANDLER_FAILED	917
26.6.32	CB_CMDS_OPTION4_NETWORK_SET_HANDLER_FAILED	917
26.6.33	CB_CMDS_OPTION4_POOL_DEL_HANDLER_FAILED	917
26.6.34	CB_CMDS_OPTION4_POOL_SET_HANDLER_FAILED	917
26.6.35	CB_CMDS_OPTION4_SUBNET_DEL_HANDLER_FAILED	918
26.6.36	CB_CMDS_OPTION4_SUBNET_SET_HANDLER_FAILED	918
26.6.37	CB_CMDS_OPTION6_GLOBAL_DEL_HANDLER_FAILED	918
26.6.38	CB_CMDS_OPTION6_GLOBAL_GET_ALL_HANDLER_FAILED	918
26.6.39	CB_CMDS_OPTION6_GLOBAL_GET_HANDLER_FAILED	918
26.6.40	CB_CMDS_OPTION6_GLOBAL_SET_HANDLER_FAILED	918
26.6.41	CB_CMDS_OPTION6_NETWORK_DEL_HANDLER_FAILED	918
26.6.42	CB_CMDS_OPTION6_NETWORK_SET_HANDLER_FAILED	919
26.6.43	CB_CMDS_OPTION6_PD_POOL_DEL_HANDLER_FAILED	919
26.6.44	CB_CMDS_OPTION6_PD_POOL_SET_HANDLER_FAILED	919

26.6.45	CB_CMDS_OPTION6_POOL_DEL_HANDLER_FAILED	919
26.6.46	CB_CMDS_OPTION6_POOL_SET_HANDLER_FAILED	919
26.6.47	CB_CMDS_OPTION6_SUBNET_DEL_HANDLER_FAILED	919
26.6.48	CB_CMDS_OPTION6_SUBNET_SET_HANDLER_FAILED	919
26.6.49	CB_CMDS_OPTION_DEF4_DEL_HANDLER_FAILED	920
26.6.50	CB_CMDS_OPTION_DEF4_GET_ALL_HANDLER_FAILED	920
26.6.51	CB_CMDS_OPTION_DEF4_GET_HANDLER_FAILED	920
26.6.52	CB_CMDS_OPTION_DEF4_SET_HANDLER_FAILED	920
26.6.53	CB_CMDS_OPTION_DEF6_DEL_HANDLER_FAILED	920
26.6.54	CB_CMDS_OPTION_DEF6_GET_ALL_HANDLER_FAILED	920
26.6.55	CB_CMDS_OPTION_DEF6_GET_HANDLER_FAILED	920
26.6.56	CB_CMDS_OPTION_DEF6_SET_HANDLER_FAILED	921
26.6.57	CB_CMDS_SERVER4_DEL_HANDLER_FAILED	921
26.6.58	CB_CMDS_SERVER4_GET_ALL_HANDLER_FAILED	921
26.6.59	CB_CMDS_SERVER4_GET_HANDLER_FAILED	921
26.6.60	CB_CMDS_SERVER4_SET_HANDLER_FAILED	921
26.6.61	CB_CMDS_SERVER6_DEL_HANDLER_FAILED	921
26.6.62	CB_CMDS_SERVER6_GET_ALL_HANDLER_FAILED	921
26.6.63	CB_CMDS_SERVER6_GET_HANDLER_FAILED	922
26.6.64	CB_CMDS_SERVER6_SET_HANDLER_FAILED	922
26.6.65	CB_CMDS_SUBNET4_DEL_BY_ID_HANDLER_FAILED	922
26.6.66	CB_CMDS_SUBNET4_DEL_BY_PREFIX_HANDLER_FAILED	922
26.6.67	CB_CMDS_SUBNET4_GET_BY_ID_HANDLER_FAILED	922
26.6.68	CB_CMDS_SUBNET4_GET_BY_PREFIX_HANDLER_FAILED	922
26.6.69	CB_CMDS_SUBNET4_LIST_HANDLER_FAILED	922
26.6.70	CB_CMDS_SUBNET4_SET_HANDLER_FAILED	923
26.6.71	CB_CMDS_SUBNET6_DEL_BY_ID_HANDLER_FAILED	923
26.6.72	CB_CMDS_SUBNET6_DEL_BY_PREFIX_HANDLER_FAILED	923
26.6.73	CB_CMDS_SUBNET6_GET_BY_ID_HANDLER_FAILED	923
26.6.74	CB_CMDS_SUBNET6_GET_BY_PREFIX_HANDLER_FAILED	923
26.6.75	CB_CMDS_SUBNET6_LIST_HANDLER_FAILED	923
26.6.76	CB_CMDS_SUBNET6_SET_HANDLER_FAILED	923
26.7	CLASS	924
26.7.1	CLASS_CMDS_CLASS_ADD	924
26.7.2	CLASS_CMDS_CLASS_ADD_FAILED	924
26.7.3	CLASS_CMDS_CLASS_ADD_HANDLER_FAILED	924
26.7.4	CLASS_CMDS_CLASS_DEL	924
26.7.5	CLASS_CMDS_CLASS_DEL_EMPTY	924
26.7.6	CLASS_CMDS_CLASS_DEL_FAILED	924
26.7.7	CLASS_CMDS_CLASS_DEL_HANDLER_FAILED	924
26.7.8	CLASS_CMDS_CLASS_GET	924
26.7.9	CLASS_CMDS_CLASS_GET_EMPTY	925
26.7.10	CLASS_CMDS_CLASS_GET_FAILED	925
26.7.11	CLASS_CMDS_CLASS_GET_HANDLER_FAILED	925
26.7.12	CLASS_CMDS_CLASS_LIST	925
26.7.13	CLASS_CMDS_CLASS_LIST_EMPTY	925
26.7.14	CLASS_CMDS_CLASS_LIST_FAILED	925
26.7.15	CLASS_CMDS_CLASS_LIST_HANDLER_FAILED	925
26.7.16	CLASS_CMDS_CLASS_UPDATE	925
26.7.17	CLASS_CMDS_CLASS_UPDATE_EMPTY	926
26.7.18	CLASS_CMDS_CLASS_UPDATE_FAILED	926
26.7.19	CLASS_CMDS_CLASS_UPDATE_HANDLER_FAILED	926
26.7.20	CLASS_CMDS_DEINIT_OK	926
26.7.21	CLASS_CMDS_INIT_FAILED	926

26.7.22	CLASS_CMDS_INIT_OK	926
26.8	COMMAND	926
26.8.1	COMMAND_ACCEPTOR_START	926
26.8.2	COMMAND_DEREGISTERED	926
26.8.3	COMMAND_EXTENDED_REGISTERED	927
26.8.4	COMMAND_HTTP_LISTENER_COMMAND_REJECTED	927
26.8.5	COMMAND_HTTP_LISTENER_STARTED	927
26.8.6	COMMAND_HTTP_LISTENER_STOPPED	927
26.8.7	COMMAND_HTTP_LISTENER_STOPPING	927
26.8.8	COMMAND_PROCESS_ERROR1	927
26.8.9	COMMAND_PROCESS_ERROR2	928
26.8.10	COMMAND_RECEIVED	928
26.8.11	COMMAND_REGISTERED	928
26.8.12	COMMAND_RESPONSE_ERROR	928
26.8.13	COMMAND_SOCKET_ACCEPT_FAIL	928
26.8.14	COMMAND_SOCKET_CLOSED_BY_FOREIGN_HOST	928
26.8.15	COMMAND_SOCKET_CONNECTION_CANCEL_FAIL	928
26.8.16	COMMAND_SOCKET_CONNECTION_CLOSED	929
26.8.17	COMMAND_SOCKET_CONNECTION_CLOSE_FAIL	929
26.8.18	COMMAND_SOCKET_CONNECTION_OPENED	929
26.8.19	COMMAND_SOCKET_CONNECTION_SHUTDOWN_FAIL	929
26.8.20	COMMAND_SOCKET_CONNECTION_TIMEOUT	929
26.8.21	COMMAND_SOCKET_READ	929
26.8.22	COMMAND_SOCKET_READ_FAIL	929
26.8.23	COMMAND_SOCKET_WRITE	929
26.8.24	COMMAND_SOCKET_WRITE_FAIL	930
26.8.25	COMMAND_WATCH_SOCKET_CLEAR_ERROR	930
26.8.26	COMMAND_WATCH_SOCKET_CLOSE_ERROR	930
26.8.27	COMMAND_WATCH_SOCKET_MARK_READY_ERROR	930
26.9	CONFIG	930
26.9.1	CONFIG_BACKENDS_REGISTERED	930
26.10	CTRL	930
26.10.1	CTRL_AGENT_COMMAND_FORWARDED	930
26.10.2	CTRL_AGENT_COMMAND_FORWARD_BEGIN	930
26.10.3	CTRL_AGENT_COMMAND_FORWARD_FAILED	931
26.10.4	CTRL_AGENT_COMMAND_RECEIVED	931
26.10.5	CTRL_AGENT_CONFIG_CHECK_FAIL	931
26.10.6	CTRL_AGENT_CONFIG_FAIL	931
26.10.7	CTRL_AGENT_CONFIG_SYNTAX_WARNING	931
26.10.8	CTRL_AGENT_FAILED	931
26.10.9	CTRL_AGENT_HTTPS_SERVICE_REUSE_FAILED	931
26.10.10	CTRL_AGENT_HTTPS_SERVICE_STARTED	932
26.10.11	CTRL_AGENT_HTTPS_SERVICE_UPDATED	932
26.10.12	CTRL_AGENT_HTTP_SERVICE_REUSE_FAILED	932
26.10.13	CTRL_AGENT_HTTP_SERVICE_STARTED	932
26.10.14	CTRL_AGENT_HTTP_SERVICE_UPDATED	932
26.10.15	CTRL_AGENT_IS_DEPRECATED	932
26.10.16	CTRL_AGENT_RUN_EXIT	932
26.10.17	CTRL_AGENT_SECURITY_CHECKS_DISABLED	932
26.10.18	CTRL_AGENT_STARTED	933
26.11	DATABASE	933
26.11.1	DATABASE_INVALID_ACCESS	933
26.11.2	DATABASE_MYSQL_COMMIT	933
26.11.3	DATABASE_MYSQL_FATAL_ERROR	933

26.11.4	DATABASE_MYSQL_INITIALIZE_SCHEMA	933
26.11.5	DATABASE_MYSQL_ROLLBACK	933
26.11.6	DATABASE_MYSQL_START_TRANSACTION	934
26.11.7	DATABASE_PGSQL_COMMIT	934
26.11.8	DATABASE_PGSQL_CREATE_SAVEPOINT	934
26.11.9	DATABASE_PGSQL_DEALLOC_ERROR	934
26.11.10	DATABASE_PGSQL_FATAL_ERROR	934
26.11.11	DATABASE_PGSQL_INITIALIZE_SCHEMA	934
26.11.12	DATABASE_PGSQL_ROLLBACK	935
26.11.13	DATABASE_PGSQL_ROLLBACK_SAVEPOINT	935
26.11.14	DATABASE_PGSQL_START_TRANSACTION	935
26.11.15	DATABASE_PGSQL_TCP_USER_TIMEOUT_UNSUPPORTED	935
26.11.16	DATABASE_TO_JSON_BOOLEAN_ERROR	935
26.11.17	DATABASE_TO_JSON_INTEGER_ERROR	935
26.11.18	DATABASE_TO_JSON_UNKNOWN_TYPE_ERROR	935
26.12	DCTL	936
26.12.1	DCTL_ALREADY_RUNNING	936
26.12.2	DCTL_CFG_FILE_RELOAD_ERROR	936
26.12.3	DCTL_CFG_FILE_RELOAD_SIGNAL_RECVD	936
26.12.4	DCTL_CONFIG_CHECK_COMPLETE	936
26.12.5	DCTL_CONFIG_COMPLETE	936
26.12.6	DCTL_CONFIG_DEPRECATED	936
26.12.7	DCTL_CONFIG_FETCH	937
26.12.8	DCTL_CONFIG_FILE_LOAD_FAIL	937
26.12.9	DCTL_CONFIG_START	937
26.12.10	DCTL_DB_OPEN_CONNECTION_WITH_RETRY_FAILED	937
26.12.11	DCTL_DEPRECATED_OUTPUT_OPTIONS	937
26.12.12	DCTL_DEVELOPMENT_VERSION	937
26.12.13	DCTL_INIT_PROCESS	937
26.12.14	DCTL_INIT_PROCESS_FAIL	938
26.12.15	DCTL_NOT_RUNNING	938
26.12.16	DCTL_OPEN_CONFIG_DB	938
26.12.17	DCTL_PARSER_FAIL	938
26.12.18	DCTL_PID_FILE_ERROR	938
26.12.19	DCTL_PROCESS_FAILED	938
26.12.20	DCTL_RUN_PROCESS	938
26.12.21	DCTL_SHUTDOWN	939
26.12.22	DCTL_SHUTDOWN_SIGNAL_RECVD	939
26.12.23	DCTL_STANDALONE	939
26.12.24	DCTL_STARTING	939
26.12.25	DCTL_UNLOAD_LIBRARIES_ERROR	939
26.12.26	DCTL_UNSUPPORTED_SIGNAL	939
26.13	DDNS	939
26.13.1	DDNS_TUNING4_CALCULATED_HOSTNAME	939
26.13.2	DDNS_TUNING4_PROCESS_ERROR	940
26.13.3	DDNS_TUNING4_SKIPPING_DDNS	940
26.13.4	DDNS_TUNING6_CALCULATED_HOSTNAME	940
26.13.5	DDNS_TUNING6_PROCESS_ERROR	940
26.13.6	DDNS_TUNING6_SKIPPING_DDNS	940
26.13.7	DDNS_TUNING_GLOBAL_EXPR_SET	940
26.13.8	DDNS_TUNING_LOAD_ERROR	940
26.13.9	DDNS_TUNING_LOAD_OK	941
26.13.10	DDNS_TUNING_SUBNET_EXPRESSION_PARSE	941
26.13.11	DDNS_TUNING_SUBNET_EXPRESSION_PARSE_ERROR	941

26.13.12DDNS_TUNING_SUBNET_EXPR_CACHED	941
26.13.13DDNS_TUNING_UNLOAD	941
26.14 DHCP4	941
26.14.1 DHCP4_ADDITIONAL_CLASS_EVAL_ERROR	941
26.14.2 DHCP4_ADDITIONAL_CLASS_EVAL_RESULT	941
26.14.3 DHCP4_ADDITIONAL_CLASS_NO_TEST	942
26.14.4 DHCP4_ADDITIONAL_CLASS_UNDEFINED	942
26.14.5 DHCP4_ALREADY_RUNNING	942
26.14.6 DHCP4_BUFFER_RECEIVED	942
26.14.7 DHCP4_BUFFER_RECEIVE_FAIL	942
26.14.8 DHCP4_BUFFER_UNPACK	942
26.14.9 DHCP4_BUFFER_WAIT_SIGNAL	942
26.14.10DHCP4_CB_ON_DEMAND_FETCH_UPDATES_FAIL	943
26.14.11DHCP4_CB_PERIODIC_FETCH_UPDATES_FAIL	943
26.14.12DHCP4_CB_PERIODIC_FETCH_UPDATES_RETRIES_EXHAUSTED	943
26.14.13DHCP4_CLASSES_ASSIGNED	943
26.14.14DHCP4_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION	943
26.14.15DHCP4_CLASS_ASSIGNED	943
26.14.16DHCP4_CLASS_UNCONFIGURED	944
26.14.17DHCP4_CLIENTID_IGNORED_FOR_LEASES	944
26.14.18DHCP4_CLIENT_FQDN_DATA	944
26.14.19DHCP4_CLIENT_FQDN_PROCESS	944
26.14.20DHCP4_CLIENT_FQDN_SCRUBBED_EMPTY	944
26.14.21DHCP4_CLIENT_HOSTNAME_DATA	944
26.14.22DHCP4_CLIENT_HOSTNAME_MALFORMED	945
26.14.23DHCP4_CLIENT_HOSTNAME_PROCESS	945
26.14.24DHCP4_CLIENT_HOSTNAME_SCRUBBED_EMPTY	945
26.14.25DHCP4_CLIENT_NAME_PROC_FAIL	945
26.14.26DHCP4_CONFIG_COMPLETE	945
26.14.27DHCP4_CONFIG_LOAD_FAIL	945
26.14.28DHCP4_CONFIG_PACKET_QUEUE	945
26.14.29DHCP4_CONFIG_RECEIVED	946
26.14.30DHCP4_CONFIG_START	946
26.14.31DHCP4_CONFIG_SYNTAX_WARNING	946
26.14.32DHCP4_CONFIG_UNRECOVERABLE_ERROR	946
26.14.33DHCP4_CONFIG_UNSUPPORTED_OBJECT	946
26.14.34DHCP4_DB_RECONNECT_DISABLED	946
26.14.35DHCP4_DB_RECONNECT_FAILED	946
26.14.36DHCP4_DB_RECONNECT_LOST_CONNECTION	947
26.14.37DHCP4_DB_RECONNECT_NO_DB_CTL	947
26.14.38DHCP4_DB_RECONNECT_SUCCEEDED	947
26.14.39DHCP4_DDNS_REQUEST_SEND_FAILED	947
26.14.40DHCP4_DECLINE_FAIL	947
26.14.41DHCP4_DECLINE_LEASE	947
26.14.42DHCP4_DECLINE_LEASE_MISMATCH	947
26.14.43DHCP4_DECLINE_LEASE_NOT_FOUND	948
26.14.44DHCP4_DEFERRED_OPTION_MISSING	948
26.14.45DHCP4_DEFERRED_OPTION_UNPACK_FAIL	948
26.14.46DHCP4_DEVELOPMENT_VERSION	948
26.14.47DHCP4_DHCP4O6_BAD_PACKET	948
26.14.48DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_DROP	948
26.14.49DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_SKIP	948
26.14.50DHCP4_DHCP4O6_PACKET_RECEIVED	949
26.14.51DHCP4_DHCP4O6_PACKET_SEND	949

26.14.52DHCP4_DHCP4O6_PACKET_SEND_FAIL	949
26.14.53DHCP4_DHCP4O6_RECEIVE_FAIL	949
26.14.54DHCP4_DHCP4O6_RECEIVING	949
26.14.55DHCP4_DHCP4O6_RESPONSE_DATA	949
26.14.56DHCP4_DHCP4O6_SUBNET_DATA	949
26.14.57DHCP4_DHCP4O6_SUBNET_SELECTED	950
26.14.58DHCP4_DHCP4O6_SUBNET_SELECTION_FAILED	950
26.14.59DHCP4_DISCOVER	950
26.14.60DHCP4_DYNAMIC_RECONFIGURATION	950
26.14.61DHCP4_DYNAMIC_RECONFIGURATION_FAIL	950
26.14.62DHCP4_DYNAMIC_RECONFIGURATION_SUCCESS	950
26.14.63DHCP4_EMPTY_HOSTNAME	950
26.14.64DHCP4_FLEX_ID	951
26.14.65DHCP4_GENERATE_FQDN	951
26.14.66DHCP4_HOOK_BUFFER_RCVD_DROP	951
26.14.67DHCP4_HOOK_BUFFER_RCVD_SKIP	951
26.14.68DHCP4_HOOK_BUFFER_SEND_SKIP	951
26.14.69DHCP4_HOOK_DDNS_UPDATE	951
26.14.70DHCP4_HOOK_DECLINE_SKIP	952
26.14.71DHCP4_HOOK_LEASE4_OFFER_ARGUMENT_MISSING	952
26.14.72DHCP4_HOOK_LEASE4_OFFER_DROP	952
26.14.73DHCP4_HOOK_LEASE4_OFFER_PARK	952
26.14.74DHCP4_HOOK_LEASE4_OFFER_PARKING_LOT_FULL	952
26.14.75DHCP4_HOOK_LEASE4_RELEASE_SKIP	952
26.14.76DHCP4_HOOK_LEASES4_COMMITTED_DROP	952
26.14.77DHCP4_HOOK_LEASES4_COMMITTED_PARK	953
26.14.78DHCP4_HOOK_LEASES4_COMMITTED_PARKING_LOT_FULL	953
26.14.79DHCP4_HOOK_PACKET_RCVD_SKIP	953
26.14.80DHCP4_HOOK_PACKET_SEND_DROP	953
26.14.81DHCP4_HOOK_PACKET_SEND_SKIP	953
26.14.82DHCP4_HOOK_SUBNET4_SELECT_4O6_PARKING_LOT_FULL	953
26.14.83DHCP4_HOOK_SUBNET4_SELECT_DROP	954
26.14.84DHCP4_HOOK_SUBNET4_SELECT_PARK	954
26.14.85DHCP4_HOOK_SUBNET4_SELECT_PARKING_LOT_FULL	954
26.14.86DHCP4_HOOK_SUBNET4_SELECT_SKIP	954
26.14.87DHCP4_HOOK_SUBNET6_SELECT_PARKING_LOT_FULL	954
26.14.88DHCP4_INFORM_DIRECT_REPLY	954
26.14.89DHCP4_INIT_FAIL	955
26.14.90DHCP4_INIT_REBOOT	955
26.14.91DHCP4_LEASE_ALLOC	955
26.14.92DHCP4_LEASE_OFFER	955
26.14.93DHCP4_LEASE_QUERY_PACKET_PACK_FAILED	955
26.14.94DHCP4_LEASE_QUERY_PACKET_UNPACK_FAILED	955
26.14.95DHCP4_LEASE_QUERY_PROCESS_FAILED	955
26.14.96DHCP4_LEASE_QUERY_RECEIVED	956
26.14.97DHCP4_LEASE_QUERY_RESPONSE_SENT	956
26.14.98DHCP4_LEASE_QUERY_SEND_FAILED	956
26.14.99DHCP4_LEASE_REUSE	956
26.14.100DHCP4_MULTI_THREADING_INFO	956
26.14.101DHCP4_NCR_CREATION_FAILED	956
26.14.102DHCP4_NOT_RUNNING	956
26.14.103DHCP4_NO_LEASE_INIT_REBOOT	957
26.14.104DHCP4_OPEN_SOCKET	957
26.14.105DHCP4_OPEN_SOCKETS_FAILED	957

26.14.10DHCP4_OPEN_SOCKETS_NO_RECONNECT_CTL	957
26.14.10DHCP4_PACKET_DROP_0001	957
26.14.10DHCP4_PACKET_DROP_0002	957
26.14.10DHCP4_PACKET_DROP_0003	957
26.14.11DHCP4_PACKET_DROP_0004	958
26.14.11DHCP4_PACKET_DROP_0005	958
26.14.11DHCP4_PACKET_DROP_0006	958
26.14.11DHCP4_PACKET_DROP_0007	958
26.14.11DHCP4_PACKET_DROP_0008	958
26.14.11DHCP4_PACKET_DROP_0009	958
26.14.11DHCP4_PACKET_DROP_0010	959
26.14.11DHCP4_PACKET_DROP_0011	959
26.14.11DHCP4_PACKET_DROP_0012	959
26.14.11DHCP4_PACKET_DROP_0013	959
26.14.12DHCP4_PACKET_DROP_0014	959
26.14.12DHCP4_PACKET_NAK_0001	959
26.14.12DHCP4_PACKET_NAK_0002	960
26.14.12DHCP4_PACKET_NAK_0003	960
26.14.12DHCP4_PACKET_NAK_0004	960
26.14.12DHCP4_PACKET_OPTIONS_SKIPPED	960
26.14.12DHCP4_PACKET_PACK	960
26.14.12DHCP4_PACKET_PACK_FAIL	960
26.14.12DHCP4_PACKET_PROCESS_EXCEPTION	961
26.14.12DHCP4_PACKET_PROCESS_EXCEPTION_MAIN	961
26.14.13DHCP4_PACKET_PROCESS_STD_EXCEPTION	961
26.14.13DHCP4_PACKET_PROCESS_STD_EXCEPTION_MAIN	961
26.14.13DHCP4_PACKET_QUEUE_FULL	961
26.14.13DHCP4_PACKET_RECEIVED	961
26.14.13DHCP4_PACKET_SEND	961
26.14.13DHCP4_PACKET_SEND_FAIL	962
26.14.13DHCP4_PARSER_COMMIT_EXCEPTION	962
26.14.13DHCP4_PARSER_COMMIT_FAIL	962
26.14.13DHCP4_PARSER_EXCEPTION	962
26.14.13DHCP4_PARSER_FAIL	962
26.14.14DHCP4_POST_ALLOCATION_NAME_UPDATE_FAIL	962
26.14.14DHCP4_QUERY_DATA	963
26.14.14DHCP4_QUERY_LABEL	963
26.14.14DHCP4_RECLAIM_EXPIRED_LEASES_FAIL	963
26.14.14DHCP4_RECLAIM_EXPIRED_LEASES_SKIPPED	963
26.14.14DHCP4_RECOVERED_STASHED_RELAY_AGENT_INFO	963
26.14.14DHCP4_RELEASE	963
26.14.14DHCP4_RELEASE_DELETED	963
26.14.14DHCP4_RELEASE_EXCEPTION	964
26.14.14DHCP4_RELEASE_EXPIRED	964
26.14.15DHCP4_RELEASE_FAIL	964
26.14.15DHCP4_RELEASE_FAIL_NO_LEASE	964
26.14.15DHCP4_RELEASE_FAIL_WRONG_CLIENT	964
26.14.15DHCP4_REQUEST	964
26.14.15DHCP4_RESERVATIONS_LOOKUP_FIRST_ENABLED	965
26.14.15DHCP4_RESERVED_HOSTNAME_ASSIGNED	965
26.14.15DHCP4_RESPONSE_DATA	965
26.14.15DHCP4_RESPONSE_FQDN_DATA	965
26.14.15DHCP4_RESPONSE_HOSTNAME_DATA	965
26.14.15DHCP4_RESPONSE_HOSTNAME_GENERATE	965

26.14.16	DHCP4_SECURITY_CHECKS_DISABLED	966
26.14.16	DHCP4_SERVER_FAILED	966
26.14.16	DHCP4_SERVER_INITIATED_DECLINE	966
26.14.16	DHCP4_SERVER_INITIATED_DECLINE_ADD_FAILED	966
26.14.16	DHCP4_SERVER_INITIATED_DECLINE_RESOURCE_BUSY	966
26.14.16	DHCP4_SERVER_INITIATED_DECLINE_UPDATE_FAILED	966
26.14.16	DHCP4_SHUTDOWN	967
26.14.16	DHCP4_SHUTDOWN_REQUEST	967
26.14.16	DHCP4_SRV_CONSTRUCT_ERROR	967
26.14.16	DHCP4_SRV_D2STOP_ERROR	967
26.14.17	DHCP4_SRV_DHCP4O6_ERROR	967
26.14.17	DHCP4_SRV_UNLOAD_LIBRARIES_ERROR	967
26.14.17	DHCP4_STARTED	967
26.14.17	DHCP4_STARTING	968
26.14.17	DHCP4_START_INFO	968
26.14.17	DHCP4_SUBNET_DATA	968
26.14.17	DHCP4_SUBNET_DYNAMICALLY_CHANGED	968
26.14.17	DHCP4_SUBNET_SELECTED	968
26.14.17	DHCP4_SUBNET_SELECTION_FAILED	968
26.14.17	DHCP4_TESTING_MODE_SEND_TO_SOURCE_ENABLED	968
26.14.18	DHCP4_UNKNOWN_ADDRESS_REQUESTED	969
26.14.18	DHCP4_V6_ONLY_PREFERRED_MISSING_IN_ACK	969
26.14.18	DHCP4_V6_ONLY_PREFERRED_MISSING_IN_OFFER	969
26.15	DHCP6	969
26.15.1	DHCP6_ADDITIONAL_CLASS_EVAL_ERROR	969
26.15.2	DHCP6_ADDITIONAL_CLASS_EVAL_RESULT	969
26.15.3	DHCP6_ADDITIONAL_CLASS_NO_TEST	969
26.15.4	DHCP6_ADDITIONAL_CLASS_UNDEFINED	970
26.15.5	DHCP6_ADDR_REG_INFORM_CLIENT_CHANGE	970
26.15.6	DHCP6_ADDR_REG_INFORM_FAIL	970
26.15.7	DHCP6_ADD_GLOBAL_STATUS_CODE	970
26.15.8	DHCP6_ADD_STATUS_CODE_FOR_IA	970
26.15.9	DHCP6_ALREADY_RUNNING	970
26.15.10	DHCP6_BUFFER_RECEIVED	971
26.15.11	DHCP6_BUFFER_UNPACK	971
26.15.12	DHCP6_BUFFER_WAIT_SIGNAL	971
26.15.13	DHCP6_CB_ON_DEMAND_FETCH_UPDATES_FAIL	971
26.15.14	DHCP6_CB_PERIODIC_FETCH_UPDATES_FAIL	971
26.15.15	DHCP6_CB_PERIODIC_FETCH_UPDATES_RETRIES_EXHAUSTED	971
26.15.16	DHCP6_CLASSES_ASSIGNED	972
26.15.17	DHCP6_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION	972
26.15.18	DHCP6_CLASS_ASSIGNED	972
26.15.19	DHCP6_CLASS_UNCONFIGURED	972
26.15.20	DHCP6_CLIENT_FQDN_SCRUBBED_EMPTY	972
26.15.21	DHCP6_CONFIG_COMPLETE	972
26.15.22	DHCP6_CONFIG_LOAD_FAIL	973
26.15.23	DHCP6_CONFIG_PACKET_QUEUE	973
26.15.24	DHCP6_CONFIG_RECEIVED	973
26.15.25	DHCP6_CONFIG_START	973
26.15.26	DHCP6_CONFIG_SYNTAX_WARNING	973
26.15.27	DHCP6_CONFIG_UNRECOVERABLE_ERROR	973
26.15.28	DHCP6_CONFIG_UNSUPPORTED_OBJECT	973
26.15.29	DHCP6_DATA_DIRECTORY_DEPRECATED	974
26.15.30	DHCP6_DB_RECONNECT_DISABLED	974

26.15.31DHCP6_DB_RECONNECT_FAILED	974
26.15.32DHCP6_DB_RECONNECT_LOST_CONNECTION	974
26.15.33DHCP6_DB_RECONNECT_NO_DB_CTL	974
26.15.34DHCP6_DB_RECONNECT_SUCCEEDED	974
26.15.35DHCP6_DDNS_CREATE_ADD_NAME_CHANGE_REQUEST	974
26.15.36DHCP6_DDNS_FQDN_GENERATED	975
26.15.37DHCP6_DDNS_GENERATED_FQDN_UPDATE_FAIL	975
26.15.38DHCP6_DDNS_GENERATE_FQDN	975
26.15.39DHCP6_DDNS_RECEIVE_FQDN	975
26.15.40DHCP6_DDNS_REMOVE_OLD_LEASE_FQDN	975
26.15.41DHCP6_DDNS_REQUEST_SEND_FAILED	975
26.15.42DHCP6_DDNS_RESPONSE_FQDN_DATA	976
26.15.43DHCP6_DECLINE_FAIL	976
26.15.44DHCP6_DECLINE_FAIL_DUID_MISMATCH	976
26.15.45DHCP6_DECLINE_FAIL_IAID_MISMATCH	976
26.15.46DHCP6_DECLINE_FAIL_LEASE_WITHOUT_DUID	976
26.15.47DHCP6_DECLINE_FAIL_NO_LEASE	976
26.15.48DHCP6_DECLINE_LEASE	977
26.15.49DHCP6_DECLINE_PROCESS_IA	977
26.15.50DHCP6_DEVELOPMENT_VERSION	977
26.15.51DHCP6_DHCP4O6_PACKET_RECEIVED	977
26.15.52DHCP6_DHCP4O6_RECEIVE_FAIL	977
26.15.53DHCP6_DHCP4O6_RECEIVING	977
26.15.54DHCP6_DHCP4O6_RESPONSE_DATA	977
26.15.55DHCP6_DHCP4O6_SEND_FAIL	978
26.15.56DHCP6_DYNAMIC_RECONFIGURATION	978
26.15.57DHCP6_DYNAMIC_RECONFIGURATION_FAIL	978
26.15.58DHCP6_DYNAMIC_RECONFIGURATION_SUCCESS	978
26.15.59DHCP6_FLEX_ID	978
26.15.60DHCP6_HOOK_ADDR6_REGISTER_DROP	978
26.15.61DHCP6_HOOK_ADDR6_REGISTER_SKIP	978
26.15.62DHCP6_HOOK_BUFFER_RCVD_DROP	979
26.15.63DHCP6_HOOK_BUFFER_RCVD_SKIP	979
26.15.64DHCP6_HOOK_BUFFER_SEND_SKIP	979
26.15.65DHCP6_HOOK_DDNS_UPDATE	979
26.15.66DHCP6_HOOK_DECLINE_DROP	979
26.15.67DHCP6_HOOK_DECLINE_SKIP	979
26.15.68DHCP6_HOOK_LEASE6_RELEASE_NA_SKIP	980
26.15.69DHCP6_HOOK_LEASE6_RELEASE_PD_SKIP	980
26.15.70DHCP6_HOOK_LEASES6_COMMITTED_DROP	980
26.15.71DHCP6_HOOK_LEASES6_COMMITTED_PARK	980
26.15.72DHCP6_HOOK_LEASES6_PARKING_LOT_FULL	980
26.15.73DHCP6_HOOK_PACKET_RCVD_SKIP	980
26.15.74DHCP6_HOOK_PACKET_SEND_DROP	981
26.15.75DHCP6_HOOK_PACKET_SEND_SKIP	981
26.15.76DHCP6_HOOK_SUBNET6_SELECT_DROP	981
26.15.77DHCP6_HOOK_SUBNET6_SELECT_PARK	981
26.15.78DHCP6_HOOK_SUBNET6_SELECT_SKIP	981
26.15.79DHCP6_INIT_FAIL	981
26.15.80DHCP6_LEASE_ADVERT	982
26.15.81DHCP6_LEASE_ADVERT_FAIL	982
26.15.82DHCP6_LEASE_ALLOC	982
26.15.83DHCP6_LEASE_ALLOC_FAIL	982
26.15.84DHCP6_LEASE_DATA	982

26.15.85DHCP6_LEASE_NA_WITHOUT_DUID	982
26.15.86DHCP6_LEASE_PD_WITHOUT_DUID	983
26.15.87DHCP6_LEASE_QUERY_ERROR_GETTING_RELAY_INFO	983
26.15.88DHCP6_LEASE_QUERY_PACKET_PACK	983
26.15.89DHCP6_LEASE_QUERY_PACKET_PACK_FAILED	983
26.15.90DHCP6_LEASE_QUERY_PACKET_UNPACK_FAILED	983
26.15.91DHCP6_LEASE_QUERY_PREFIX_LENGTH_LIST	983
26.15.92DHCP6_LEASE_QUERY_PROCESS_FAILED	984
26.15.93DHCP6_LEASE_QUERY_RECEIVED	984
26.15.94DHCP6_LEASE_QUERY_REPLY_SEND_FAILED	984
26.15.95DHCP6_LEASE_QUERY_REPLY_SENT	984
26.15.96DHCP6_LEASE_RENEW	984
26.15.97DHCP6_LEASE_REUSE	984
26.15.98DHCP6_MULTI_THREADING_INFO	984
26.15.99DHCP6_NOT_RUNNING	985
26.15.100DHCP6_NO_INTERFACES	985
26.15.101DHCP6_OPEN_SOCKET	985
26.15.102DHCP6_OPEN_SOCKETS_FAILED	985
26.15.103DHCP6_OPEN_SOCKETS_NO_RECONNECT_CTL	985
26.15.104DHCP6_PACKET_DROP_DHCP_DISABLED	985
26.15.105DHCP6_PACKET_DROP_DROP_CLASS	985
26.15.106DHCP6_PACKET_DROP_DROP_CLASS2	985
26.15.107DHCP6_PACKET_DROP_DROP_CLASS_EARLY	986
26.15.108DHCP6_PACKET_DROP_DUPLICATE	986
26.15.109DHCP6_PACKET_DROP_PARSE_FAIL	986
26.15.110DHCP6_PACKET_DROP_SERVERID_MISMATCH	986
26.15.111DHCP6_PACKET_DROP_UNICAST	986
26.15.112DHCP6_PACKET_OPTIONS_SKIPPED	986
26.15.113DHCP6_PACKET_PROCESS_EXCEPTION	986
26.15.114DHCP6_PACKET_PROCESS_EXCEPTION_MAIN	987
26.15.115DHCP6_PACKET_PROCESS_FAIL	987
26.15.116DHCP6_PACKET_PROCESS_STD_EXCEPTION	987
26.15.117DHCP6_PACKET_PROCESS_STD_EXCEPTION_MAIN	987
26.15.118DHCP6_PACKET_QUEUE_FULL	987
26.15.119DHCP6_PACKET_RECEIVED	987
26.15.120DHCP6_PACKET_RECEIVE_FAIL	987
26.15.121DHCP6_PACKET_SEND	988
26.15.122DHCP6_PACKET_SEND_FAIL	988
26.15.123DHCP6_PACK_FAIL	988
26.15.124DHCP6_PARSER_COMMIT_EXCEPTION	988
26.15.125DHCP6_PARSER_COMMIT_FAIL	988
26.15.126DHCP6_PARSER_EXCEPTION	988
26.15.127DHCP6_PARSER_FAIL	988
26.15.128DHCP6_PD_LEASE_ADVERT	989
26.15.129DHCP6_PD_LEASE_ADVERT_FAIL	989
26.15.130DHCP6_PD_LEASE_ALLOC	989
26.15.131DHCP6_PD_LEASE_ALLOC_FAIL	989
26.15.132DHCP6_PD_LEASE_RENEW	989
26.15.133DHCP6_PD_LEASE_REUSE	989
26.15.134DHCP6_PROCESS_IA_NA_EXTEND	990
26.15.135DHCP6_PROCESS_IA_NA_RELEASE	990
26.15.136DHCP6_PROCESS_IA_NA_REQUEST	990
26.15.137DHCP6_PROCESS_IA_NA_SOLICIT	990
26.15.138DHCP6_PROCESS_IA_PD_EXTEND	990

26.15.13	DHCP6_PROCESS_IA_PD_REQUEST	990
26.15.14	DHCP6_PROCESS_IA_PD_SOLICIT	991
26.15.14	DHCP6_QUERY_DATA	991
26.15.14	DHCP6_QUERY_LABEL	991
26.15.14	DHCP6_RAPID_COMMIT	991
26.15.14	DHCP6_RECLAIM_EXPIRED_LEASES_FAIL	991
26.15.14	DHCP6_RECLAIM_EXPIRED_LEASES_SKIPPED	991
26.15.14	DHCP6_REGISTERED_LEASE_ADD_FAIL	991
26.15.14	DHCP6_REGISTERED_LEASE_UPDATE_FAIL	992
26.15.14	DHCP6_RELEASE_NA	992
26.15.14	DHCP6_RELEASE_NA_DELETED	992
26.15.15	DHCP6_RELEASE_NA_EXPIRED	992
26.15.15	DHCP6_RELEASE_NA_FAIL	992
26.15.15	DHCP6_RELEASE_NA_FAIL_WRONG_DUID	992
26.15.15	DHCP6_RELEASE_NA_FAIL_WRONG_IAID	993
26.15.15	DHCP6_RELEASE_PD	993
26.15.15	DHCP6_RELEASE_PD_DELETED	993
26.15.15	DHCP6_RELEASE_PD_EXPIRED	993
26.15.15	DHCP6_RELEASE_PD_FAIL	993
26.15.15	DHCP6_RELEASE_PD_FAIL_WRONG_DUID	993
26.15.15	DHCP6_RELEASE_PD_FAIL_WRONG_IAID	994
26.15.16	DHCP6_REQUIRED_OPTIONS_CHECK_FAIL	994
26.15.16	DHCP6_RESERVATIONS_LOOKUP_FIRST_ENABLED	994
26.15.16	DHCP6_RESPONSE_DATA	994
26.15.16	DHCP6_SECURITY_CHECKS_DISABLED	994
26.15.16	DHCP6_SERVER_FAILED	994
26.15.16	DHCP6_SHUTDOWN	995
26.15.16	DHCP6_SHUTDOWN_REQUEST	995
26.15.16	DHCP6_SRV_CONSTRUCT_ERROR	995
26.15.16	DHCP6_SRV_D2STOP_ERROR	995
26.15.16	DHCP6_SRV_UNLOAD_LIBRARIES_ERROR	995
26.15.17	DHCP6_STARTED	995
26.15.17	DHCP6_STARTING	995
26.15.17	DHCP6_START_INFO	996
26.15.17	DHCP6_SUBNET_DATA	996
26.15.17	DHCP6_SUBNET_DYNAMICALLY_CHANGED	996
26.15.17	DHCP6_SUBNET_SELECTED	996
26.15.17	DHCP6_SUBNET_SELECTION_FAILED	996
26.15.17	DHCP6_UNKNOWN_MSG_RECEIVED	996
26.15.17	DHCP6_USING_SERVERID	997
26.16	DHCPSRV	997
26.16.1	DHCPSRV_CFGMGR_ADD_IFACE	997
26.16.2	DHCPSRV_CFGMGR_ADD_SUBNET4	997
26.16.3	DHCPSRV_CFGMGR_ADD_SUBNET6	997
26.16.4	DHCPSRV_CFGMGR_ALL_IFACES_ACTIVE	997
26.16.5	DHCPSRV_CFGMGR_CFG_DHCP_DDNS	997
26.16.6	DHCPSRV_CFGMGR_CONFIG4_MERGED	997
26.16.7	DHCPSRV_CFGMGR_CONFIG6_MERGED	998
26.16.8	DHCPSRV_CFGMGR_CONFIGURE_SERVERID	998
26.16.9	DHCPSRV_CFGMGR_DEL_SUBNET4	998
26.16.10	DHCPSRV_CFGMGR_DEL_SUBNET6	998
26.16.11	DHCPSRV_CFGMGR_FLQ_POPULATE_FREE_ADDRESS_LEASES	998
26.16.12	DHCPSRV_CFGMGR_FLQ_POPULATE_FREE_ADDRESS_LEASES_DONE	998
26.16.13	DHCPSRV_CFGMGR_FLQ_POPULATE_FREE_PREFIX_LEASES	998

26.16.14DHCP_SRV_CFGMGR_FLQ_POPULATE_FREE_PREFIX_LEASES_DONE	999
26.16.15DHCP_SRV_CFGMGR_IPV4_RESERVATIONS_NON_UNIQUE_IGNORED	999
26.16.16DHCP_SRV_CFGMGR_IPV6_RESERVATIONS_NON_UNIQUE_IGNORED	999
26.16.17DHCP_SRV_CFGMGR_IP_RESERVATIONS_UNIQUE_DUPLICATES_DETECTED	999
26.16.18DHCP_SRV_CFGMGR_IP_RESERVATIONS_UNIQUE_DUPLICATES_POSSIBLE	999
26.16.19DHCP_SRV_CFGMGR_NEW_SUBNET4	999
26.16.20DHCP_SRV_CFGMGR_NEW_SUBNET6	1000
26.16.21DHCP_SRV_CFGMGR_OPTION_DUPLICATE	1000
26.16.22DHCP_SRV_CFGMGR_RENEW_GTR_REBIND	1000
26.16.23DHCP_SRV_CFGMGR_SOCKET_RAW_UNSUPPORTED	1000
26.16.24DHCP_SRV_CFGMGR_SOCKET_TYPE_DEFAULT	1000
26.16.25DHCP_SRV_CFGMGR_SOCKET_TYPE_SELECT	1000
26.16.26DHCP_SRV_CFGMGR_SUBNET4	1001
26.16.27DHCP_SRV_CFGMGR_SUBNET4_ADDR	1001
26.16.28DHCP_SRV_CFGMGR_SUBNET4_IFACE	1001
26.16.29DHCP_SRV_CFGMGR_SUBNET4_RELAY	1001
26.16.30DHCP_SRV_CFGMGR_SUBNET6	1001
26.16.31DHCP_SRV_CFGMGR_SUBNET6_IFACE	1001
26.16.32DHCP_SRV_CFGMGR_SUBNET6_IFACE_ID	1001
26.16.33DHCP_SRV_CFGMGR_SUBNET6_RELAY	1002
26.16.34DHCP_SRV_CFGMGR_UNICAST_LINK_LOCAL	1002
26.16.35DHCP_SRV_CFGMGR_UPDATE_SUBNET4	1002
26.16.36DHCP_SRV_CFGMGR_UPDATE_SUBNET6	1002
26.16.37DHCP_SRV_CFGMGR_USE_ADDRESS	1002
26.16.38DHCP_SRV_CFGMGR_USE_ALLOCATOR	1002
26.16.39DHCP_SRV_CFGMGR_USE_UNICAST	1002
26.16.40DHCP_SRV_CLASS_WITH_ADDITIONAL_AND_LIFETIMES	1002
26.16.41DHCP_SRV_CLIENT_CLASS_DEPRECATED	1003
26.16.42DHCP_SRV_CLOSE_DB	1003
26.16.43DHCP_SRV_DDNS_TTL_TOO_LARGE	1003
26.16.44DHCP_SRV_DDNS_TTL_TOO_SMALL	1003
26.16.45DHCP_SRV_DHCP4O6_RECEIVED_BAD_PACKET	1003
26.16.46DHCP_SRV_DHCP_DDNS_ERROR_EXCEPTION	1003
26.16.47DHCP_SRV_DHCP_DDNS_HANDLER_NULL	1003
26.16.48DHCP_SRV_DHCP_DDNS_NCR_REJECTED	1004
26.16.49DHCP_SRV_DHCP_DDNS_NCR_SENT	1004
26.16.50DHCP_SRV_DHCP_DDNS_SENDER_STARTED	1004
26.16.51DHCP_SRV_DHCP_DDNS_SENDER_STOPPED	1004
26.16.52DHCP_SRV_DHCP_DDNS_SUSPEND_UPDATES	1004
26.16.53DHCP_SRV_EVAL_ERROR	1004
26.16.54DHCP_SRV_EVAL_RESULT	1004
26.16.55DHCP_SRV_FORENSIC_BACKENDS_REGISTERED	1005
26.16.56DHCP_SRV_FORENSIC_BACKEND_DEREGISTER	1005
26.16.57DHCP_SRV_FORENSIC_BACKEND_REGISTER	1005
26.16.58DHCP_SRV_HOOK_LEASE4_RECOVER_SKIP	1005
26.16.59DHCP_SRV_HOOK_LEASE4_RENEW_SKIP	1005
26.16.60DHCP_SRV_HOOK_LEASE4_SELECT_SKIP	1005
26.16.61DHCP_SRV_HOOK_LEASE6_EXTEND_SKIP	1005
26.16.62DHCP_SRV_HOOK_LEASE6_RECOVER_SKIP	1006
26.16.63DHCP_SRV_HOOK_LEASE6_SELECT_SKIP	1006
26.16.64DHCP_SRV_HOST_MGR_DB_OPEN_CONNECTION_WITH_RETRY_FAILED	1006
26.16.65DHCP_SRV_LEASE4_EXTENDED_INFO_SANITY_FAIL	1006
26.16.66DHCP_SRV_LEASE4_EXTENDED_INFO_UPGRADED	1006
26.16.67DHCP_SRV_LEASE6_EXTENDED_INFO_SANITY_FAIL	1006

26.16.68DHCP	SRV_LEASE6_EXTENDED_INFO_UPGRADED	1006
26.16.69DHCP	SRV_LEASE_MGR_BACKENDS_REGISTERED	1007
26.16.70DHCP	SRV_LEASE_MGR_BACKEND_DEREGISTER	1007
26.16.71DHCP	SRV_LEASE_MGR_BACKEND_REGISTER	1007
26.16.72DHCP	SRV_LEASE_MGR_CALLBACK_EXCEPTION	1007
26.16.73DHCP	SRV_LEASE_MGR_CALLBACK_UNKNOWN_EXCEPTION	1007
26.16.74DHCP	SRV_LEASE_MGR_DB_OPEN_CONNECTION_WITH_RETRY_FAILED	1007
26.16.75DHCP	SRV_LEASE_SANITY_FAIL	1008
26.16.76DHCP	SRV_LEASE_SANITY_FAIL_DISCARD	1008
26.16.77DHCP	SRV_LEASE_SANITY_FIXED	1008
26.16.78DHCP	SRV_MEMFILE_ADD_ADDR4	1008
26.16.79DHCP	SRV_MEMFILE_ADD_ADDR6	1008
26.16.80DHCP	SRV_MEMFILE_BEGIN_BUILD_EXTENDED_INFO_TABLES6	1008
26.16.81DHCP	SRV_MEMFILE_BEGIN_EXTRACT_EXTENDED_INFO4	1008
26.16.82DHCP	SRV_MEMFILE_BUILD_EXTENDED_INFO_TABLES6	1009
26.16.83DHCP	SRV_MEMFILE_BUILD_EXTENDED_INFO_TABLES6_ERROR	1009
26.16.84DHCP	SRV_MEMFILE_COMMIT	1009
26.16.85DHCP	SRV_MEMFILE_CONVERTING_LEASE_FILES	1009
26.16.86DHCP	SRV_MEMFILE_DB	1009
26.16.87DHCP	SRV_MEMFILE_DELETE_ADDR4	1009
26.16.88DHCP	SRV_MEMFILE_DELETE_ADDR6	1009
26.16.89DHCP	SRV_MEMFILE_DELETE_EXPIRED_RECLAIMED4	1010
26.16.90DHCP	SRV_MEMFILE_DELETE_EXPIRED_RECLAIMED6	1010
26.16.91DHCP	SRV_MEMFILE_DELETE_EXPIRED_RECLAIMED_START	1010
26.16.92DHCP	SRV_MEMFILE_EXTRACT_EXTENDED_INFO4	1010
26.16.93DHCP	SRV_MEMFILE_EXTRACT_EXTENDED_INFO4_ERROR	1010
26.16.94DHCP	SRV_MEMFILE_FAILED_TO_OPEN	1010
26.16.95DHCP	SRV_MEMFILE_GET4	1010
26.16.96DHCP	SRV_MEMFILE_GET6	1011
26.16.97DHCP	SRV_MEMFILE_GET6_DUID	1011
26.16.98DHCP	SRV_MEMFILE_GET_ADDR4	1011
26.16.99DHCP	SRV_MEMFILE_GET_ADDR6	1011
26.16.100DHCP	SRV_MEMFILE_GET_CLIENTID	1011
26.16.100DHCP	SRV_MEMFILE_GET_EXPIRED4	1011
26.16.100DHCP	SRV_MEMFILE_GET_EXPIRED6	1011
26.16.100DHCP	SRV_MEMFILE_GET_HOSTNAME4	1012
26.16.100DHCP	SRV_MEMFILE_GET_HOSTNAME6	1012
26.16.100DHCP	SRV_MEMFILE_GET_HWADDR	1012
26.16.100DHCP	SRV_MEMFILE_GET_IAID_DUID	1012
26.16.100DHCP	SRV_MEMFILE_GET_IAID_SUBID_DUID	1012
26.16.100DHCP	SRV_MEMFILE_GET_PAGE4	1012
26.16.100DHCP	SRV_MEMFILE_GET_PAGE6	1012
26.16.110DHCP	SRV_MEMFILE_GET_RELAYID4	1013
26.16.110DHCP	SRV_MEMFILE_GET_RELAYID6	1013
26.16.110DHCP	SRV_MEMFILE_GET_REMOTEID4	1013
26.16.110DHCP	SRV_MEMFILE_GET_REMOTEID6	1013
26.16.110DHCP	SRV_MEMFILE_GET_SUBID4	1013
26.16.110DHCP	SRV_MEMFILE_GET_SUBID6	1013
26.16.110DHCP	SRV_MEMFILE_GET_SUBID_CLIENTID	1013
26.16.110DHCP	SRV_MEMFILE_GET_SUBID_HWADDR	1014
26.16.110DHCP	SRV_MEMFILE_GET_SUBID_PAGE6	1014
26.16.110DHCP	SRV_MEMFILE_LEASE_FILE_LOAD	1014
26.16.120DHCP	SRV_MEMFILE_LEASE_LOAD	1014
26.16.12DHCP	SRV_MEMFILE_LEASE_LOAD_ROW_ERROR	1014

26.16.12	DHCP_SRV_MEMFILE_LFC_EXECUTE	1014
26.16.12	DHCP_SRV_MEMFILE_LFC_LEASE_FILE_RENAME_FAIL	1014
26.16.12	DHCP_SRV_MEMFILE_LFC_LEASE_FILE_REOPEN_FAIL	1015
26.16.12	DHCP_SRV_MEMFILE_LFC_SETUP	1015
26.16.12	DHCP_SRV_MEMFILE_LFC_SPAWN_FAIL	1015
26.16.12	DHCP_SRV_MEMFILE_LFC_START	1015
26.16.12	DHCP_SRV_MEMFILE_LFC_UNREGISTER_TIMER_FAILED	1015
26.16.12	DHCP_SRV_MEMFILE_NEEDS_DOWNGRADING	1015
26.16.13	DHCP_SRV_MEMFILE_NEEDS_UPGRADING	1015
26.16.13	DHCP_SRV_MEMFILE_NO_STORAGE	1016
26.16.13	DHCP_SRV_MEMFILE_READ_HWADDR_FAIL	1016
26.16.13	DHCP_SRV_MEMFILE_ROLLBACK	1016
26.16.13	DHCP_SRV_MEMFILE_UPDATE_ADDR4	1016
26.16.13	DHCP_SRV_MEMFILE_UPDATE_ADDR6	1016
26.16.13	DHCP_SRV_MEMFILE_WIPE_LEASES4	1016
26.16.13	DHCP_SRV_MEMFILE_WIPE_LEASES4_FINISHED	1016
26.16.13	DHCP_SRV_MEMFILE_WIPE_LEASES6	1017
26.16.13	DHCP_SRV_MEMFILE_WIPE_LEASES6_FINISHED	1017
26.16.14	DHCP_SRV_MT_DISABLED_QUEUE_CONTROL	1017
26.16.14	DHCP_SRV_MULTIPLE_RAW_SOCKETS_PER_IFACE	1017
26.16.14	DHCP_SRV_NOTYPE_DB	1017
26.16.14	DHCP_SRV_NO_SOCKETS_OPEN	1017
26.16.14	DHCP_SRV_ONLY_IF_REQUIRED_DEPRECATED	1017
26.16.14	DHCP_SRV_OPEN_SOCKET_FAIL	1018
26.16.14	DHCP_SRV_QUEUE_NCR	1018
26.16.14	DHCP_SRV_QUEUE_NCR_FAILED	1018
26.16.14	DHCP_SRV_QUEUE_NCR_SKIP	1018
26.16.14	DHCP_SRV_REQUIRE_CLIENT_CLASSES_DEPRECATED	1018
26.16.15	DHCP_SRV_SUBNET4O6_SELECT_FAILED	1018
26.16.15	DHCP_SRV_SUBNET4_SELECT_BY_ADDRESS_NO_MATCH	1018
26.16.15	DHCP_SRV_SUBNET4_SELECT_BY_INTERFACE_NO_MATCH	1019
26.16.15	DHCP_SRV_SUBNET4_SELECT_BY_RELAY_ADDRESS_NO_MATCH	1019
26.16.15	DHCP_SRV_SUBNET4_SELECT_NO_RA_OPTIONS	1019
26.16.15	DHCP_SRV_SUBNET4_SELECT_NO_RELAY_ADDRESS	1019
26.16.15	DHCP_SRV_SUBNET4_SELECT_NO_USABLE_ADDRESS	1019
26.16.15	DHCP_SRV_SUBNET6_SELECT_BY_ADDRESS_NO_MATCH	1019
26.16.15	DHCP_SRV_SUBNET6_SELECT_BY_INTERFACE_ID_NO_MATCH	1019
26.16.15	DHCP_SRV_SUBNET6_SELECT_BY_INTERFACE_NO_MATCH	1019
26.16.16	DHCP_SRV_TEMPLATE_EVAL_ERROR	1020
26.16.16	DHCP_SRV_TEMPLATE_EVAL_RESULT	1020
26.16.16	DHCP_SRV_TIMERMGR_CALLBACK_FAILED	1020
26.16.16	DHCP_SRV_TIMERMGR_REGISTER_TIMER	1020
26.16.16	DHCP_SRV_TIMERMGR_RUN_TIMER_OPERATION	1020
26.16.16	DHCP_SRV_TIMERMGR_START_TIMER	1020
26.16.16	DHCP_SRV_TIMERMGR_STOP_TIMER	1020
26.16.16	DHCP_SRV_TIMERMGR_UNREGISTER_ALL_TIMERS	1021
26.16.16	DHCP_SRV_TIMERMGR_UNREGISTER_TIMER	1021
26.16.16	DHCP_SRV_UNKNOWN_DB	1021
26.17	DHCP	1021
26.17.1	DHCP_DDNS_ADD_FAILED	1021
26.17.2	DHCP_DDNS_ADD_SUCCEEDED	1021
26.17.3	DHCP_DDNS_AT_MAX_TRANSACTIONS	1021
26.17.4	DHCP_DDNS_CLEARED_FOR_SHUTDOWN	1021
26.17.5	DHCP_DDNS_CONFIGURE	1022

26.17.6	DHCP_DDNS_CONFIGURED_CALLOUT_DROP	1022
26.17.7	DHCP_DDNS_CONFIG_CHECK_FAIL	1022
26.17.8	DHCP_DDNS_CONFIG_FAIL	1022
26.17.9	DHCP_DDNS_CONFIG_SYNTAX_WARNING	1022
26.17.10	DHCP_DDNS_FAILED	1022
26.17.11	DHCP_DDNS_FORWARD_ADD_BAD_DNSCLIENT_STATUS	1022
26.17.12	DHCP_DDNS_FORWARD_ADD_BUILD_FAILURE	1023
26.17.13	DHCP_DDNS_FORWARD_ADD_IO_ERROR	1023
26.17.14	DHCP_DDNS_FORWARD_ADD_REJECTED	1023
26.17.15	DHCP_DDNS_FORWARD_ADD_RESP_CORRUPT	1023
26.17.16	DHCP_DDNS_FORWARD_ADD_TIMEOUT	1023
26.17.17	DHCP_DDNS_FORWARD_REMOVE_ADDRS_BAD_DNSCLIENT_STATUS	1023
26.17.18	DHCP_DDNS_FORWARD_REMOVE_ADDRS_BUILD_FAILURE	1024
26.17.19	DHCP_DDNS_FORWARD_REMOVE_ADDRS_IO_ERROR	1024
26.17.20	DHCP_DDNS_FORWARD_REMOVE_ADDRS_REJECTED	1024
26.17.21	DHCP_DDNS_FORWARD_REMOVE_ADDRS_RESP_CORRUPT	1024
26.17.22	DHCP_DDNS_FORWARD_REMOVE_ADDRS_TIMEOUT	1024
26.17.23	DHCP_DDNS_FORWARD_REMOVE_RRS_BAD_DNSCLIENT_STATUS	1024
26.17.24	DHCP_DDNS_FORWARD_REMOVE_RRS_BUILD_FAILURE	1025
26.17.25	DHCP_DDNS_FORWARD_REMOVE_RRS_IO_ERROR	1025
26.17.26	DHCP_DDNS_FORWARD_REMOVE_RRS_REJECTED	1025
26.17.27	DHCP_DDNS_FORWARD_REMOVE_RRS_RESP_CORRUPT	1025
26.17.28	DHCP_DDNS_FORWARD_REMOVE_RRS_TIMEOUT	1025
26.17.29	DHCP_DDNS_FORWARD_REPLACE_BAD_DNSCLIENT_STATUS	1025
26.17.30	DHCP_DDNS_FORWARD_REPLACE_BUILD_FAILURE	1026
26.17.31	DHCP_DDNS_FORWARD_REPLACE_IO_ERROR	1026
26.17.32	DHCP_DDNS_FORWARD_REPLACE_REJECTED	1026
26.17.33	DHCP_DDNS_FORWARD_REPLACE_RESP_CORRUPT	1026
26.17.34	DHCP_DDNS_FORWARD_REPLACE_TIMEOUT	1026
26.17.35	DHCP_DDNS_FWD_REQUEST_IGNORED	1026
26.17.36	DHCP_DDNS_INVALID_NCR	1027
26.17.37	DHCP_DDNS_INVALID_RESPONSE	1027
26.17.38	DHCP_DDNS_LISTENING_ON_ALL_INTERFACES	1027
26.17.39	DHCP_DDNS_NCR_FLUSH_IO_ERROR	1027
26.17.40	DHCP_DDNS_NCR_LISTEN_CLOSE_ERROR	1027
26.17.41	DHCP_DDNS_NCR_RECV_NEXT_ERROR	1027
26.17.42	DHCP_DDNS_NCR_SEND_CLOSE_ERROR	1028
26.17.43	DHCP_DDNS_NCR_SEND_NEXT_ERROR	1028
26.17.44	DHCP_DDNS_NCR_UDP_CLEAR_READY_ERROR	1028
26.17.45	DHCP_DDNS_NCR_UDP_RECV_CANCELED	1028
26.17.46	DHCP_DDNS_NCR_UDP_RECV_ERROR	1028
26.17.47	DHCP_DDNS_NCR_UDP_SEND_CANCELED	1028
26.17.48	DHCP_DDNS_NCR_UDP_SEND_ERROR	1028
26.17.49	DHCP_DDNS_NOT_ON_LOOPBACK	1029
26.17.50	DHCP_DDNS_NO_ELIGIBLE_JOBS	1029
26.17.51	DHCP_DDNS_NO_FWD_MATCH_ERROR	1029
26.17.52	DHCP_DDNS_NO_MATCH	1029
26.17.53	DHCP_DDNS_NO_REV_MATCH_ERROR	1029
26.17.54	DHCP_DDNS_QUEUE_MGR_QUEUE_FULL	1029
26.17.55	DHCP_DDNS_QUEUE_MGR_QUEUE_RECEIVE	1030
26.17.56	DHCP_DDNS_QUEUE_MGR_RECONFIGURING	1030
26.17.57	DHCP_DDNS_QUEUE_MGR_RECOVERING	1030
26.17.58	DHCP_DDNS_QUEUE_MGR_RECV_ERROR	1030
26.17.59	DHCP_DDNS_QUEUE_MGR_RESUME_ERROR	1030

26.17.60DHCP_DDNS_QUEUE_MGR_RESUMING	1030
26.17.61DHCP_DDNS_QUEUE_MGR_STARTED	1030
26.17.62DHCP_DDNS_QUEUE_MGR_START_ERROR	1031
26.17.63DHCP_DDNS_QUEUE_MGR_STOPPED	1031
26.17.64DHCP_DDNS_QUEUE_MGR_STOPPING	1031
26.17.65DHCP_DDNS_QUEUE_MGR_STOP_ERROR	1031
26.17.66DHCP_DDNS_QUEUE_MGR_UNEXPECTED_HANDLER_ERROR	1031
26.17.67DHCP_DDNS_QUEUE_MGR_UNEXPECTED_STOP	1031
26.17.68DHCP_DDNS_REMOVE_FAILED	1031
26.17.69DHCP_DDNS_REMOVE_SUCCEEDED	1032
26.17.70DHCP_DDNS_REQUEST_DROPPED	1032
26.17.71DHCP_DDNS_REVERSE_REMOVE_BAD_DNSCLIENT_STATUS	1032
26.17.72DHCP_DDNS_REVERSE_REMOVE_BUILD_FAILURE	1032
26.17.73DHCP_DDNS_REVERSE_REMOVE_IO_ERROR	1032
26.17.74DHCP_DDNS_REVERSE_REMOVE_REJECTED	1032
26.17.75DHCP_DDNS_REVERSE_REMOVE_RESP_CORRUPT	1033
26.17.76DHCP_DDNS_REVERSE_REMOVE_TIMEOUT	1033
26.17.77DHCP_DDNS_REVERSE_REPLACE_BAD_DNSCLIENT_STATUS	1033
26.17.78DHCP_DDNS_REVERSE_REPLACE_BUILD_FAILURE	1033
26.17.79DHCP_DDNS_REVERSE_REPLACE_IO_ERROR	1033
26.17.80DHCP_DDNS_REVERSE_REPLACE_REJECTED	1033
26.17.81DHCP_DDNS_REVERSE_REPLACE_RESP_CORRUPT	1034
26.17.82DHCP_DDNS_REVERSE_REPLACE_TIMEOUT	1034
26.17.83DHCP_DDNS_REV_REQUEST_IGNORED	1034
26.17.84DHCP_DDNS_RUN_EXIT	1034
26.17.85DHCP_DDNS_SECURITY_CHECKS_DISABLED	1034
26.17.86DHCP_DDNS_SHUTDOWN_COMMAND	1034
26.17.87DHCP_DDNS_STARTED	1034
26.17.88DHCP_DDNS_STARTING_TRANSACTION	1035
26.17.89DHCP_DDNS_STATE_MODEL_UNEXPECTED_ERROR	1035
26.17.90DHCP_DDNS_TRANS_SEND_ERROR	1035
26.17.91DHCP_DDNS_UDP_SENDER_WATCH_SOCKET_CLOSE_ERROR	1035
26.17.92DHCP_DDNS_UNCAUGHT_NCR_RECV_HANDLER_ERROR	1035
26.17.93DHCP_DDNS_UNCAUGHT_NCR_SEND_HANDLER_ERROR	1035
26.17.94DHCP_DDNS_UPDATE_REQUEST_SENT	1035
26.17.95DHCP_DDNS_UPDATE_RESPONSE_RECEIVED	1036
26.18 EVAL	1036
26.18.1 EVAL_DEBUG_AND	1036
26.18.2 EVAL_DEBUG_BRANCH	1036
26.18.3 EVAL_DEBUG_CONCAT	1036
26.18.4 EVAL_DEBUG_EQUAL	1036
26.18.5 EVAL_DEBUG_HEXSTRING	1036
26.18.6 EVAL_DEBUG_IFELSE_FALSE	1036
26.18.7 EVAL_DEBUG_IFELSE_TRUE	1037
26.18.8 EVAL_DEBUG_INT16TOTEXT	1037
26.18.9 EVAL_DEBUG_INT32TOTEXT	1037
26.18.10EVAL_DEBUG_INT8TOTEXT	1037
26.18.11EVAL_DEBUG_IPADDRESS	1037
26.18.12EVAL_DEBUG_IPADDRESSSTOTEXT	1037
26.18.13EVAL_DEBUG_LCASE	1037
26.18.14EVAL_DEBUG_MATCH	1038
26.18.15EVAL_DEBUG_MATCH_ERROR	1038
26.18.16EVAL_DEBUG_MEMBER	1038
26.18.17EVAL_DEBUG_NOT	1038

26.18.18	EVAL_DEBUG_OPTION	1038
26.18.19	EVAL_DEBUG_OR	1038
26.18.20	EVAL_DEBUG_PKT	1038
26.18.21	EVAL_DEBUG_PKT4	1039
26.18.22	EVAL_DEBUG_PKT6	1039
26.18.23	EVAL_DEBUG_POP_AND_BRANCH_FALSE	1039
26.18.24	EVAL_DEBUG_POP_OR_BRANCH_FALSE	1039
26.18.25	EVAL_DEBUG_POP_OR_BRANCH_TRUE	1039
26.18.26	EVAL_DEBUG_RELAY6	1039
26.18.27	EVAL_DEBUG_RELAY6_RANGE	1039
26.18.28	EVAL_DEBUG_SPLIT	1040
26.18.29	EVAL_DEBUG_SPLIT_DELIM_EMPTY	1040
26.18.30	EVAL_DEBUG_SPLIT_EMPTY	1040
26.18.31	EVAL_DEBUG_SPLIT_FIELD_OUT_OF_RANGE	1040
26.18.32	EVAL_DEBUG_STRING	1040
26.18.33	EVAL_DEBUG_SUBSTRING	1040
26.18.34	EVAL_DEBUG_SUBSTRING_EMPTY	1040
26.18.35	EVAL_DEBUG_SUBSTRING_RANGE	1041
26.18.36	EVAL_DEBUG_SUB_OPTION	1041
26.18.37	EVAL_DEBUG_SUB_OPTION_NO_OPTION	1041
26.18.38	EVAL_DEBUG_TOHEXSTRING	1041
26.18.39	EVAL_DEBUG_UCASE	1041
26.18.40	EVAL_DEBUG_UINT16TOTEXT	1041
26.18.41	EVAL_DEBUG_UINT32TOTEXT	1041
26.18.42	EVAL_DEBUG_UINT8TOTEXT	1042
26.18.43	EVAL_DEBUG_VENDOR_CLASS_DATA	1042
26.18.44	EVAL_DEBUG_VENDOR_CLASS_DATA_NOT_FOUND	1042
26.18.45	EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID	1042
26.18.46	EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID_MISMATCH	1042
26.18.47	EVAL_DEBUG_VENDOR_CLASS_EXISTS	1042
26.18.48	EVAL_DEBUG_VENDOR_CLASS_NO_OPTION	1042
26.18.49	EVAL_DEBUG_VENDOR_ENTERPRISE_ID	1043
26.18.50	EVAL_DEBUG_VENDOR_ENTERPRISE_ID_MISMATCH	1043
26.18.51	EVAL_DEBUG_VENDOR_EXISTS	1043
26.18.52	EVAL_DEBUG_VENDOR_NO_OPTION	1043
26.19	FLEX	1043
26.19.1	FLEX_ID_EXPRESSION_EMPTY	1043
26.19.2	FLEX_ID_EXPRESSION_EVALUATED	1043
26.19.3	FLEX_ID_EXPRESSION_EVALUATED_NP	1043
26.19.4	FLEX_ID_EXPRESSION_HEX	1044
26.19.5	FLEX_ID_EXPRESSION_INVALID_JSON_TYPE	1044
26.19.6	FLEX_ID_EXPRESSION_NOT_DEFINED	1044
26.19.7	FLEX_ID_EXPRESSION_PARSE_FAILED	1044
26.19.8	FLEX_ID_IGNORE_IAID_APPLIED_ON_NA	1044
26.19.9	FLEX_ID_IGNORE_IAID_APPLIED_ON_PD	1044
26.19.10	FLEX_ID_IGNORE_IAID_ENABLED	1044
26.19.11	FLEX_ID_IGNORE_IAID_JSON_TYPE	1045
26.19.12	FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_NA	1045
26.19.13	FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_PD	1045
26.19.14	FLEX_ID_LOAD_ERROR	1045
26.19.15	FLEX_ID_REPLACE_CLIENT_ID_JSON_TYPE	1045
26.19.16	FLEX_ID_RESTORE_CLIENT_ID	1045
26.19.17	FLEX_ID_RESTORE_DUID	1045
26.19.18	FLEX_ID_UNLOAD	1046

26.19.19	FLEX_ID_USED_AS_CLIENT_ID	1046
26.19.20	FLEX_ID_USED_AS_DUID	1046
26.19.21	FLEX_OPTION_LOAD_ERROR	1046
26.19.22	FLEX_OPTION_PROCESS_ADD	1046
26.19.23	FLEX_OPTION_PROCESS_CLIENT_CLASS	1046
26.19.24	FLEX_OPTION_PROCESS_ERROR	1046
26.19.25	FLEX_OPTION_PROCESS_REMOVE	1047
26.19.26	FLEX_OPTION_PROCESS_SUB_ADD	1047
26.19.27	FLEX_OPTION_PROCESS_SUB_CLIENT_CLASS	1047
26.19.28	FLEX_OPTION_PROCESS_SUB_REMOVE	1047
26.19.29	FLEX_OPTION_PROCESS_SUB_SUPERSEDE	1047
26.19.30	FLEX_OPTION_PROCESS_SUPERSEDE	1047
26.19.31	FLEX_OPTION_PROCESS_VENDOR_ID_MISMATCH	1047
26.19.32	FLEX_OPTION_UNLOAD	1048
26.20	FUZZ	1048
26.20.1	FUZZ_DATA_READ	1048
26.20.2	FUZZ_INIT_COMPLETE	1048
26.20.3	FUZZ_INIT_FAIL	1048
26.20.4	FUZZ_READ_FAIL	1048
26.20.5	FUZZ_SEND	1048
26.20.6	FUZZ_SEND_ERROR	1048
26.20.7	FUZZ_SHORT_SEND	1049
26.20.8	FUZZ_SOCKET_CREATE_FAIL	1049
26.21	GSS	1049
26.21.1	GSS_TSIG_COMMAND_PROCESSED_FAILED	1049
26.21.2	GSS_TSIG_LOAD_FAILED	1049
26.21.3	GSS_TSIG_LOAD_OK	1049
26.21.4	GSS_TSIG_MANAGER_STARTED	1049
26.21.5	GSS_TSIG_MANAGER_STOPPED	1049
26.21.6	GSS_TSIG_MANAGER_STOP_ERROR	1050
26.21.7	GSS_TSIG_MANAGER_STOP_GENERAL_ERROR	1050
26.21.8	GSS_TSIG_NEW_KEY	1050
26.21.9	GSS_TSIG_NEW_KEY_SETUP_FAILED	1050
26.21.10	GSS_TSIG_NEW_KEY_SETUP_SUCCEED	1050
26.21.11	GSS_TSIG_OLD_KEY_REMOVED	1050
26.21.12	GSS_TSIG_UNLOAD_OK	1050
26.21.13	GSS_TSIG_VERIFIED	1050
26.21.14	GSS_TSIG_VERIFY_FAILED	1051
26.22	HA	1051
26.22.1	HA_BUFFER4_RECEIVE_FAILED	1051
26.22.2	HA_BUFFER4_RECEIVE_NOT_FOR_US	1051
26.22.3	HA_BUFFER4_RECEIVE_PACKET_OPTIONS_SKIPPED	1051
26.22.4	HA_BUFFER4_RECEIVE_UNPACK_FAILED	1051
26.22.5	HA_BUFFER6_RECEIVE_FAILED	1051
26.22.6	HA_BUFFER6_RECEIVE_NOT_FOR_US	1051
26.22.7	HA_BUFFER6_RECEIVE_PACKET_OPTIONS_SKIPPED	1052
26.22.8	HA_BUFFER6_RECEIVE_UNPACK_FAILED	1052
26.22.9	HA_COMMAND_PROCESSED_FAILED	1052
26.22.10	HA_COMMUNICATION_INTERRUPTED	1052
26.22.11	HA_COMMUNICATION_INTERRUPTED_CLIENT4	1052
26.22.12	HA_COMMUNICATION_INTERRUPTED_CLIENT4_UNACKED	1052
26.22.13	HA_COMMUNICATION_INTERRUPTED_CLIENT6	1053
26.22.14	HA_COMMUNICATION_INTERRUPTED_CLIENT6_UNACKED	1053
26.22.15	HA_CONFIGURATION_FAILED	1053

26.22.16	HA_CONFIGURATION_SUCCESSFUL	1053
26.22.17	HA_CONFIG_AUTO_FAILOVER_DISABLED	1053
26.22.18	HA_CONFIG_DHCP_MT_DISABLED	1054
26.22.19	HA_CONFIG_DHCP_MT_DISABLED_AND_KEA_MT_ENABLED	1054
26.22.20	HA_CONFIG_LEASE_SYNCING_DISABLED	1054
26.22.21	HA_CONFIG_LEASE_SYNCING_DISABLED_REMINDER	1054
26.22.22	HA_CONFIG_LEASE_UPDATES_AND_SYNCING_DIFFER	1054
26.22.23	HA_CONFIG_LEASE_UPDATES_DISABLED	1054
26.22.24	HA_CONFIG_LEASE_UPDATES_DISABLED_REMINDER	1055
26.22.25	HA_CONFIG_SYSTEM_MT_UNSUPPORTED	1055
26.22.26	HA_CONTINUE_HANDLER_FAILED	1055
26.22.27	HA_DEINIT_OK	1055
26.22.28	HA_DHCP4_START_SERVICE_FAILED	1055
26.22.29	HA_DHCP6_START_SERVICE_FAILED	1055
26.22.30	HA_DHCP_DISABLE_COMMUNICATIONS_FAILED	1055
26.22.31	HA_DHCP_DISABLE_FAILED	1056
26.22.32	HA_DHCP_ENABLE_COMMUNICATIONS_FAILED	1056
26.22.33	HA_DHCP_ENABLE_FAILED	1056
26.22.34	HA_HEARTBEAT_COMMUNICATIONS_FAILED	1056
26.22.35	HA_HEARTBEAT_FAILED	1056
26.22.36	HA_HEARTBEAT_HANDLER_FAILED	1056
26.22.37	HA_HIGH_CLOCK_SKEW	1056
26.22.38	HA_HIGH_CLOCK_SKEW_CAUSED_TERMINATION	1057
26.22.39	HA_INIT_OK	1057
26.22.40	HA_INVALID_PARTNER_STATE_COMMUNICATION_RECOVERY	1057
26.22.41	HA_INVALID_PARTNER_STATE_HOT_STANDBY	1057
26.22.42	HA_INVALID_PARTNER_STATE_LOAD_BALANCING	1057
26.22.43	HA_LEASE4_EXPIRE_FAILED	1057
26.22.44	HA_LEASE4_EXPIRE_INVALID_HA_SERVER_NAME	1057
26.22.45	HA_LEASE4_EXPIRE_RECLAMATION_SKIP	1058
26.22.46	HA_LEASE4_SERVER_DECLINE_FAILED	1058
26.22.47	HA_LEASE6_EXPIRE_FAILED	1058
26.22.48	HA_LEASE6_EXPIRE_INVALID_HA_SERVER_NAME	1058
26.22.49	HA_LEASE6_EXPIRE_RECLAMATION_SKIP	1058
26.22.50	HA_LEASES4_COMMITTED_FAILED	1058
26.22.51	HA_LEASES4_COMMITTED_NOTHING_TO_UPDATE	1058
26.22.52	HA_LEASES4_COMMITTED_NO_RELATIONSHIP	1059
26.22.53	HA_LEASES6_COMMITTED_FAILED	1059
26.22.54	HA_LEASES6_COMMITTED_NOTHING_TO_UPDATE	1059
26.22.55	HA_LEASES6_COMMITTED_NO_RELATIONSHIP	1059
26.22.56	HA_LEASES_BACKLOG_COMMUNICATIONS_FAILED	1059
26.22.57	HA_LEASES_BACKLOG_FAILED	1059
26.22.58	HA_LEASES_BACKLOG_NOTHING_TO_SEND	1060
26.22.59	HA_LEASES_BACKLOG_START	1060
26.22.60	HA_LEASES_BACKLOG_SUCCESS	1060
26.22.61	HA_LEASES_SYNC_APPLIED_LEASES	1060
26.22.62	HA_LEASES_SYNC_COMMUNICATIONS_FAILED	1060
26.22.63	HA_LEASES_SYNC_FAILED	1060
26.22.64	HA_LEASES_SYNC_LEASE_PAGE_RECEIVED	1060
26.22.65	HA_LEASE_SYNC_FAILED	1061
26.22.66	HA_LEASE_SYNC_STALE_LEASE4_SKIP	1061
26.22.67	HA_LEASE_SYNC_STALE_LEASE6_SKIP	1061
26.22.68	HA_LEASE_UPDATES_DISABLED	1061
26.22.69	HA_LEASE_UPDATES_ENABLED	1061

26.22.70	HA_LEASE_UPDATE_COMMUNICATIONS_FAILED	1061
26.22.71	HA_LEASE_UPDATE_CONFLICT	1062
26.22.72	HA_LEASE_UPDATE_CREATE_UPDATE_FAILED_ON_PEER	1062
26.22.73	HA_LEASE_UPDATE_DELETE_FAILED_ON_PEER	1062
26.22.74	HA_LEASE_UPDATE_FAILED	1062
26.22.75	HA_LEASE_UPDATE_REJECTS_CAUSED_TERMINATION	1062
26.22.76	HA_LOAD_BALANCING_DUID_MISSING	1062
26.22.77	HA_LOAD_BALANCING_IDENTIFIER_MISSING	1062
26.22.78	HA_LOAD_BALANCING_LEASE_DUID_MISSING	1063
26.22.79	HA_LOAD_BALANCING_LEASE_IDENTIFIER_MISSING	1063
26.22.80	HA_LOCAL_DHCP_DISABLE	1063
26.22.81	HA_LOCAL_DHCP_ENABLE	1063
26.22.82	HA_MAINTENANCE_CANCEL_HANDLER_FAILED	1063
26.22.83	HA_MAINTENANCE_NOTIFY_CANCEL_COMMUNICATIONS_FAILED	1063
26.22.84	HA_MAINTENANCE_NOTIFY_CANCEL_FAILED	1063
26.22.85	HA_MAINTENANCE_NOTIFY_COMMUNICATIONS_FAILED	1064
26.22.86	HA_MAINTENANCE_NOTIFY_FAILED	1064
26.22.87	HA_MAINTENANCE_NOTIFY_HANDLER_FAILED	1064
26.22.88	HA_MAINTENANCE_SHUTDOWN_SAFE	1064
26.22.89	HA_MAINTENANCE_STARTED	1064
26.22.90	HA_MAINTENANCE_STARTED_IN_PARTNER_DOWN	1064
26.22.91	HA_MAINTENANCE_START_HANDLER_FAILED	1065
26.22.92	HA_MISSING_CONFIGURATION	1065
26.22.93	HA_PAUSE_CLIENT_LISTENER_FAILED	1065
26.22.94	HA_PAUSE_CLIENT_LISTENER_ILLEGAL	1065
26.22.95	HA_RESET_COMMUNICATIONS_FAILED	1065
26.22.96	HA_RESET_FAILED	1065
26.22.97	HA_RESET_HANDLER_FAILED	1065
26.22.98	HA_RESUME_CLIENT_LISTENER_FAILED	1066
26.22.99	HA_SCOPES_HANDLER_FAILED	1066
26.22.100	HA_SERVICE_STARTED	1066
26.22.101	HA_STATE_MACHINE_CONTINUED	1066
26.22.102	HA_STATE_MACHINE_PAUSED	1066
26.22.103	HA_STATE_TRANSITION	1066
26.22.104	HA_STATE_TRANSITION_PASSIVE_BACKUP	1066
26.22.105	HA_SUBNET4_SELECT_FAILED	1067
26.22.106	HA_SUBNET4_SELECT_INVALID_HA_SERVER_NAME	1067
26.22.107	HA_SUBNET4_SELECT_NOT_FOR_US	1067
26.22.108	HA_SUBNET4_SELECT_NO_RELATIONSHIP_FOR_SUBNET	1067
26.22.109	HA_SUBNET4_SELECT_NO_RELATIONSHIP_SELECTOR_FOR_SUBNET	1067
26.22.110	HA_SUBNET4_SELECT_NO_SUBNET_SELECTED	1067
26.22.111	HA_SUBNET6_SELECT_FAILED	1068
26.22.112	HA_SUBNET6_SELECT_INVALID_HA_SERVER_NAME	1068
26.22.113	HA_SUBNET6_SELECT_NOT_FOR_US	1068
26.22.114	HA_SUBNET6_SELECT_NO_RELATIONSHIP_FOR_SUBNET	1068
26.22.115	HA_SUBNET6_SELECT_NO_RELATIONSHIP_SELECTOR_FOR_SUBNET	1068
26.22.116	HA_SUBNET6_SELECT_NO_SUBNET_SELECTED	1068
26.22.117	HA_SYNC_COMPLETE_NOTIFY_COMMUNICATIONS_FAILED	1069
26.22.118	HA_SYNC_COMPLETE_NOTIFY_FAILED	1069
26.22.119	HA_SYNC_COMPLETE_NOTIFY_HANDLER_FAILED	1069
26.22.120	HA_SYNC_FAILED	1069
26.22.121	HA_SYNC_HANDLER_FAILED	1069
26.22.122	HA_SYNC_START	1069
26.22.123	HA_SYNC_SUCCESSFUL	1069

26.22.12	HA_TERMINATED	1070
26.22.12	HA_TERMINATED_PARTNER_DID_NOT_RESTART	1070
26.22.12	HA_TERMINATED_RESTART_PARTNER	1070
26.23	HOOKS	1070
26.23.1	HOOKS_ALL_CALLOUTS_DEREGISTERED	1070
26.23.2	HOOKS_CALLOUTS_BEGIN	1070
26.23.3	HOOKS_CALLOUTS_COMPLETE	1070
26.23.4	HOOKS_CALLOUTS_REMOVED	1071
26.23.5	HOOKS_CALLOUT_CALLED	1071
26.23.6	HOOKS_CALLOUT_DEREGISTERED	1071
26.23.7	HOOKS_CALLOUT_ERROR	1071
26.23.8	HOOKS_CALLOUT_EXCEPTION	1071
26.23.9	HOOKS_CALLOUT_REGISTRATION	1071
26.23.10	HOOKS_CLOSE_ERROR	1072
26.23.11	HOOKS_HOOK_LIST_RESET	1072
26.23.12	HOOKS_INCORRECT_VERSION	1072
26.23.13	HOOKS_LIBRARY_CLOSED	1072
26.23.14	HOOKS_LIBRARY_LOADED	1072
26.23.15	HOOKS_LIBRARY_LOADING	1072
26.23.16	HOOKS_LIBRARY_MULTI_THREADING_COMPATIBLE	1072
26.23.17	HOOKS_LIBRARY_MULTI_THREADING_NOT_COMPATIBLE	1073
26.23.18	HOOKS_LIBRARY_UNLOADED	1073
26.23.19	HOOKS_LIBRARY_UNLOADING	1073
26.23.20	HOOKS_LIBRARY_VERSION	1073
26.23.21	HOOKS_LOAD_ERROR	1073
26.23.22	HOOKS_LOAD_EXCEPTION	1073
26.23.23	HOOKS_LOAD_FRAMEWORK_EXCEPTION	1073
26.23.24	HOOKS_LOAD_SUCCESS	1074
26.23.25	HOOKS_MULTI_THREADING_COMPATIBLE_EXCEPTION	1074
26.23.26	HOOKS_NO_LOAD	1074
26.23.27	HOOKS_NO_UNLOAD	1074
26.23.28	HOOKS_NO_VERSION	1074
26.23.29	HOOKS_OPEN_ERROR	1074
26.23.30	HOOKS_STD_CALLOUT_REGISTERED	1074
26.23.31	HOOKS_UNLOAD_ERROR	1075
26.23.32	HOOKS_UNLOAD_EXCEPTION	1075
26.23.33	HOOKS_UNLOAD_FRAMEWORK_EXCEPTION	1075
26.23.34	HOOKS_UNLOAD_SUCCESS	1075
26.23.35	HOOKS_VERSION_EXCEPTION	1075
26.24	HOSTS	1075
26.24.1	HOSTS_BACKENDS_REGISTERED	1075
26.24.2	HOSTS_BACKEND_DEREGISTER	1075
26.24.3	HOSTS_BACKEND_REGISTER	1076
26.24.4	HOSTS_CFG_ADD_HOST	1076
26.24.5	HOSTS_CFG_CACHE_HOST_DATA_SOURCE	1076
26.24.6	HOSTS_CFG_CLOSE_HOST_DATA_SOURCE	1076
26.24.7	HOSTS_CFG_DEL	1076
26.24.8	HOSTS_CFG_DEL4	1076
26.24.9	HOSTS_CFG_DEL6	1076
26.24.10	HOSTS_CFG_DEL_ALL_SUBNET4	1077
26.24.11	HOSTS_CFG_DEL_ALL_SUBNET6	1077
26.24.12	HOSTS_CFG_GET_ALL	1077
26.24.13	HOSTS_CFG_GET_ALL_ADDRESS4	1077
26.24.14	HOSTS_CFG_GET_ALL_ADDRESS4_COUNT	1077

26.24.15HOSTS_CFG_GET_ALL_ADDRESS4_HOST	1077
26.24.16HOSTS_CFG_GET_ALL_ADDRESS6	1077
26.24.17HOSTS_CFG_GET_ALL_ADDRESS6_COUNT	1078
26.24.18HOSTS_CFG_GET_ALL_ADDRESS6_HOST	1078
26.24.19HOSTS_CFG_GET_ALL_COUNT	1078
26.24.20HOSTS_CFG_GET_ALL_HOST	1078
26.24.21HOSTS_CFG_GET_ALL_HOSTNAME	1078
26.24.22HOSTS_CFG_GET_ALL_HOSTNAME_COUNT	1078
26.24.23HOSTS_CFG_GET_ALL_HOSTNAME_HOST	1078
26.24.24HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4	1079
26.24.25HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_COUNT	1079
26.24.26HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_HOST	1079
26.24.27HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6	1079
26.24.28HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_COUNT	1079
26.24.29HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_HOST	1079
26.24.30HOSTS_CFG_GET_ALL_IDENTIFIER	1079
26.24.31HOSTS_CFG_GET_ALL_IDENTIFIER_COUNT	1080
26.24.32HOSTS_CFG_GET_ALL_IDENTIFIER_HOST	1080
26.24.33HOSTS_CFG_GET_ALL_SUBNET_ID4	1080
26.24.34HOSTS_CFG_GET_ALL_SUBNET_ID4_COUNT	1080
26.24.35HOSTS_CFG_GET_ALL_SUBNET_ID4_HOST	1080
26.24.36HOSTS_CFG_GET_ALL_SUBNET_ID6	1080
26.24.37HOSTS_CFG_GET_ALL_SUBNET_ID6_COUNT	1080
26.24.38HOSTS_CFG_GET_ALL_SUBNET_ID6_HOST	1081
26.24.39HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4	1081
26.24.40HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_COUNT	1081
26.24.41HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_HOST	1081
26.24.42HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6	1081
26.24.43HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_COUNT	1081
26.24.44HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_HOST	1081
26.24.45HOSTS_CFG_GET_ONE_PREFIX	1082
26.24.46HOSTS_CFG_GET_ONE_PREFIX_HOST	1082
26.24.47HOSTS_CFG_GET_ONE_PREFIX_NULL	1082
26.24.48HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4	1082
26.24.49HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_HOST	1082
26.24.50HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_NULL	1082
26.24.51HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6	1082
26.24.52HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_HOST	1083
26.24.53HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_NULL	1083
26.24.54HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER	1083
26.24.55HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_HOST	1083
26.24.56HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_NULL	1083
26.24.57HOSTS_CFG_UPDATE_ADD	1083
26.24.58HOSTS_CFG_UPDATE_DEL4	1083
26.24.59HOSTS_CFG_UPDATE_DEL6	1084
26.24.60HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_ADDRESS4	1084
26.24.61HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER	1084
26.24.62HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_HOST	1084
26.24.63HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_NULL	1084
26.24.64HOSTS_MGR_ALTERNATE_GET6_PREFIX	1084
26.24.65HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_ADDRESS6	1084
26.24.66HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER	1085
26.24.67HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_HOST	1085
26.24.68HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_NULL	1085

26.24.69	HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS4	1085
26.24.70	HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS6	1085
26.24.71	HOSTS_MGR_NON_UNIQUE_IP_UNSUPPORTED	1085
26.25	HOST	1085
26.25.1	HOST_CACHE_ADD	1085
26.25.2	HOST_CACHE_ADD_DUPLICATE	1086
26.25.3	HOST_CACHE_COMMAND_CLEAR	1086
26.25.4	HOST_CACHE_COMMAND_CLEAR_FAILED	1086
26.25.5	HOST_CACHE_COMMAND_FLUSH	1086
26.25.6	HOST_CACHE_COMMAND_FLUSH_FAILED	1086
26.25.7	HOST_CACHE_COMMAND_GET	1086
26.25.8	HOST_CACHE_COMMAND_GET_BY_ID	1086
26.25.9	HOST_CACHE_COMMAND_GET_BY_ID_FAILED	1086
26.25.10	HOST_CACHE_COMMAND_GET_FAILED	1087
26.25.11	HOST_CACHE_COMMAND_INSERT	1087
26.25.12	HOST_CACHE_COMMAND_INSERT_FAILED	1087
26.25.13	HOST_CACHE_COMMAND_LOAD	1087
26.25.14	HOST_CACHE_COMMAND_LOAD_FAILED	1087
26.25.15	HOST_CACHE_COMMAND_REMOVE	1087
26.25.16	HOST_CACHE_COMMAND_REMOVE_FAILED	1087
26.25.17	HOST_CACHE_COMMAND_SIZE	1087
26.25.18	HOST_CACHE_COMMAND_SIZE_FAILED	1088
26.25.19	HOST_CACHE_COMMAND_WRITE	1088
26.25.20	HOST_CACHE_COMMAND_WRITE_FAILED	1088
26.25.21	HOST_CACHE_CONFIGURATION_FAILED	1088
26.25.22	HOST_CACHE_DEINIT_OK	1088
26.25.23	HOST_CACHE_DEL_SUBNET_ID_ADDRESS4	1088
26.25.24	HOST_CACHE_DEL_SUBNET_ID_ADDRESS6	1088
26.25.25	HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER4	1088
26.25.26	HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER6	1089
26.25.27	HOST_CACHE_GET_ONE_PREFIX	1089
26.25.28	HOST_CACHE_GET_ONE_PREFIX_HOST	1089
26.25.29	HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4	1089
26.25.30	HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4_HOST	1089
26.25.31	HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6	1089
26.25.32	HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6_HOST	1089
26.25.33	HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER	1090
26.25.34	HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER_HOST	1090
26.25.35	HOST_CACHE_INIT_OK	1090
26.25.36	HOST_CMDS_DEINIT_OK	1090
26.25.37	HOST_CMDS_INIT_FAILED	1090
26.25.38	HOST_CMDS_INIT_OK	1090
26.25.39	HOST_CMDS_RESERV_ADD	1090
26.25.40	HOST_CMDS_RESERV_ADD_FAILED	1090
26.25.41	HOST_CMDS_RESERV_ADD_SUCCESS	1091
26.25.42	HOST_CMDS_RESERV_DEL	1091
26.25.43	HOST_CMDS_RESERV_DEL_FAILED	1091
26.25.44	HOST_CMDS_RESERV_DEL_SUCCESS	1091
26.25.45	HOST_CMDS_RESERV_GET	1091
26.25.46	HOST_CMDS_RESERV_GET_ALL	1091
26.25.47	HOST_CMDS_RESERV_GET_ALL_FAILED	1091
26.25.48	HOST_CMDS_RESERV_GET_ALL_SUCCESS	1091
26.25.49	HOST_CMDS_RESERV_GET_BY_ADDRESS	1091
26.25.50	HOST_CMDS_RESERV_GET_BY_ADDRESS_FAILED	1092

26.25.51	HOST_CMDS_RESERV_GET_BY_ADDRESS_SUCCESS	1092
26.25.52	HOST_CMDS_RESERV_GET_BY_HOSTNAME	1092
26.25.53	HOST_CMDS_RESERV_GET_BY_HOSTNAME_FAILED	1092
26.25.54	HOST_CMDS_RESERV_GET_BY_HOSTNAME_SUCCESS	1092
26.25.55	HOST_CMDS_RESERV_GET_BY_ID	1092
26.25.56	HOST_CMDS_RESERV_GET_BY_ID_FAILED	1092
26.25.57	HOST_CMDS_RESERV_GET_BY_ID_SUCCESS	1092
26.25.58	HOST_CMDS_RESERV_GET_FAILED	1092
26.25.59	HOST_CMDS_RESERV_GET_PAGE	1093
26.25.60	HOST_CMDS_RESERV_GET_PAGE_FAILED	1093
26.25.61	HOST_CMDS_RESERV_GET_PAGE_SUCCESS	1093
26.25.62	HOST_CMDS_RESERV_GET_SUCCESS	1093
26.25.63	HOST_CMDS_RESERV_UPDATE	1093
26.25.64	HOST_CMDS_RESERV_UPDATE_FAILED	1093
26.25.65	HOST_CMDS_RESERV_UPDATE_SUCCESS	1093
26.26	HTTPS	1093
26.26.1	HTTPS_REQUEST_RECEIVE_START	1093
26.27	HTTP	1094
26.27.1	HTTP_BAD_CLIENT_REQUEST_RECEIVED	1094
26.27.2	HTTP_BAD_CLIENT_REQUEST_RECEIVED_DETAILS	1094
26.27.3	HTTP_BAD_SERVER_RESPONSE_RECEIVED	1094
26.27.4	HTTP_BAD_SERVER_RESPONSE_RECEIVED_DETAILS	1094
26.27.5	HTTP_CLIENT_MT_STARTED	1094
26.27.6	HTTP_CLIENT_QUEUE_SIZE_GROWING	1094
26.27.7	HTTP_CLIENT_REQUEST_AUTHORIZED	1095
26.27.8	HTTP_CLIENT_REQUEST_BAD_AUTH_HEADER	1095
26.27.9	HTTP_CLIENT_REQUEST_NOT_AUTHORIZED	1095
26.27.10	HTTP_CLIENT_REQUEST_NO_AUTH_HEADER	1095
26.27.11	HTTP_CLIENT_REQUEST_RECEIVED	1095
26.27.12	HTTP_CLIENT_REQUEST_RECEIVED_DETAILS	1095
26.27.13	HTTP_CLIENT_REQUEST_SEND	1095
26.27.14	HTTP_CLIENT_REQUEST_SEND_DETAILS	1096
26.27.15	HTTP_CLIENT_REQUEST_TIMEOUT_OCCURRED	1096
26.27.16	HTTP_COMMAND_MGR_HTTPS_SERVICE_REUSE_FAILED	1096
26.27.17	HTTP_COMMAND_MGR_HTTPS_SERVICE_UPDATED	1096
26.27.18	HTTP_COMMAND_MGR_HTTP_SERVICE_REUSE_FAILED	1096
26.27.19	HTTP_COMMAND_MGR_HTTP_SERVICE_UPDATED	1096
26.27.20	HTTP_COMMAND_MGR_SERVICE_STARTED	1096
26.27.21	HTTP_COMMAND_MGR_SERVICE_STOPPING	1097
26.27.22	HTTP_CONNECTION_CLOSE_CALLBACK_FAILED	1097
26.27.23	HTTP_CONNECTION_HANDSHAKE_FAILED	1097
26.27.24	HTTP_CONNECTION_HANDSHAKE_START	1097
26.27.25	HTTP_CONNECTION_SHUTDOWN	1097
26.27.26	HTTP_CONNECTION_SHUTDOWN_FAILED	1097
26.27.27	HTTP_CONNECTION_STOP	1097
26.27.28	HTTP_CONNECTION_STOP_FAILED	1098
26.27.29	HTTP_CONNECTION_WATCH_SOCKET_CLEAR_ERROR	1098
26.27.30	HTTP_CONNECTION_WATCH_SOCKET_CLOSE_ERROR	1098
26.27.31	HTTP_CONNECTION_WATCH_SOCKET_MARK_READY_ERROR	1098
26.27.32	HTTP_DATA_RECEIVED	1098
26.27.33	HTTP_IDLE_CONNECTION_TIMEOUT_OCCURRED	1098
26.27.34	HTTP_PREMATURE_CONNECTION_TIMEOUT_OCCURRED	1098
26.27.35	HTTP_REQUEST_RECEIVE_START	1099
26.27.36	HTTP_SERVER_RESPONSE_RECEIVED	1099

26.27.37	HTTP_SERVER_RESPONSE_RECEIVED_DETAILS	1099
26.27.38	HTTP_SERVER_RESPONSE_SEND	1099
26.27.39	HTTP_SERVER_RESPONSE_SEND_DETAILS	1099
26.28	KEY	1099
26.28.1	KEY_LOOKUP_DISABLED	1099
26.28.2	KEY_LOOKUP_FOUND	1100
26.28.3	KEY_LOOKUP_NONE	1100
26.28.4	KEY_PROCESSING_FAILED	1100
26.28.5	KEY_PROCESSING_FAILED_UNSPECIFIED_ERROR	1100
26.29	LEASE	1100
26.29.1	LEASE_CMDS_ADD4	1100
26.29.2	LEASE_CMDS_ADD4_CONFLICT	1100
26.29.3	LEASE_CMDS_ADD4_FAILED	1100
26.29.4	LEASE_CMDS_ADD6	1101
26.29.5	LEASE_CMDS_ADD6_CONFLICT	1101
26.29.6	LEASE_CMDS_ADD6_FAILED	1101
26.29.7	LEASE_CMDS_BULK_APPLY6	1101
26.29.8	LEASE_CMDS_BULK_APPLY6_FAILED	1101
26.29.9	LEASE_CMDS_DEINIT_OK	1101
26.29.10	LEASE_CMDS_DEL4	1101
26.29.11	LEASE_CMDS_DEL4_FAILED	1101
26.29.12	LEASE_CMDS_DEL6	1102
26.29.13	LEASE_CMDS_DEL6_FAILED	1102
26.29.14	LEASE_CMDS_GET4_FAILED	1102
26.29.15	LEASE_CMDS_GET6_FAILED	1102
26.29.16	LEASE_CMDS_INIT_OK	1102
26.29.17	LEASE_CMDS_LEASE4_OFFER_FAILED	1102
26.29.18	LEASE_CMDS_LEASES4_COMMITTED_FAILED	1102
26.29.19	LEASE_CMDS_LEASES6_COMMITTED_CONFLICT	1102
26.29.20	LEASE_CMDS_LEASES6_COMMITTED_FAILED	1103
26.29.21	LEASE_CMDS_LEASES6_COMMITTED_LEASE_ERROR	1103
26.29.22	LEASE_CMDS_LOAD_ERROR	1103
26.29.23	LEASE_CMDS_RESEND_DDNS4	1103
26.29.24	LEASE_CMDS_RESEND_DDNS4_FAILED	1103
26.29.25	LEASE_CMDS_RESEND_DDNS6	1103
26.29.26	LEASE_CMDS_RESEND_DDNS6_FAILED	1103
26.29.27	LEASE_CMDS_UPDATE4	1103
26.29.28	LEASE_CMDS_UPDATE4_CONFLICT	1104
26.29.29	LEASE_CMDS_UPDATE4_FAILED	1104
26.29.30	LEASE_CMDS_UPDATE6	1104
26.29.31	LEASE_CMDS_UPDATE6_CONFLICT	1104
26.29.32	LEASE_CMDS_UPDATE6_FAILED	1104
26.29.33	LEASE_CMDS_WIPE4	1104
26.29.34	LEASE_CMDS_WIPE4_FAILED	1104
26.29.35	LEASE_CMDS_WIPE6	1104
26.29.36	LEASE_CMDS_WIPE6_FAILED	1105
26.29.37	LEASE_QUERY_LOAD_FAILED	1105
26.29.38	LEASE_QUERY_LOAD_OK	1105
26.29.39	LEASE_QUERY_UNLOAD_OK	1105
26.30	LEGAL	1105
26.30.1	LEGAL_LOG_COMMAND_NO_LEGAL_STORE	1105
26.30.2	LEGAL_LOG_COMMAND_WRITE_ERROR	1105
26.30.3	LEGAL_LOG_DB_OPEN_CONNECTION_WITH_RETRY_FAILED	1105
26.30.4	LEGAL_LOG_LEASE4_NO_LEGAL_STORE	1106

26.30.5	LEGAL_LOG_LEASE4_WRITE_ERROR	1106
26.30.6	LEGAL_LOG_LEASE6_NO_LEGAL_STORE	1106
26.30.7	LEGAL_LOG_LEASE6_WRITE_ERROR	1106
26.30.8	LEGAL_LOG_LOAD_ERROR	1106
26.30.9	LEGAL_LOG_MYSQL_COMMIT	1106
26.30.10	LEGAL_LOG_MYSQL_DB_RECONNECT_ATTEMPT_FAILED	1106
26.30.11	LEGAL_LOG_MYSQL_DB_RECONNECT_ATTEMPT_SCHEDULE	1106
26.30.12	LEGAL_LOG_MYSQL_DB_RECONNECT_FAILED	1107
26.30.13	LEGAL_LOG_MYSQL_FATAL_ERROR	1107
26.30.14	LEGAL_LOG_MYSQL_GET_VERSION	1107
26.30.15	LEGAL_LOG_MYSQL_INSERT_LOG	1107
26.30.16	LEGAL_LOG_MYSQL_INVALID_ACCESS	1107
26.30.17	LEGAL_LOG_MYSQL_NO_TLS	1107
26.30.18	LEGAL_LOG_MYSQL_ROLLBACK	1107
26.30.19	LEGAL_LOG_MYSQL_START_TRANSACTION	1107
26.30.20	LEGAL_LOG_MYSQL_TLS_CIPHER	1108
26.30.21	LEGAL_LOG_PGSQL_COMMIT	1108
26.30.22	LEGAL_LOG_PGSQL_DB_RECONNECT_ATTEMPT_FAILED	1108
26.30.23	LEGAL_LOG_PGSQL_DB_RECONNECT_ATTEMPT_SCHEDULE	1108
26.30.24	LEGAL_LOG_PGSQL_DB_RECONNECT_FAILED	1108
26.30.25	LEGAL_LOG_PGSQL_DEALLOC_ERROR	1108
26.30.26	LEGAL_LOG_PGSQL_FATAL_ERROR	1108
26.30.27	LEGAL_LOG_PGSQL_GET_VERSION	1109
26.30.28	LEGAL_LOG_PGSQL_INSERT_LOG	1109
26.30.29	LEGAL_LOG_PGSQL_INVALID_ACCESS	1109
26.30.30	LEGAL_LOG_PGSQL_ROLLBACK	1109
26.30.31	LEGAL_LOG_PGSQL_START_TRANSACTION	1109
26.30.32	LEGAL_LOG_STORE_CLOSED	1109
26.30.33	LEGAL_LOG_STORE_CLOSE_ERROR	1109
26.30.34	LEGAL_LOG_STORE_OPEN	1109
26.30.35	LEGAL_LOG_STORE_OPENED	1110
26.30.36	LEGAL_LOG_SYSLOG	1110
26.30.37	LEGAL_LOG_SYSLOG_STORE_OPEN	1110
26.30.38	LEGAL_LOG_UNLOAD_ERROR	1110
26.31	LFC	1110
26.31.1	LFC_FAIL_PID_CREATE	1110
26.31.2	LFC_FAIL_PID_DEL	1110
26.31.3	LFC_FAIL_PROCESS	1110
26.31.4	LFC_FAIL_ROTATE	1110
26.31.5	LFC_PROCESSING	1111
26.31.6	LFC_READ_STATS	1111
26.31.7	LFC_ROTATING	1111
26.31.8	LFC_RUNNING	1111
26.31.9	LFC_START	1111
26.31.10	LFC_TERMINATE	1111
26.31.11	LFC_WRITE_STATS	1111
26.32	LIMITS	1111
26.32.1	LIMITS_CONFIGURATION_LEASE_BACKEND_NOT_AVAILABLE	1111
26.32.2	LIMITS_CONFIGURATION_LEASE_BACKEND_SHOULD_HAVE_BEEN_AVAILABLE	1112
26.32.3	LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_CLIENT_CLASS	1112
26.32.4	LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_SUBNET	1112
26.32.5	LIMITS_CONFIGURED_PREFIX_LIMIT_BY_CLIENT_CLASS	1112
26.32.6	LIMITS_CONFIGURED_PREFIX_LIMIT_BY_SUBNET	1112
26.32.7	LIMITS_CONFIGURED_RATE_LIMIT_BY_CLIENT_CLASS	1112

26.32.8	LIMITS_CONFIGURED_RATE_LIMIT_BY_SUBNET	1112
26.32.9	LIMITS_LEASE_LIMIT_EXCEEDED	1113
26.32.10	LIMITS_LEASE_WITHIN_LIMITS	1113
26.32.11	LIMITS_PACKET_WIIH_SUBNET_ID_RATE_NO_SUBNET	1113
26.32.12	LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_DROPPED	1113
26.32.13	LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_HONORED	1113
26.32.14	LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_DROPPED	1113
26.32.15	LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_HONORED	1113
26.33	LOGIMPL	1114
26.33.1	LOGIMPL_ABOVE_MAX_DEBUG	1114
26.33.2	LOGIMPL_BAD_DEBUG_STRING	1114
26.33.3	LOGIMPL_BELOW_MIN_DEBUG	1114
26.34	LOG	1114
26.34.1	LOG_BAD_DESTINATION	1114
26.34.2	LOG_BAD_SEVERITY	1114
26.34.3	LOG_BAD_STREAM	1114
26.34.4	LOG_DUPLICATE_MESSAGE_ID	1115
26.34.5	LOG_DUPLICATE_NAMESPACE	1115
26.34.6	LOG_INPUT_OPEN_FAIL	1115
26.34.7	LOG_INVALID_MESSAGE_ID	1115
26.34.8	LOG_LOCK_TEST_MESSAGE	1115
26.34.9	LOG_NAMESPACE_EXTRA_ARGS	1115
26.34.10	LOG_NAMESPACE_INVALID_ARG	1115
26.34.11	LOG_NAMESPACE_NO_ARGS	1116
26.34.12	LOG_NO_MESSAGE_ID	1116
26.34.13	LOG_NO_MESSAGE_TEXT	1116
26.34.14	LOG_NO_SUCH_MESSAGE	1116
26.34.15	LOG_OPEN_OUTPUT_FAIL	1116
26.34.16	LOG_PREFIX_EXTRA_ARGS	1116
26.34.17	LOG_PREFIX_INVALID_ARG	1117
26.34.18	LOG_READING_LOCAL_FILE	1117
26.34.19	LOG_READ_ERROR	1117
26.34.20	LOG_UNRECOGNIZED_DIRECTIVE	1117
26.34.21	LOG_WRITE_ERROR	1117
26.35	MT	1117
26.35.1	MT_TCP_LISTENER_MGR_STARTED	1117
26.35.2	MT_TCP_LISTENER_MGR_STOPPED	1117
26.35.3	MT_TCP_LISTENER_MGR_STOPPING	1118
26.36	MYSQL	1118
26.36.1	MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4	1118
26.36.2	MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6	1118
26.36.3	MYSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6	1118
26.36.4	MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4	1118
26.36.5	MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6	1118
26.36.6	MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS4	1118
26.36.7	MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS6	1118
26.36.8	MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4	1119
26.36.9	MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6	1119
26.36.10	MYSQL_CB_CREATE_UPDATE_OPTION4	1119
26.36.11	MYSQL_CB_CREATE_UPDATE_OPTION6	1119
26.36.12	MYSQL_CB_CREATE_UPDATE_OPTION_DEF4	1119
26.36.13	MYSQL_CB_CREATE_UPDATE_OPTION_DEF6	1119
26.36.14	MYSQL_CB_CREATE_UPDATE_SERVER4	1119
26.36.15	MYSQL_CB_CREATE_UPDATE_SERVER6	1119

26.36.16	MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK4	1120
26.36.17	MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK6	1120
26.36.18	MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4	1120
26.36.19	MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6	1120
26.36.20	MYSQL_CB_CREATE_UPDATE_SUBNET4	1120
26.36.21	MYSQL_CB_CREATE_UPDATE_SUBNET6	1120
26.36.22	MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4	1120
26.36.23	MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT	1120
26.36.24	MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6	1121
26.36.25	MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT	1121
26.36.26	MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4	1121
26.36.27	MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT	1121
26.36.28	MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6	1121
26.36.29	MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT	1121
26.36.30	MYSQL_CB_DELETE_ALL_OPTION_DEFS4	1121
26.36.31	MYSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT	1121
26.36.32	MYSQL_CB_DELETE_ALL_OPTION_DEFS6	1121
26.36.33	MYSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT	1122
26.36.34	MYSQL_CB_DELETE_ALL_SERVERS4	1122
26.36.35	MYSQL_CB_DELETE_ALL_SERVERS4_RESULT	1122
26.36.36	MYSQL_CB_DELETE_ALL_SERVERS6	1122
26.36.37	MYSQL_CB_DELETE_ALL_SERVERS6_RESULT	1122
26.36.38	MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4	1122
26.36.39	MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT	1122
26.36.40	MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6	1122
26.36.41	MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT	1122
26.36.42	MYSQL_CB_DELETE_ALL_SUBNETS4	1123
26.36.43	MYSQL_CB_DELETE_ALL_SUBNETS4_RESULT	1123
26.36.44	MYSQL_CB_DELETE_ALL_SUBNETS6	1123
26.36.45	MYSQL_CB_DELETE_ALL_SUBNETS6_RESULT	1123
26.36.46	MYSQL_CB_DELETE_BY_POOL_OPTION4	1123
26.36.47	MYSQL_CB_DELETE_BY_POOL_OPTION4_RESULT	1123
26.36.48	MYSQL_CB_DELETE_BY_POOL_OPTION6	1123
26.36.49	MYSQL_CB_DELETE_BY_POOL_OPTION6_RESULT	1123
26.36.50	MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6	1123
26.36.51	MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT	1124
26.36.52	MYSQL_CB_DELETE_BY_PREFIX_SUBNET4	1124
26.36.53	MYSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT	1124
26.36.54	MYSQL_CB_DELETE_BY_PREFIX_SUBNET6	1124
26.36.55	MYSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT	1124
26.36.56	MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4	1124
26.36.57	MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT	1124
26.36.58	MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6	1124
26.36.59	MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT	1124
26.36.60	MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4	1125
26.36.61	MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT	1125
26.36.62	MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6	1125
26.36.63	MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT	1125
26.36.64	MYSQL_CB_DELETE_CLIENT_CLASS4	1125
26.36.65	MYSQL_CB_DELETE_CLIENT_CLASS4_RESULT	1125
26.36.66	MYSQL_CB_DELETE_CLIENT_CLASS6	1125
26.36.67	MYSQL_CB_DELETE_CLIENT_CLASS6_RESULT	1125
26.36.68	MYSQL_CB_DELETE_GLOBAL_PARAMETER4	1125
26.36.69	MYSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT	1126

26.36.70	MYSQL_CB_DELETE_GLOBAL_PARAMETER6	1126
26.36.71	MYSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT	1126
26.36.72	MYSQL_CB_DELETE_OPTION4	1126
26.36.73	MYSQL_CB_DELETE_OPTION4_RESULT	1126
26.36.74	MYSQL_CB_DELETE_OPTION6	1126
26.36.75	MYSQL_CB_DELETE_OPTION6_RESULT	1126
26.36.76	MYSQL_CB_DELETE_OPTION_DEF4	1126
26.36.77	MYSQL_CB_DELETE_OPTION_DEF4_RESULT	1126
26.36.78	MYSQL_CB_DELETE_OPTION_DEF6	1127
26.36.79	MYSQL_CB_DELETE_OPTION_DEF6_RESULT	1127
26.36.80	MYSQL_CB_DELETE_SERVER4	1127
26.36.81	MYSQL_CB_DELETE_SERVER4_RESULT	1127
26.36.82	MYSQL_CB_DELETE_SERVER6	1127
26.36.83	MYSQL_CB_DELETE_SERVER6_RESULT	1127
26.36.84	MYSQL_CB_DELETE_SHARED_NETWORK4	1127
26.36.85	MYSQL_CB_DELETE_SHARED_NETWORK4_RESULT	1127
26.36.86	MYSQL_CB_DELETE_SHARED_NETWORK6	1127
26.36.87	MYSQL_CB_DELETE_SHARED_NETWORK6_RESULT	1128
26.36.88	MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4	1128
26.36.89	MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT	1128
26.36.90	MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6	1128
26.36.91	MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT	1128
26.36.92	MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4	1128
26.36.93	MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT	1128
26.36.94	MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6	1128
26.36.95	MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT	1128
26.36.96	MYSQL_CB_GET_ALL_CLIENT_CLASSES4	1129
26.36.97	MYSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT	1129
26.36.98	MYSQL_CB_GET_ALL_CLIENT_CLASSES6	1129
26.36.99	MYSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT	1129
26.36.100	MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4	1129
26.36.101	MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT	1129
26.36.102	MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6	1129
26.36.103	MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT	1129
26.36.104	MYSQL_CB_GET_ALL_OPTIONS4	1129
26.36.105	MYSQL_CB_GET_ALL_OPTIONS4_RESULT	1130
26.36.106	MYSQL_CB_GET_ALL_OPTIONS6	1130
26.36.107	MYSQL_CB_GET_ALL_OPTIONS6_RESULT	1130
26.36.108	MYSQL_CB_GET_ALL_OPTION_DEFS4	1130
26.36.109	MYSQL_CB_GET_ALL_OPTION_DEFS4_RESULT	1130
26.36.110	MYSQL_CB_GET_ALL_OPTION_DEFS6	1130
26.36.111	MYSQL_CB_GET_ALL_OPTION_DEFS6_RESULT	1130
26.36.112	MYSQL_CB_GET_ALL_SERVERS4	1130
26.36.113	MYSQL_CB_GET_ALL_SERVERS4_RESULT	1130
26.36.114	MYSQL_CB_GET_ALL_SERVERS6	1131
26.36.115	MYSQL_CB_GET_ALL_SERVERS6_RESULT	1131
26.36.116	MYSQL_CB_GET_ALL_SHARED_NETWORKS4	1131
26.36.117	MYSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT	1131
26.36.118	MYSQL_CB_GET_ALL_SHARED_NETWORKS6	1131
26.36.119	MYSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT	1131
26.36.120	MYSQL_CB_GET_ALL_SUBNETS4	1131
26.36.121	MYSQL_CB_GET_ALL_SUBNETS4_RESULT	1131
26.36.122	MYSQL_CB_GET_ALL_SUBNETS6	1131
26.36.123	MYSQL_CB_GET_ALL_SUBNETS6_RESULT	1132

26.36.12	MYSQL_CB_GET_CLIENT_CLASS4	1132
26.36.12	MYSQL_CB_GET_CLIENT_CLASS6	1132
26.36.12	MYSQL_CB_GET_GLOBAL_PARAMETER4	1132
26.36.12	MYSQL_CB_GET_GLOBAL_PARAMETER6	1132
26.36.12	MYSQL_CB_GET_HOST4	1132
26.36.12	MYSQL_CB_GET_HOST6	1132
26.36.13	MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4	1132
26.36.13	MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT	1133
26.36.13	MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6	1133
26.36.13	MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT	1133
26.36.13	MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4	1133
26.36.13	MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT	1133
26.36.13	MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6	1133
26.36.13	MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT	1133
26.36.13	MYSQL_CB_GET_MODIFIED_OPTIONS4	1134
26.36.13	MYSQL_CB_GET_MODIFIED_OPTIONS4_RESULT	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTIONS6	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTIONS6_RESULT	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTION_DEFS4	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTION_DEFS6	1134
26.36.14	MYSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT	1135
26.36.14	MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4	1135
26.36.14	MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT	1135
26.36.14	MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6	1135
26.36.14	MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT	1135
26.36.15	MYSQL_CB_GET_MODIFIED_SUBNETS4	1135
26.36.15	MYSQL_CB_GET_MODIFIED_SUBNETS4_RESULT	1135
26.36.15	MYSQL_CB_GET_MODIFIED_SUBNETS6	1136
26.36.15	MYSQL_CB_GET_MODIFIED_SUBNETS6_RESULT	1136
26.36.15	MYSQL_CB_GET_OPTION4	1136
26.36.15	MYSQL_CB_GET_OPTION6	1136
26.36.15	MYSQL_CB_GET_OPTION_DEF4	1136
26.36.15	MYSQL_CB_GET_OPTION_DEF6	1136
26.36.15	MYSQL_CB_GET_PORT4	1136
26.36.15	MYSQL_CB_GET_PORT6	1136
26.36.16	MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4	1137
26.36.16	MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT	1137
26.36.16	MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6	1137
26.36.16	MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT	1137
26.36.16	MYSQL_CB_GET_SERVER4	1137
26.36.16	MYSQL_CB_GET_SERVER6	1137
26.36.16	MYSQL_CB_GET_SHARED_NETWORK4	1137
26.36.16	MYSQL_CB_GET_SHARED_NETWORK6	1137
26.36.16	MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4	1138
26.36.16	MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT	1138
26.36.17	MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6	1138
26.36.17	MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT	1138
26.36.17	MYSQL_CB_GET_SUBNET4_BY_PREFIX	1138
26.36.17	MYSQL_CB_GET_SUBNET4_BY_SUBNET_ID	1138
26.36.17	MYSQL_CB_GET_SUBNET6_BY_PREFIX	1138
26.36.17	MYSQL_CB_GET_SUBNET6_BY_SUBNET_ID	1138
26.36.17	MYSQL_CB_GET_TYPE4	1138
26.36.17	MYSQL_CB_GET_TYPE6	1139

26.36.17	MYSQL_CB_NO_TLS	1139
26.36.17	MYSQL_CB_RECONNECT_ATTEMPT_FAILED4	1139
26.36.18	MYSQL_CB_RECONNECT_ATTEMPT_FAILED6	1139
26.36.18	MYSQL_CB_RECONNECT_ATTEMPT_SCHEDULE4	1139
26.36.18	MYSQL_CB_RECONNECT_ATTEMPT_SCHEDULE6	1139
26.36.18	MYSQL_CB_RECONNECT_FAILED4	1139
26.36.18	MYSQL_CB_RECONNECT_FAILED6	1139
26.36.18	MYSQL_CB_REGISTER_BACKEND_TYPE4	1140
26.36.18	MYSQL_CB_REGISTER_BACKEND_TYPE6	1140
26.36.18	MYSQL_CB_TLS_CIPHER	1140
26.36.18	MYSQL_CB_UNREGISTER_BACKEND_TYPE4	1140
26.36.18	MYSQL_CB_UNREGISTER_BACKEND_TYPE6	1140
26.36.19	MYSQL_DEINIT_OK	1140
26.36.19	MYSQL_FB_DB	1140
26.36.19	MYSQL_HB_DB	1140
26.36.19	MYSQL_HB_DB_GET_VERSION	1141
26.36.19	MYSQL_HB_DB_READONLY	1141
26.36.19	MYSQL_HB_DB_RECONNECT_ATTEMPT_FAILED	1141
26.36.19	MYSQL_HB_DB_RECONNECT_ATTEMPT_SCHEDULE	1141
26.36.19	MYSQL_HB_DB_RECONNECT_FAILED	1141
26.36.19	MYSQL_HB_NO_TLS	1141
26.36.19	MYSQL_HB_TLS_CIPHER	1141
26.36.20	MYSQL_INIT_OK	1141
26.36.20	MYSQL_LB_ADD_ADDR4	1142
26.36.20	MYSQL_LB_ADD_ADDR6	1142
26.36.20	MYSQL_LB_COMMIT	1142
26.36.20	MYSQL_LB_DB	1142
26.36.20	MYSQL_LB_DB_RECONNECT_ATTEMPT_FAILED	1142
26.36.20	MYSQL_LB_DB_RECONNECT_ATTEMPT_SCHEDULE	1142
26.36.20	MYSQL_LB_DB_RECONNECT_FAILED	1142
26.36.20	MYSQL_LB_DELETED_EXPIRED_RECLAIMED	1143
26.36.20	MYSQL_LB_DELETED_SUBNET4_ID	1143
26.36.21	MYSQL_LB_DELETED_SUBNET6_ID	1143
26.36.21	MYSQL_LB_DELETE_ADDR4	1143
26.36.21	MYSQL_LB_DELETE_ADDR6	1143
26.36.21	MYSQL_LB_DELETE_EXPIRED_RECLAIMED4	1143
26.36.21	MYSQL_LB_DELETE_EXPIRED_RECLAIMED6	1143
26.36.21	MYSQL_LB_GET4	1144
26.36.21	MYSQL_LB_GET6	1144
26.36.21	MYSQL_LB_GET_ADDR4	1144
26.36.21	MYSQL_LB_GET_ADDR6	1144
26.36.21	MYSQL_LB_GET_CLIENTID	1144
26.36.22	MYSQL_LB_GET_DUID	1144
26.36.22	MYSQL_LB_GET_EXPIRED4	1144
26.36.22	MYSQL_LB_GET_EXPIRED6	1145
26.36.22	MYSQL_LB_GET_HOSTNAME4	1145
26.36.22	MYSQL_LB_GET_HOSTNAME6	1145
26.36.22	MYSQL_LB_GET_HWADDR	1145
26.36.22	MYSQL_LB_GET_IAID_DUID	1145
26.36.22	MYSQL_LB_GET_IAID_SUBID_DUID	1145
26.36.22	MYSQL_LB_GET_PAGE4	1145
26.36.22	MYSQL_LB_GET_PAGE6	1146
26.36.23	MYSQL_LB_GET_RELAYID4	1146
26.36.23	MYSQL_LB_GET_RELAYID6	1146

26.36.23	MYSQL_LB_GET_REMOTEID4	1146
26.36.23	MYSQL_LB_GET_REMOTEID6	1146
26.36.23	MYSQL_LB_GET_SUBID4	1146
26.36.23	MYSQL_LB_GET_SUBID6	1146
26.36.23	MYSQL_LB_GET_SUBID_CLIENTID	1147
26.36.23	MYSQL_LB_GET_SUBID_HWADDR	1147
26.36.23	MYSQL_LB_GET_SUBID_PAGE6	1147
26.36.23	MYSQL_LB_GET_VERSION	1147
26.36.24	MYSQL_LB_NEGATIVE_LEASES_STAT	1147
26.36.24	MYSQL_LB_NO_TLS	1147
26.36.24	MYSQL_LB_ROLLBACK	1147
26.36.24	MYSQL_LB_TLS_CIPHER	1148
26.36.24	MYSQL_LB_UPDATE_ADDR4	1148
26.36.24	MYSQL_LB_UPDATE_ADDR6	1148
26.36.24	MYSQL_LB_UPGRADE_EXTENDED_INFO4	1148
26.36.24	MYSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR	1148
26.36.24	MYSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE	1148
26.36.24	MYSQL_LB_UPGRADE_EXTENDED_INFO6	1148
26.36.25	MYSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR	1149
26.36.25	MYSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE	1149
26.37	NETCONF	1149
26.37.1	NETCONF_BOOT_UPDATE_COMPLETED	1149
26.37.2	NETCONF_CONFIG_CHANGED_DETAIL	1149
26.37.3	NETCONF_CONFIG_CHANGE_EVENT	1149
26.37.4	NETCONF_CONFIG_CHECK_FAIL	1149
26.37.5	NETCONF_CONFIG_FAIL	1149
26.37.6	NETCONF_CONFIG_SYNTAX_WARNING	1150
26.37.7	NETCONF_FAILED	1150
26.37.8	NETCONF_GET_CONFIG	1150
26.37.9	NETCONF_GET_CONFIG_FAILED	1150
26.37.10	NETCONF_GET_CONFIG_STARTED	1150
26.37.11	NETCONF_MODULE_CHANGE_INTERNAL_ERROR	1150
26.37.12	NETCONF_MODULE_MISSING_ERR	1150
26.37.13	NETCONF_MODULE_MISSING_WARN	1151
26.37.14	NETCONF_MODULE_REVISION_ERR	1151
26.37.15	NETCONF_MODULE_REVISION_WARN	1151
26.37.16	NETCONF_NOTIFICATION_INTERNAL_ERROR	1151
26.37.17	NETCONF_NOTIFICATION_RECEIVED	1151
26.37.18	NETCONF_NOT_SUBSCRIBED_TO_NOTIFICATIONS	1151
26.37.19	NETCONF_RUN_EXIT	1151
26.37.20	NETCONF_SET_CONFIG	1152
26.37.21	NETCONF_SET_CONFIG_FAILED	1152
26.37.22	NETCONF_SET_CONFIG_STARTED	1152
26.37.23	NETCONF_STARTED	1152
26.37.24	NETCONF_SUBSCRIBE_CONFIG	1152
26.37.25	NETCONF_SUBSCRIBE_CONFIG_FAILED	1152
26.37.26	NETCONF_SUBSCRIBE_NOTIFICATIONS	1152
26.37.27	NETCONF_UPDATE_CONFIG	1153
26.37.28	NETCONF_UPDATE_CONFIG_COMPLETED	1153
26.37.29	NETCONF_UPDATE_CONFIG_FAILED	1153
26.37.30	NETCONF_UPDATE_CONFIG_STARTED	1153
26.37.31	NETCONF_VALIDATE_CONFIG	1153
26.37.32	NETCONF_VALIDATE_CONFIG_COMPLETED	1153
26.37.33	NETCONF_VALIDATE_CONFIG_FAILED	1153

26.37.34	NETCONF_VALIDATE_CONFIG_REJECTED	1153
26.37.35	NETCONF_VALIDATE_CONFIG_STARTED	1154
26.38	PERFMON	1154
26.38.1	PERFMON_ALARM_CLEARED	1154
26.38.2	PERFMON_ALARM_TRIGGERED	1154
26.38.3	PERFMON_CMDS_CONTROL_ERROR	1154
26.38.4	PERFMON_CMDS_CONTROL_OK	1154
26.38.5	PERFMON_CMDS_GET_ALL_DURATIONS_ERROR	1154
26.38.6	PERFMON_CMDS_GET_ALL_DURATIONS_OK	1155
26.38.7	PERFMON_DEINIT_OK	1155
26.38.8	PERFMON_DHCP4_PKT_EVENTS	1155
26.38.9	PERFMON_DHCP4_PKT_PROCESS_ERROR	1155
26.38.10	PERFMON_DHCP4_SOCKET_RECEIVED_TIME_SUPPORT	1155
26.38.11	PERFMON_DHCP6_PKT_EVENTS	1155
26.38.12	PERFMON_DHCP6_PKT_PROCESS_ERROR	1155
26.38.13	PERFMON_DHCP6_SOCKET_RECEIVED_TIME_SUPPORT	1156
26.38.14	PERFMON_INIT_FAILED	1156
26.38.15	PERFMON_INIT_OK	1156
26.39	PGSQL	1156
26.39.1	PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4	1156
26.39.2	PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6	1156
26.39.3	PGSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6	1156
26.39.4	PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4	1156
26.39.5	PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6	1156
26.39.6	PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS4	1157
26.39.7	PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS6	1157
26.39.8	PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4	1157
26.39.9	PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6	1157
26.39.10	PGSQL_CB_CREATE_UPDATE_OPTION4	1157
26.39.11	PGSQL_CB_CREATE_UPDATE_OPTION6	1157
26.39.12	PGSQL_CB_CREATE_UPDATE_OPTION_DEF4	1157
26.39.13	PGSQL_CB_CREATE_UPDATE_OPTION_DEF6	1157
26.39.14	PGSQL_CB_CREATE_UPDATE_SERVER4	1157
26.39.15	PGSQL_CB_CREATE_UPDATE_SERVER6	1158
26.39.16	PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK4	1158
26.39.17	PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK6	1158
26.39.18	PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4	1158
26.39.19	PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6	1158
26.39.20	PGSQL_CB_CREATE_UPDATE_SUBNET4	1158
26.39.21	PGSQL_CB_CREATE_UPDATE_SUBNET6	1158
26.39.22	PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4	1158
26.39.23	PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT	1159
26.39.24	PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6	1159
26.39.25	PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT	1159
26.39.26	PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4	1159
26.39.27	PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT	1159
26.39.28	PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6	1159
26.39.29	PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT	1159
26.39.30	PGSQL_CB_DELETE_ALL_OPTION_DEFS4	1159
26.39.31	PGSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT	1159
26.39.32	PGSQL_CB_DELETE_ALL_OPTION_DEFS6	1160
26.39.33	PGSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT	1160
26.39.34	PGSQL_CB_DELETE_ALL_SERVERS4	1160
26.39.35	PGSQL_CB_DELETE_ALL_SERVERS4_RESULT	1160

26.39.36PGSQL_CB_DELETE_ALL_SERVERS6	1160
26.39.37PGSQL_CB_DELETE_ALL_SERVERS6_RESULT	1160
26.39.38PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4	1160
26.39.39PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT	1160
26.39.40PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6	1160
26.39.41PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT	1161
26.39.42PGSQL_CB_DELETE_ALL_SUBNETS4	1161
26.39.43PGSQL_CB_DELETE_ALL_SUBNETS4_RESULT	1161
26.39.44PGSQL_CB_DELETE_ALL_SUBNETS6	1161
26.39.45PGSQL_CB_DELETE_ALL_SUBNETS6_RESULT	1161
26.39.46PGSQL_CB_DELETE_BY_POOL_OPTION4	1161
26.39.47PGSQL_CB_DELETE_BY_POOL_OPTION4_RESULT	1161
26.39.48PGSQL_CB_DELETE_BY_POOL_OPTION6	1161
26.39.49PGSQL_CB_DELETE_BY_POOL_OPTION6_RESULT	1161
26.39.50PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6	1162
26.39.51PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT	1162
26.39.52PGSQL_CB_DELETE_BY_PREFIX_SUBNET4	1162
26.39.53PGSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT	1162
26.39.54PGSQL_CB_DELETE_BY_PREFIX_SUBNET6	1162
26.39.55PGSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT	1162
26.39.56PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4	1162
26.39.57PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT	1162
26.39.58PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6	1162
26.39.59PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT	1163
26.39.60PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4	1163
26.39.61PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT	1163
26.39.62PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6	1163
26.39.63PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT	1163
26.39.64PGSQL_CB_DELETE_CLIENT_CLASS4	1163
26.39.65PGSQL_CB_DELETE_CLIENT_CLASS4_RESULT	1163
26.39.66PGSQL_CB_DELETE_CLIENT_CLASS6	1163
26.39.67PGSQL_CB_DELETE_CLIENT_CLASS6_RESULT	1163
26.39.68PGSQL_CB_DELETE_GLOBAL_PARAMETER4	1164
26.39.69PGSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT	1164
26.39.70PGSQL_CB_DELETE_GLOBAL_PARAMETER6	1164
26.39.71PGSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT	1164
26.39.72PGSQL_CB_DELETE_OPTION4	1164
26.39.73PGSQL_CB_DELETE_OPTION4_RESULT	1164
26.39.74PGSQL_CB_DELETE_OPTION6	1164
26.39.75PGSQL_CB_DELETE_OPTION6_RESULT	1164
26.39.76PGSQL_CB_DELETE_OPTION_DEF4	1164
26.39.77PGSQL_CB_DELETE_OPTION_DEF4_RESULT	1165
26.39.78PGSQL_CB_DELETE_OPTION_DEF6	1165
26.39.79PGSQL_CB_DELETE_OPTION_DEF6_RESULT	1165
26.39.80PGSQL_CB_DELETE_SERVER4	1165
26.39.81PGSQL_CB_DELETE_SERVER4_RESULT	1165
26.39.82PGSQL_CB_DELETE_SERVER6	1165
26.39.83PGSQL_CB_DELETE_SERVER6_RESULT	1165
26.39.84PGSQL_CB_DELETE_SHARED_NETWORK4	1165
26.39.85PGSQL_CB_DELETE_SHARED_NETWORK4_RESULT	1165
26.39.86PGSQL_CB_DELETE_SHARED_NETWORK6	1166
26.39.87PGSQL_CB_DELETE_SHARED_NETWORK6_RESULT	1166
26.39.88PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4	1166
26.39.89PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT	1166

26.39.90PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6	1166
26.39.91PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT	1166
26.39.92PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4	1166
26.39.93PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT	1166
26.39.94PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6	1166
26.39.95PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT	1167
26.39.96PGSQL_CB_GET_ALL_CLIENT_CLASSES4	1167
26.39.97PGSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT	1167
26.39.98PGSQL_CB_GET_ALL_CLIENT_CLASSES6	1167
26.39.99PGSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT	1167
26.39.10PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4	1167
26.39.10PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT	1167
26.39.10PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6	1167
26.39.10PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT	1167
26.39.10PGSQL_CB_GET_ALL_OPTIONS4	1168
26.39.10PGSQL_CB_GET_ALL_OPTIONS4_RESULT	1168
26.39.10PGSQL_CB_GET_ALL_OPTIONS6	1168
26.39.10PGSQL_CB_GET_ALL_OPTIONS6_RESULT	1168
26.39.10PGSQL_CB_GET_ALL_OPTION_DEFS4	1168
26.39.10PGSQL_CB_GET_ALL_OPTION_DEFS4_RESULT	1168
26.39.11PGSQL_CB_GET_ALL_OPTION_DEFS6	1168
26.39.11PGSQL_CB_GET_ALL_OPTION_DEFS6_RESULT	1168
26.39.11PGSQL_CB_GET_ALL_SERVERS4	1168
26.39.11PGSQL_CB_GET_ALL_SERVERS4_RESULT	1169
26.39.11PGSQL_CB_GET_ALL_SERVERS6	1169
26.39.11PGSQL_CB_GET_ALL_SERVERS6_RESULT	1169
26.39.11PGSQL_CB_GET_ALL_SHARED_NETWORKS4	1169
26.39.11PGSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT	1169
26.39.11PGSQL_CB_GET_ALL_SHARED_NETWORKS6	1169
26.39.11PGSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT	1169
26.39.12PGSQL_CB_GET_ALL_SUBNETS4	1169
26.39.12PGSQL_CB_GET_ALL_SUBNETS4_RESULT	1169
26.39.12PGSQL_CB_GET_ALL_SUBNETS6	1170
26.39.12PGSQL_CB_GET_ALL_SUBNETS6_RESULT	1170
26.39.12PGSQL_CB_GET_CLIENT_CLASS4	1170
26.39.12PGSQL_CB_GET_CLIENT_CLASS6	1170
26.39.12PGSQL_CB_GET_GLOBAL_PARAMETER4	1170
26.39.12PGSQL_CB_GET_GLOBAL_PARAMETER6	1170
26.39.12PGSQL_CB_GET_HOST4	1170
26.39.12PGSQL_CB_GET_HOST6	1170
26.39.13PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4	1170
26.39.13PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT	1171
26.39.13PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6	1171
26.39.13PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT	1171
26.39.13PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4	1171
26.39.13PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT	1171
26.39.13PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6	1171
26.39.13PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT	1171
26.39.13PGSQL_CB_GET_MODIFIED_OPTIONS4	1172
26.39.13PGSQL_CB_GET_MODIFIED_OPTIONS4_RESULT	1172
26.39.14PGSQL_CB_GET_MODIFIED_OPTIONS6	1172
26.39.14PGSQL_CB_GET_MODIFIED_OPTIONS6_RESULT	1172
26.39.14PGSQL_CB_GET_MODIFIED_OPTION_DEFS4	1172
26.39.14PGSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT	1172

26.39.14PGSQL_CB_GET_MODIFIED_OPTION_DEFS6	1172
26.39.14PGSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT	1173
26.39.14PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4	1173
26.39.14PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT	1173
26.39.14PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6	1173
26.39.14PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT	1173
26.39.15PGSQL_CB_GET_MODIFIED_SUBNETS4	1173
26.39.15PGSQL_CB_GET_MODIFIED_SUBNETS4_RESULT	1173
26.39.15PGSQL_CB_GET_MODIFIED_SUBNETS6	1174
26.39.15PGSQL_CB_GET_MODIFIED_SUBNETS6_RESULT	1174
26.39.15PGSQL_CB_GET_OPTION4	1174
26.39.15PGSQL_CB_GET_OPTION6	1174
26.39.15PGSQL_CB_GET_OPTION_DEF4	1174
26.39.15PGSQL_CB_GET_OPTION_DEF6	1174
26.39.15PGSQL_CB_GET_PORT4	1174
26.39.15PGSQL_CB_GET_PORT6	1174
26.39.16PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4	1175
26.39.16PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT	1175
26.39.16PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6	1175
26.39.16PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT	1175
26.39.16PGSQL_CB_GET_SERVER4	1175
26.39.16PGSQL_CB_GET_SERVER6	1175
26.39.16PGSQL_CB_GET_SHARED_NETWORK4	1175
26.39.16PGSQL_CB_GET_SHARED_NETWORK6	1175
26.39.16PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4	1176
26.39.16PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT	1176
26.39.17PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6	1176
26.39.17PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT	1176
26.39.17PGSQL_CB_GET_SUBNET4_BY_PREFIX	1176
26.39.17PGSQL_CB_GET_SUBNET4_BY_SUBNET_ID	1176
26.39.17PGSQL_CB_GET_SUBNET6_BY_PREFIX	1176
26.39.17PGSQL_CB_GET_SUBNET6_BY_SUBNET_ID	1176
26.39.17PGSQL_CB_GET_TYPE4	1176
26.39.17PGSQL_CB_GET_TYPE6	1177
26.39.17PGSQL_CB_RECONNECT_ATTEMPT_FAILED4	1177
26.39.17PGSQL_CB_RECONNECT_ATTEMPT_FAILED6	1177
26.39.18PGSQL_CB_RECONNECT_ATTEMPT_SCHEDULE4	1177
26.39.18PGSQL_CB_RECONNECT_ATTEMPT_SCHEDULE6	1177
26.39.18PGSQL_CB_RECONNECT_FAILED4	1177
26.39.18PGSQL_CB_RECONNECT_FAILED6	1177
26.39.18PGSQL_CB_REGISTER_BACKEND_TYPE4	1177
26.39.18PGSQL_CB_REGISTER_BACKEND_TYPE6	1178
26.39.18PGSQL_CB_UNREGISTER_BACKEND_TYPE4	1178
26.39.18PGSQL_CB_UNREGISTER_BACKEND_TYPE6	1178
26.39.18PGSQL_DEINIT_OK	1178
26.39.18PGSQL_FB_DB	1178
26.39.19PGSQL_HB_DB	1178
26.39.19PGSQL_HB_DB_GET_VERSION	1178
26.39.19PGSQL_HB_DB_READONLY	1178
26.39.19PGSQL_HB_DB_RECONNECT_ATTEMPT_FAILED	1179
26.39.19PGSQL_HB_DB_RECONNECT_ATTEMPT_SCHEDULE	1179
26.39.19PGSQL_HB_DB_RECONNECT_FAILED	1179
26.39.19PGSQL_INIT_OK	1179
26.39.19PGSQL_LB_ADD_ADDR4	1179

26.39.19	PGSQL_LB_ADD_ADDR6	1179
26.39.19	PGSQL_LB_COMMIT	1179
26.39.20	PGSQL_LB_DB	1179
26.39.20	PGSQL_LB_DB_RECONNECT_ATTEMPT_FAILED	1180
26.39.20	PGSQL_LB_DB_RECONNECT_ATTEMPT_SCHEDULE	1180
26.39.20	PGSQL_LB_DB_RECONNECT_FAILED	1180
26.39.20	PGSQL_LB_DELETED_SUBNET4_ID	1180
26.39.20	PGSQL_LB_DELETED_SUBNET6_ID	1180
26.39.20	PGSQL_LB_DELETE_ADDR4	1180
26.39.20	PGSQL_LB_DELETE_ADDR6	1180
26.39.20	PGSQL_LB_DELETE_EXPIRED_RECLAIMED4	1180
26.39.20	PGSQL_LB_DELETE_EXPIRED_RECLAIMED6	1181
26.39.21	PGSQL_LB_GET4	1181
26.39.21	PGSQL_LB_GET6	1181
26.39.21	PGSQL_LB_GET_ADDR4	1181
26.39.21	PGSQL_LB_GET_ADDR6	1181
26.39.21	PGSQL_LB_GET_CLIENTID	1181
26.39.21	PGSQL_LB_GET_DUID	1181
26.39.21	PGSQL_LB_GET_EXPIRED4	1182
26.39.21	PGSQL_LB_GET_EXPIRED6	1182
26.39.21	PGSQL_LB_GET_HOSTNAME4	1182
26.39.21	PGSQL_LB_GET_HOSTNAME6	1182
26.39.22	PGSQL_LB_GET_HWADDR	1182
26.39.22	PGSQL_LB_GET_IAID_DUID	1182
26.39.22	PGSQL_LB_GET_IAID_SUBID_DUID	1182
26.39.22	PGSQL_LB_GET_PAGE4	1183
26.39.22	PGSQL_LB_GET_PAGE6	1183
26.39.22	PGSQL_LB_GET_RELAYID4	1183
26.39.22	PGSQL_LB_GET_RELAYID6	1183
26.39.22	PGSQL_LB_GET_REMOTEID4	1183
26.39.22	PGSQL_LB_GET_REMOTEID6	1183
26.39.22	PGSQL_LB_GET_SUBID4	1183
26.39.23	PGSQL_LB_GET_SUBID6	1184
26.39.23	PGSQL_LB_GET_SUBID_CLIENTID	1184
26.39.23	PGSQL_LB_GET_SUBID_HWADDR	1184
26.39.23	PGSQL_LB_GET_SUBID_PAGE6	1184
26.39.23	PGSQL_LB_GET_VERSION	1184
26.39.23	PGSQL_LB_NEGATIVE_LEASES_STAT	1184
26.39.23	PGSQL_LB_ROLLBACK	1184
26.39.23	PGSQL_LB_UPDATE_ADDR4	1185
26.39.23	PGSQL_LB_UPDATE_ADDR6	1185
26.39.23	PGSQL_LB_UPGRADE_EXTENDED_INFO4	1185
26.39.24	PGSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR	1185
26.39.24	PGSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE	1185
26.39.24	PGSQL_LB_UPGRADE_EXTENDED_INFO6	1185
26.39.24	PGSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR	1185
26.39.24	PGSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE	1186
26.40	PING	1186
26.40.1	PING_CHECK_CB4_UPDATE_FAILED	1186
26.40.2	PING_CHECK_CHANNEL_ECHO_REPLY_RECEIVED	1186
26.40.3	PING_CHECK_CHANNEL_ECHO_REQUEST_SENT	1186
26.40.4	PING_CHECK_CHANNEL_MALFORMED_PACKET_RECEIVED	1186
26.40.5	PING_CHECK_CHANNEL_NETWORK_WRITE_ERROR	1186
26.40.6	PING_CHECK_CHANNEL_SOCKET_CLOSED	1186

26.40.7	PING_CHECK_CHANNEL_SOCKET_CLOSE_ERROR	1187
26.40.8	PING_CHECK_CHANNEL_SOCKET_OPENED	1187
26.40.9	PING_CHECK_CHANNEL_SOCKET_READ_FAILED	1187
26.40.10	PING_CHECK_CHANNEL_SOCKET_WRITE_FAILED	1187
26.40.11	PING_CHECK_CHANNEL_STOP	1187
26.40.12	PING_CHECK_CHANNEL_WATCH_SOCKET_CLEAR_ERROR	1187
26.40.13	PING_CHECK_CHANNEL_WATCH_SOCKET_CLOSE_ERROR	1187
26.40.14	PING_CHECK_DHCP4_SRV_CONFIGURED_FAILED	1188
26.40.15	PING_CHECK_DUPLICATE_CHECK	1188
26.40.16	PING_CHECK_LEASE4_OFFER_FAILED	1188
26.40.17	PING_CHECK_LOAD_ERROR	1188
26.40.18	PING_CHECK_LOAD_OK	1188
26.40.19	PING_CHECK_MGR_CHANNEL_DOWN	1188
26.40.20	PING_CHECK_MGR_LEASE_FREE_TO_USE	1188
26.40.21	PING_CHECK_MGR_NEXT_ECHO_SCHEDULED	1188
26.40.22	PING_CHECK_MGR_RECEIVED_ECHO_REPLY	1189
26.40.23	PING_CHECK_MGR_RECEIVED_UNEXPECTED_ECHO_REPLY	1189
26.40.24	PING_CHECK_MGR_RECEIVED_UNEXPECTED_UNREACHABLE_MSG	1189
26.40.25	PING_CHECK_MGR_RECEIVED_UNREACHABLE_MSG	1189
26.40.26	PING_CHECK_MGR_REPLY_RECEIVED_ERROR	1189
26.40.27	PING_CHECK_MGR_REPLY_TIMEOUT_EXPIRED	1189
26.40.28	PING_CHECK_MGR_SEND_COMPLETED_ERROR	1189
26.40.29	PING_CHECK_MGR_STARTED	1190
26.40.30	PING_CHECK_MGR_STARTED_SINGLE_THREADED	1190
26.40.31	PING_CHECK_MGR_START_PING_CHECK	1190
26.40.32	PING_CHECK_MGR_STOPPED	1190
26.40.33	PING_CHECK_MGR_STOPPING	1190
26.40.34	PING_CHECK_MGR_SUBNET_CONFIG_FAILED	1190
26.40.35	PING_CHECK_NO_LEASE_OR_LEASE_REUSED	1190
26.40.36	PING_CHECK_PAUSE_FAILED	1191
26.40.37	PING_CHECK_PAUSE_ILLEGAL	1191
26.40.38	PING_CHECK_PAUSE_PERMISSIONS_FAILED	1191
26.40.39	PING_CHECK_RESUME_FAILED	1191
26.40.40	PING_CHECK_UNEXPECTED_READ_ERROR	1191
26.40.41	PING_CHECK_UNEXPECTED_WRITE_ERROR	1191
26.40.42	PING_CHECK_UNLOAD	1191
26.41	RADIUS	1192
26.41.1	RADIUS_ACCESS_BUILD_FAILED	1192
26.41.2	RADIUS_ACCESS_CACHE_GET	1192
26.41.3	RADIUS_ACCESS_CACHE_INSERT	1192
26.41.4	RADIUS_ACCESS_CONFLICT	1192
26.41.5	RADIUS_ACCESS_DROP_PARKED_QUERY	1192
26.41.6	RADIUS_ACCESS_ERROR	1192
26.41.7	RADIUS_ACCESS_GET_IDENTIFIER	1192
26.41.8	RADIUS_ACCESS_GET_IDENTIFIER_FAILED	1192
26.41.9	RADIUS_ACCESS_HOST_BACKEND_ERROR	1193
26.41.10	RADIUS_ACCESS_MAX_PENDING_REQUESTS	1193
26.41.11	RADIUS_ACCESS_NO_HOST_CACHE	1193
26.41.12	RADIUS_ACCESS_ORPHAN	1193
26.41.13	RADIUS_ACCESS_RESUME_PARKED_QUERY	1193
26.41.14	RADIUS_ACCESS_SUBNET_RESELECT	1193
26.41.15	RADIUS_ACCESS_TERMINATE_ERROR	1193
26.41.16	RADIUS_ACCOUNTING_ASYNC	1194
26.41.17	RADIUS_ACCOUNTING_ASYNC_FAILED	1194

26.41.18RADIUS_ACCOUNTING_ASYNC_SUCCEED	1194
26.41.19RADIUS_ACCOUNTING_ERROR	1194
26.41.20RADIUS_ACCOUNTING_HISTORY_UPDATE_FAILED	1194
26.41.21RADIUS_ACCOUNTING_NO_HISTORY	1194
26.41.22RADIUS_ACCOUNTING_SYNC	1194
26.41.23RADIUS_ACCOUNTING_SYNC_FAILED	1194
26.41.24RADIUS_ACCOUNTING_SYNC_SUCCEED	1195
26.41.25RADIUS_AUTHENTICATION_ASYNC	1195
26.41.26RADIUS_AUTHENTICATION_ASYNC_ACCEPTED	1195
26.41.27RADIUS_AUTHENTICATION_ASYNC_FAILED	1195
26.41.28RADIUS_AUTHENTICATION_ASYNC_REJECTED	1195
26.41.29RADIUS_AUTHENTICATION_SYNC	1195
26.41.30RADIUS_AUTHENTICATION_SYNC_ACCEPTED	1195
26.41.31RADIUS_AUTHENTICATION_SYNC_FAILED	1195
26.41.32RADIUS_AUTHENTICATION_SYNC_REJECTED	1196
26.41.33RADIUS_BACKEND_GET4	1196
26.41.34RADIUS_BACKEND_GET6	1196
26.41.35RADIUS_CLEANUP_EXCEPTION	1196
26.41.36RADIUS_CONFIGURATION_FAILED	1196
26.41.37RADIUS_DECODE_MESSAGE	1196
26.41.38RADIUS_DEINIT_OK	1196
26.41.39RADIUS_ENCODE_MESSAGE	1197
26.41.40RADIUS_EXCHANGE_FAILED	1197
26.41.41RADIUS_EXCHANGE_OPEN_FAILED	1197
26.41.42RADIUS_EXCHANGE_RECEIVED	1197
26.41.43RADIUS_EXCHANGE_RECEIVED_ACCESS_ACCEPT	1197
26.41.44RADIUS_EXCHANGE_RECEIVED_ACCESS_REJECT	1197
26.41.45RADIUS_EXCHANGE_RECEIVED_ACCOUNTING_RESPONSE	1197
26.41.46RADIUS_EXCHANGE_RECEIVED_BAD_RESPONSE	1198
26.41.47RADIUS_EXCHANGE_RECEIVED_MISMATCH	1198
26.41.48RADIUS_EXCHANGE_RECEIVED_RESPONSE	1198
26.41.49RADIUS_EXCHANGE_RECEIVED_UNEXPECTED	1198
26.41.50RADIUS_EXCHANGE_RECEIVE_FAILED	1198
26.41.51RADIUS_EXCHANGE_SEND_FAILED	1198
26.41.52RADIUS_EXCHANGE_SEND_NEW	1198
26.41.53RADIUS_EXCHANGE_SEND_RETRY	1199
26.41.54RADIUS_EXCHANGE_SENT	1199
26.41.55RADIUS_EXCHANGE_START	1199
26.41.56RADIUS_EXCHANGE_SYNC_RETURN	1199
26.41.57RADIUS_EXCHANGE_TERMINATE	1199
26.41.58RADIUS_EXCHANGE_TIMEOUT	1199
26.41.59RADIUS_HOOK_FAILED	1199
26.41.60RADIUS_INIT_OK	1199
26.41.61RADIUS_INTEGER_ATTRIBUTE_FROM_BYTES_FAILED	1200
26.41.62RADIUS_INTEGER_ATTRIBUTE_FROM_TEXT_FAILED	1200
26.41.63RADIUS_IPADDR_ATTRIBUTE_FROM_BYTES_FAILED	1200
26.41.64RADIUS_IPADDR_ATTRIBUTE_FROM_TEXT_FAILED	1200
26.41.65RADIUS_IPV6ADDR_ATTRIBUTE_FROM_BYTES_FAILED	1200
26.41.66RADIUS_IPV6ADDR_ATTRIBUTE_FROM_TEXT_FAILED	1200
26.41.67RADIUS_IPV6PREFIX_ATTRIBUTE_FROM_BYTES_FAILED	1200
26.41.68RADIUS_IPV6PREFIX_ATTRIBUTE_FROM_TEXT_FAILED	1201
26.41.69RADIUS_PAUSE_FAILED	1201
26.41.70RADIUS_PAUSE_ILLEGAL	1201
26.41.71RADIUS_PAUSE_PERMISSIONS_FAILED	1201

26.41.72	RADIUS_REPLY_MESSAGE_ATTRIBUTE	1201
26.41.73	RADIUS_RESUME_FAILED	1201
26.41.74	RADIUS_SERVER_CONFIGURED	1201
26.41.75	RADIUS_SESSION_HISTORY_APPEND_FAILED	1202
26.41.76	RADIUS_SESSION_HISTORY_LOADED	1202
26.41.77	RADIUS_SESSION_HISTORY_LOAD_FAILED	1202
26.41.78	RADIUS_SESSION_HISTORY_OPENED	1202
26.41.79	RADIUS_SESSION_HISTORY_OPEN_FAILED	1202
26.41.80	RADIUS_SESSION_HISTORY_STORED	1202
26.41.81	RADIUS_SESSION_HISTORY_STORE_FAILED	1202
26.41.82	RADIUS_THREAD_POOL_STARTED	1203
26.42	RBAC	1203
26.42.1	RBAC_CONFIGURED_ACLS	1203
26.42.2	RBAC_CONFIGURED_COMMANDS	1203
26.42.3	RBAC_CONFIGURED_ROLES	1203
26.42.4	RBAC_HTTP_AUTH_ERROR	1203
26.42.5	RBAC_HTTP_AUTH_RESPONSE	1203
26.42.6	RBAC_LOAD_FAILED	1203
26.42.7	RBAC_LOAD_OK	1204
26.42.8	RBAC_READ_API_FILES	1204
26.42.9	RBAC_TRACE_HTTP_AUTH_ACCEPT	1204
26.42.10	RBAC_TRACE_HTTP_AUTH_BAD_BODY_TYPE	1204
26.42.11	RBAC_TRACE_HTTP_AUTH_BAD_COMMAND_TYPE	1204
26.42.12	RBAC_TRACE_HTTP_AUTH_COMMAND	1204
26.42.13	RBAC_TRACE_HTTP_AUTH_DISABLED	1204
26.42.14	RBAC_TRACE_HTTP_AUTH_EMPTY_BODY	1204
26.42.15	RBAC_TRACE_HTTP_AUTH_NO_COMMAND	1205
26.42.16	RBAC_TRACE_HTTP_AUTH_NO_JSON	1205
26.42.17	RBAC_TRACE_HTTP_AUTH_NO_REQUEST	1205
26.42.18	RBAC_TRACE_HTTP_AUTH_NO_TLS_REJECT	1205
26.42.19	RBAC_TRACE_HTTP_AUTH_REJECT	1205
26.42.20	RBAC_TRACE_HTTP_AUTH_RESPONSE	1205
26.42.21	RBAC_TRACE_HTTP_AUTH_ROLE	1205
26.42.22	RBAC_TRACE_HTTP_RESPONSE_BAD_BODY_TYPE	1205
26.42.23	RBAC_TRACE_HTTP_RESPONSE_CONTEXT	1206
26.42.24	RBAC_TRACE_HTTP_RESPONSE_DISABLED	1206
26.42.25	RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY	1206
26.42.26	RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY_LIST	1206
26.42.27	RBAC_TRACE_HTTP_RESPONSE_MODIFIED	1206
26.42.28	RBAC_TRACE_HTTP_RESPONSE_NO_ARGUMENTS	1206
26.42.29	RBAC_UNLOAD_OK	1206
26.43	RUN	1206
26.43.1	RUN_SCRIPT_LOAD	1206
26.43.2	RUN_SCRIPT_LOAD_ERROR	1207
26.43.3	RUN_SCRIPT_UNLOAD	1207
26.44	START	1207
26.44.1	START_REKEY_TIMER	1207
26.44.2	START_RETRY_TIMER	1207
26.45	STAT	1207
26.45.1	STAT_CMDS_DEINIT_OK	1207
26.45.2	STAT_CMDS_INIT_OK	1207
26.45.3	STAT_CMDS_LEASE4_FAILED	1208
26.45.4	STAT_CMDS_LEASE4_GET	1208
26.45.5	STAT_CMDS_LEASE4_GET_FAILED	1208

26.45.6	STAT_CMDS_LEASE4_GET_INVALID	1208
26.45.7	STAT_CMDS_LEASE4_GET_NO_SUBNETS	1208
26.45.8	STAT_CMDS_LEASE4_ORPHANED_STATS	1208
26.45.9	STAT_CMDS_LEASE6_FAILED	1208
26.45.10	STAT_CMDS_LEASE6_GET	1209
26.45.11	STAT_CMDS_LEASE6_GET_FAILED	1209
26.45.12	STAT_CMDS_LEASE6_GET_INVALID	1209
26.45.13	STAT_CMDS_LEASE6_GET_NO_SUBNETS	1209
26.45.14	STAT_CMDS_LEASE6_ORPHANED_STATS	1209
26.46	SUBNET	1209
26.46.1	SUBNET_CMDS_DEINIT_OK	1209
26.46.2	SUBNET_CMDS_INIT_FAILED	1209
26.46.3	SUBNET_CMDS_INIT_OK	1210
26.46.4	SUBNET_CMDS_NETWORK4_ADD_FAILED	1210
26.46.5	SUBNET_CMDS_NETWORK4_DEL_FAILED	1210
26.46.6	SUBNET_CMDS_NETWORK4_GET_FAILED	1210
26.46.7	SUBNET_CMDS_NETWORK4_LIST_FAILED	1210
26.46.8	SUBNET_CMDS_NETWORK4_SUBNET_ADD_FAILED	1210
26.46.9	SUBNET_CMDS_NETWORK4_SUBNET_DEL_FAILED	1210
26.46.10	SUBNET_CMDS_NETWORK6_ADD_FAILED	1211
26.46.11	SUBNET_CMDS_NETWORK6_DEL_FAILED	1211
26.46.12	SUBNET_CMDS_NETWORK6_GET_FAILED	1211
26.46.13	SUBNET_CMDS_NETWORK6_LIST_FAILED	1211
26.46.14	SUBNET_CMDS_NETWORK6_SUBNET_ADD_FAILED	1211
26.46.15	SUBNET_CMDS_NETWORK6_SUBNET_DEL_FAILED	1211
26.46.16	SUBNET_CMDS_NETWORK_ADD	1212
26.46.17	SUBNET_CMDS_NETWORK_DEL	1212
26.46.18	SUBNET_CMDS_NETWORK_GET	1212
26.46.19	SUBNET_CMDS_NETWORK_GET_EMPTY	1212
26.46.20	SUBNET_CMDS_NETWORK_LIST	1212
26.46.21	SUBNET_CMDS_NETWORK_LIST_EMPTY	1212
26.46.22	SUBNET_CMDS_NETWORK_SUBNET_ADD	1212
26.46.23	SUBNET_CMDS_NETWORK_SUBNET_DEL	1213
26.46.24	SUBNET_CMDS_SUBNET4_ADD_FAILED	1213
26.46.25	SUBNET_CMDS_SUBNET4_DELTA_ADD_FAILED	1213
26.46.26	SUBNET_CMDS_SUBNET4_DELTA_DEL_FAILED	1213
26.46.27	SUBNET_CMDS_SUBNET4_DEL_FAILED	1213
26.46.28	SUBNET_CMDS_SUBNET4_GET_FAILED	1213
26.46.29	SUBNET_CMDS_SUBNET4_LIST_FAILED	1214
26.46.30	SUBNET_CMDS_SUBNET4_UPDATE_FAILED	1214
26.46.31	SUBNET_CMDS_SUBNET6_ADD_FAILED	1214
26.46.32	SUBNET_CMDS_SUBNET6_DELTA_ADD_FAILED	1214
26.46.33	SUBNET_CMDS_SUBNET6_DELTA_DEL_FAILED	1214
26.46.34	SUBNET_CMDS_SUBNET6_DEL_FAILED	1214
26.46.35	SUBNET_CMDS_SUBNET6_GET_FAILED	1214
26.46.36	SUBNET_CMDS_SUBNET6_LIST_FAILED	1215
26.46.37	SUBNET_CMDS_SUBNET6_UPDATE_FAILED	1215
26.46.38	SUBNET_CMDS_SUBNET_ADD	1215
26.46.39	SUBNET_CMDS_SUBNET_DEL	1215
26.46.40	SUBNET_CMDS_SUBNET_GET	1215
26.46.41	SUBNET_CMDS_SUBNET_GET_EMPTY	1215
26.46.42	SUBNET_CMDS_SUBNET_LIST	1216
26.46.43	SUBNET_CMDS_SUBNET_LIST_EMPTY	1216
26.46.44	SUBNET_CMDS_SUBNET_UPDATE	1216

26.47	TCP	1216
26.47.1	TCP_CLIENT_REQUEST_RECEIVED	1216
26.47.2	TCP_CONNECTION_REJECTED_BY_FILTER	1216
26.47.3	TCP_CONNECTION_SHUTDOWN	1216
26.47.4	TCP_CONNECTION_SHUTDOWN_FAILED	1216
26.47.5	TCP_CONNECTION_STOP	1217
26.47.6	TCP_CONNECTION_STOP_FAILED	1217
26.47.7	TCP_DATA_RECEIVED	1217
26.47.8	TCP_DATA_SENT	1217
26.47.9	TCP_IDLE_CONNECTION_TIMEOUT_OCCURRED	1217
26.47.10	TCP_REQUEST_RECEIVED_FAILED	1217
26.47.11	TCP_REQUEST_RECEIVE_START	1217
26.47.12	TCP_SERVER_RESPONSE_SEND	1218
26.48	TKEY	1218
26.48.1	TKEY_EXCHANGE_ANSWER_CLASS	1218
26.48.2	TKEY_EXCHANGE_FAILED_TO_VERIFY	1218
26.48.3	TKEY_EXCHANGE_FAIL_EMPTY_IN_TOKEN	1218
26.48.4	TKEY_EXCHANGE_FAIL_EMPTY_OUT_TOKEN	1218
26.48.5	TKEY_EXCHANGE_FAIL_EMPTY_RESPONSE	1218
26.48.6	TKEY_EXCHANGE_FAIL_IO_ERROR	1218
26.48.7	TKEY_EXCHANGE_FAIL_IO_STOPPED	1218
26.48.8	TKEY_EXCHANGE_FAIL_IO_TIMEOUT	1219
26.48.9	TKEY_EXCHANGE_FAIL_NOT_SIGNED	1219
26.48.10	TKEY_EXCHANGE_FAIL_NO_RDATA	1219
26.48.11	TKEY_EXCHANGE_FAIL_NO_RESPONSE_ANSWER	1219
26.48.12	TKEY_EXCHANGE_FAIL_NULL_RESPONSE	1219
26.48.13	TKEY_EXCHANGE_FAIL_RESPONSE_ERROR	1219
26.48.14	TKEY_EXCHANGE_FAIL_TKEY_ERROR	1219
26.48.15	TKEY_EXCHANGE_FAIL_TO_INIT	1219
26.48.16	TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_ANSWER_COUNT	1220
26.48.17	TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_ANSWER_TYPE	1220
26.48.18	TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_OPCODE	1220
26.48.19	TKEY_EXCHANGE_NOT_A_RESPONSE	1220
26.48.20	TKEY_EXCHANGE_OUT_TOKEN_NOT_EMPTY	1220
26.48.21	TKEY_EXCHANGE_RDATA_COUNT	1220
26.48.22	TKEY_EXCHANGE_RECEIVE_MESSAGE	1220
26.48.23	TKEY_EXCHANGE_RESPONSE_TTL	1220
26.48.24	TKEY_EXCHANGE_SEND_MESSAGE	1221
26.48.25	TKEY_EXCHANGE_VALID	1221
26.48.26	TKEY_EXCHANGE_VERIFIED	1221
26.49	TLS	1221
26.49.1	TLS_CONNECTION_HANDSHAKE_FAILED	1221
26.49.2	TLS_CONNECTION_HANDSHAKE_START	1221
26.49.3	TLS_REQUEST_RECEIVE_START	1221
26.49.4	TLS_SERVER_RESPONSE_SEND	1221
26.50	USER	1222
26.50.1	USER_CHK_HOOK_LOAD_ERROR	1222
26.50.2	USER_CHK_HOOK_UNLOAD_ERROR	1222
26.50.3	USER_CHK_SUBNET4_SELECT_ERROR	1222
26.50.4	USER_CHK_SUBNET4_SELECT_REGISTRY_NULL	1222
26.50.5	USER_CHK_SUBNET6_SELECT_ERROR	1222
26.50.6	USER_CHK_SUBNET6_SELECT_REGISTRY_NULL	1222
26.51	Kea Debug Messages By Log Level	1222
26.51.1	Messages printed on debuglevel 0	1222

26.51.2	Messages printed on debuglevel 10	1223
26.51.3	Messages printed on debuglevel 15	1224
26.51.4	Messages printed on debuglevel 20	1225
26.51.5	Messages printed on debuglevel 40	1225
26.51.6	Messages printed on debuglevel 45	1246
26.51.7	Messages printed on debuglevel 50	1247
26.51.8	Messages printed on debuglevel 55	1253
26.51.9	Messages printed on debuglevel 70	1256
27	Configuration Templates	1257
27.1	Template: Home Network of a Power User	1257
27.1.1	Deployment Considerations	1257
27.1.2	Possible Extensions	1258
27.2	Template: Secure High Availability Kea DHCP with Multi-Threading	1271
27.2.1	Deployment Considerations	1272
27.2.2	Possible Extensions	1273
28	Kea Flow Diagrams	1287
28.1	Main Loop	1287
28.2	DHCPv4 Packet Processing	1287
28.3	DHCPREQUEST Processing	1287
28.4	DHCPv4 Subnet Selection	1287
28.5	DHCPv4 Special Case of Double-Booting	1292
28.6	DHCPv4 Lease Allocation	1292
28.7	Lease States	1294
28.8	Checking for Host Reservations	1294
28.9	Building the Options List	1294
28.10	How Kea Recognizes the Same Client In Different DHCP Messages	1294
28.11	RADIUS workflows for lease allocation	1294
29	Kea Configuration File Syntax (BNF)	1303
29.1	BNF Grammar for DHCPv4	1303
29.2	BNF Grammar for DHCPv6	1325
29.3	BNF Grammar for Control Agent	1349
29.4	BNF Grammar for DHCP-DDNS	1354
29.5	BNF Grammar for the Kea NETCONF Agent	1362
30	Acknowledgments	1367

Kea is an open source implementation of the Dynamic Host Configuration Protocol (DHCP) servers, developed and maintained by Internet Systems Consortium (ISC).

This is the reference guide for Kea version 3.0.2. Links to the most up-to-date version of this document (in PDF, HTML, and plain text formats) can be found on [Read the Docs](#). Other useful Kea information can be found in our [Knowledgebase](#).

INTRODUCTION

Kea is the next generation of DHCP software, developed by Internet Systems Consortium (ISC). It supports both the DHCPv4 and DHCPv6 protocols along with their extensions, e.g. prefix delegation and dynamic updates to DNS.

This guide covers Kea version 3.0.2.

For information about supported platforms see [Supported Platforms](#).

1.1 Supported Platforms

In general, this version of Kea builds and runs on any POSIX-compliant system with a C++ compiler (with C++14 support), the Botan cryptographic library, the log4cplus logging library and the Boost system library.

ISC regularly tests Kea on many operating systems and architectures, but lacks the resources to test all of them. Consequently, ISC is only able to offer support on a "best-effort" basis for some.

1.1.1 Regularly Tested Platforms

Kea is officially supported on Alpine, Debian, Fedora, FreeBSD, RHEL, and Ubuntu systems. Kea-3.0.2 builds have been tested on:

- Alpine — 3.18, 3.19, 3.20, 3.21, 3.22
- Debian — 10, 11, 12
- Fedora — 41, 42
- FreeBSD — 14
- RHEL — 8, 9, 10
- Ubuntu — 20.04, 22.04, 24.04

There are currently no plans to port Kea to Windows systems.

1.1.2 Best-Effort

The following are platforms on which Kea is known to build and run. ISC makes every effort to fix bugs on these platforms, but may be unable to do so quickly due to lack of hardware, less familiarity on the part of engineering staff, and other constraints.

- macOS — 13, 14

1.1.3 Community-Maintained

These systems have once been regularly tested, but official support for it has been abandoned, usually due to discontinued support on their own part. Older versions may not have the required dependencies for building Kea easily available, although it is possible in many cases to compile on those directly from source. The community and interested parties may wish to help with maintenance, and we welcome patch contributions, although we cannot guarantee that we will accept them. All contributions are assessed against the risk of adverse effect on officially supported platforms.

These include platforms past their respective EOL dates, such as:

- Alpine — 3.16 (EOL 23 May 2022) and older
- CentOS — 8 (EOL 31 December 2021) and older
- Debian — 9 (EOL 30 June 2022) and older
- Fedora — 38 (EOL 21 May 2024) and older
- FreeBSD — 12 and older
- macOS — 10.15 (EOL 12 September 2022) and older
- Ubuntu — 21.04 (EOL 20 January 2022) and older

1.1.4 Unsupported Platforms

These are platforms on which versions of Kea since 1.7 are known *not* to build or run:

- Windows (all versions)
- Windows Server (all versions)
- Any platform with OpenSSL 1.0.1 or earlier, which does not also have Botan as an alternative
- Any platform with log4cplus version 1.0.2 or earlier.
- Any platform with a compiler that doesn't support C++14.

1.2 Required Software at Runtime

Kea uses various extra software packages which may not be provided in the default installation of some operating systems, nor in the standard package collections. This required software may need to be installed separately. (For the build requirements, also see [Build Requirements](#).)

- Kea supports two cryptographic libraries: Botan and OpenSSL. Only one of them is required to be installed during compilation. Kea uses the Botan library for C++ (<https://botan.randombit.net/>), version 2.0 or later; support for Botan versions earlier than 2.0 was removed as of Kea 1.7.0. As an alternative to Botan, Kea can use the OpenSSL cryptographic library (<https://www.openssl.org/>), version 1.0.2 or later.
- Kea uses the log4cplus C++ logging library (<https://sourceforge.net/p/log4cplus/wiki/Home/>). It requires log4cplus version 1.0.3 or later.
- Kea requires the Boost system library (<https://www.boost.org/>). Building with the header-only version of Boost is no longer recommended.

Some optional features of Kea have additional dependencies.

- To store lease information in a MySQL database, Kea requires MySQL headers and libraries. This is an optional dependency; Kea can be built without MySQL support.
- To store lease information in a PostgreSQL database, Kea requires PostgreSQL headers and libraries. This is an optional dependency; Kea can be built without PostgreSQL support.

- Kea provides a NETCONF interface with the `kea-netconf` agent. This Kea module requires Sysrepo software when used. Building Kea with NETCONF support requires many dependencies to be installed, which are described in more detail in *Installing NETCONF*. This is an optional dependency; Kea can be built without NETCONF support.
- To sign and verify DNS updates the Kea DDNS server may use GSS-TSIG, which requires MIT Kerberos 5 or Heimdal libraries. The dependencies required to be installed are described in more detail in *GSS-TSIG Compilation*. This is an optional dependency; Kea can be built without GSS-TSIG support.

1.3 Kea Software

Kea is a modular DHCP server solution. This modularity is accomplished using multiple cooperating processes which, together, provide the server functionality. The following software is included with Kea:

- *keactrl* — This tool starts, stops, reconfigures, and reports the status of the Kea servers.
- *kea-dhcp4* — The DHCPv4 server process. This process responds to DHCPv4 queries from clients.
- *kea-dhcp6* — The DHCPv6 server process. This process responds to DHCPv6 queries from clients.
- *kea-dhcp-ddns* — The DHCP Dynamic DNS process. This process acts as an intermediary between the DHCP servers and external DNS servers. It receives name update requests from the DHCP servers and sends DNS update messages to the DNS servers.
- *kea-admin* — This is a useful tool for database backend maintenance (creating a new database, checking versions, upgrading, etc.).
- *kea-lfc* — This process removes redundant information from the files used to provide persistent storage for the memfile database backend. While it can be run standalone, it is normally run as and when required by the Kea DHCP servers.
- *kea-ctrl-agent* — The Kea Control Agent (CA) is a daemon that exposes a RESTful control interface for managing Kea servers.
- *kea-netconf* - `kea-netconf` is an agent that provides a YANG/NETCONF interface for configuring Kea.
- *kea-shell* — This simple text client uses the REST interface to connect to the Kea Control Agent.
- *perfdhcp* — This is a DHCP benchmarking tool which simulates multiple clients to test both DHCPv4 and DHCPv6 server performance.

The tools and modules are covered in full detail in this guide. In addition, manual pages are also provided in the default installation.

Kea also provides C++ libraries and programmer interfaces for DHCP. These include detailed developer documentation and code examples.

QUICK START

This section describes the basic steps needed to get Kea up and running. For further details, full customizations, and troubleshooting, see the respective chapters elsewhere in this Kea Administrator Reference Manual (ARM).

2.1 Quick Start Guide Using tarball

1. Install required runtime and build dependencies. See *Build Requirements* for details.
2. Download the Kea source tarball from [the main isc.org downloads page](#), [the ISC downloads site](#), or [the ISC Cloudsmith page](#).

3. Extract the tarball. For example:

```
$ tar -Jxvf kea-3.0.2.tar.xz
```

4. Go into the source directory and set up the build:

```
$ cd kea-3.0.2
$ meson setup build [your extra parameters]
```

5. Build it:

```
$ meson compile -C build
```

6. Install it (by default it will be placed in `/usr/local/`, so root privileges are likely required for this step):

```
$ meson install -C build
```

2.2 Quick Start Guide Using Native Packages

ISC provides native Alpine, deb, and RPM packages, which make Kea installation much easier than building from source. Unless specific compilation options are desired, it is usually easier to install Kea using native packages.

1. Go to [Kea on cloudsmith.io](#).
2. Choose the Cloudsmith repository e.g. kea-3-0 for Kea 3.0.2.

Note

All stable releases are stored in separate repositories, such as kea-2-6, kea-3-0, and so on. All development releases, which are not recommended for production use, are located in the kea-dev repository.

3. Click on the arrow beside the "Set Me Up" button and select the desired OS flavor: Alpine, Debian, or RedHat.
4. Follow the onscreen instructions.

Note

For example, the Debian setup instructions for Kea 3.0 can be found here: <https://cloudsmith.io/~isc/repos/kea-3-0/setup/#formats-deb>

The dropdown near the top of the page offers instructions for other operating systems.

5. Kea is split into various packages. The entire list is available on the Cloudsmith repository page under Packages, or it can be retrieved using `apk/apt/dnf`.

Debian/Ubuntu	<code>\$ apt search isc-kea</code>
Note <code>isc-kea-dhcp4-server</code> and <code>isc-kea-dhcp6-server</code> are empty transitional packages. The working server packages are <code>isc-kea-dhcp4</code> and <code>isc-kea-dhcp6</code> .	
Fedora/Redhat	<code>\$ dnf search 'isc-kea*'</code>
Alpine	<code>\$ apk search isc-kea</code>

6. Install the metapackage containing all the tools, libraries, servers, documentation, and open source hooks:

Debian/Ubuntu	<code>\$ sudo apt install isc-kea</code>
Fedora/Redhat	<code>\$ sudo dnf install isc-kea</code>
Alpine	<code># apk add isc-kea</code>

or specific packages:

Debian/Ubuntu	<code>\$ sudo apt install isc-kea-dhcp6</code>
Fedora/Redhat	<code>\$ sudo dnf install isc-kea-dhcp6</code>
Alpine	<code>\$ apk add isc-kea-dhcp6</code>

or ALL Kea-related packages, including development headers, debug symbols, and subscriber hooks (if available):

Debian/Ubuntu `$ sudo apt install 'isc-kea*'`

Fedora/Red `$ sudo dnf install 'isc-kea*'`

Alpine Installing packages via globbing (*) is not available for Alpine, but it can be simulated with the following command:
`# apk search isc-kea | sed 's/-[0-9].*//g' | xargs apk add`

or all packages with a specified version number:

Debian/Ubuntu `$ sudo apt install 'isc-kea*=2.4.0-isc20230921141113'`

Fedora/Red `$ sudo dnf install 'isc-kea*2.4.0-isc20230921141113*'`

Alpine Installing packages via globbing (*) is not available for Alpine, but it can be simulated with the following command:
`# apk search isc-kea | sed 's/-[0-9].*//g' | grep r20230921141113 | xargs apk add`

8. All installed packages should be now available.

You can start a server manually:

```
# kea-dhcp6 -c /etc/kea/kea-dhcp6.conf
```

or using systemd:

```
# systemctl restart kea-dhcp6
```

or using OpenRC on Alpine:

```
# service kea-dhcp6 restart
```

Note

keactrl is not available in packages, as similar functionality is provided by the native systemd scripts.

9. On Debian/Ubuntu systems, the service is automatically enabled at boot time when the package is installed. On Fedora/RHEL and Alpine, the service is not automatically enabled, so if desired, it must be enabled manually.

With systemd on Fedora/RedHat:

```
# systemctl enable kea-dhcp6
```

With OpenRC on Alpine:

```
# rc-update add kea-dhcp6
```

2.3 Quick Start Guide Using Docker Containers

1. Go to the [ISC docker repository on cloudsmith.io](#).
2. Create an ipvlan network attached to the client-facing host interface and assigned to the subnet that is served by Kea.

```
$ docker network create --driver ipvlan --ipv6 --subnet 2001:db8::/64 --opt parent=eth0 \
↳ ipvlan0
```

3. Pick the desired image and pull it locally.

```
$ docker pull docker.cloudsmith.io/isc/docker/kea-dhcp6
```

4. Create a container from the image. Mount the configuration volume and the data volume if needed.

```
$ docker create \
  --name kea-dhcp6 \
  --network ipvlan0 \
  --volume /local/kea/config:/etc/kea \
  --volume /local/kea/data:/var/lib/kea \
  docker.cloudsmith.io/isc/docker/kea-dhcp6
```

5. Start the docker container.

```
$ docker start kea-dhcp6
```

6. To stop the docker container, run:

```
$ docker stop kea-dhcp6
```

Note

Refer to the [kea-docker readme](#) for more complex scenarios.

2.4 Quick Start Guide for DHCPv4 and DHCPv6 Services

1. Edit the Kea configuration files, which by default are installed in the `[kea-install-dir]/etc/kea/` directory. These are: `kea-dhcp4.conf`, `kea-dhcp6.conf`, `kea-dhcp-ddns.conf` and `kea-ctrl-agent.conf`, `keactrl.conf` for the DHCPv4 server, DHCPv6 server, D2, Control Agent, and keactrl script, respectively.
2. To start the DHCPv4 server in the background, run the following command (as root):

```
# keactrl start -s dhcp4
```

Or run the following command to start the DHCPv6 server:

```
# keactrl start -s dhcp6
```

Note that it is also possible to start all servers simultaneously:

```
# keactrl start
```

3. Verify that the Kea server(s) is/are running:

```
# keactrl status
```

A server status of "inactive" may indicate a configuration error. Please check the log file (by default named `[kea-install-dir]/var/log/kea/kea-dhcp4.log`, `[kea-install-dir]/var/log/kea/kea-dhcp6.log`, `[kea-install-dir]/var/log/kea/kea-ddns.log`, or `[kea-install-dir]/var/log/kea/kea-ctrl-agent.log`) for the details of any errors.

4. If the server has started successfully, test that it is responding to DHCP queries and that the client receives a configuration from the server; for example, use the [ISC DHCP client](#).
5. To stop running the server(s):

```
# keactrl stop
```

For system-specific instructions, please read the [system-specific notes](#), available in the Kea section of [ISC's Knowledgebase](#).

The details of `keactrl` script usage can be found in [Managing Kea with keactrl](#).

Once Kea services are up and running, consider deploying a dashboard solution to monitor running services. For more details, see [Monitoring Kea With Stork](#).

2.5 Running the Kea Servers Directly

The Kea servers can be started directly, without the need to use `keactrl` or `systemctl`. To start the DHCPv4 server run the following command:

```
# kea-dhcp4 -c /path/to/your/kea4/config/file.json
```

Similarly, to start the DHCPv6 server, run the following command:

```
# kea-dhcp6 -c /path/to/your/kea6/config/file.json
```


INSTALLATION

3.1 Packages

ISC publishes native RPM, deb, and APK packages, as well as tarballs with the source code. The packages are available in ISC's [Cloudsmith repositories](#). The native packages can be downloaded and installed using the system available in a specific distribution (such as dpkg or rpm). The Kea repository can also be added to the system, making it easier to install updates. For details, please go to <https://cloudsmith.io/~isc/repos>, choose the desired Kea version, and then click the "Set Me Up" button. For detailed instructions, please refer to this [Knowledgebase article](#).

ISC maintains two types of repositories: stable and development. The stable repositories contain a single stable release (e.g., kea-2-6 or kea-3-0) along with all its maintenance updates. Separate repositories were introduced to minimize the risk of unintentionally upgrading from one stable release to another.

The development repository, kea-dev, includes current and future development releases, which ISC does not recommend for production use. Packages in the kea-dev repository are subject to cleanup, and older versions may be removed.

3.1.1 Installation From Cloudsmith Packages

ISC provides Kea packages for Alpine, Debian, Fedora, RHEL, and Ubuntu. The recommended method for installing Kea on any of these systems is to install the `isc-kea` metapackage from the Cloudsmith repository. This metapackage is included on all supported distros and installs all of the services offered by the Kea software suite.

Specific Kea components can be installed individually, with any of the following packages:

- `isc-kea-dhcp4` — Kea DHCPv4 server package
- `isc-kea-dhcp6` — Kea DHCPv6 server package
- `isc-kea-dhcp-ddns` — Kea DHCP DDNS server
- `isc-kea-ctrl-agent` — Kea Control Agent for remote configuration
- `isc-kea-admin` — Kea database administration tools
- `isc-kea-hooks` — Kea open source DHCP hooks

The Kea subscriber hook packages are not included in the `isc-kea-hooks` package. For ISC customers with access to the subscriber hooks, those packages have the `isc-kea-subscriber-` prefix. For users interested in purchasing professional support services from ISC and getting access to the subscriber hooks, please fill out our contact form at <https://www.isc.org/contact>.

Once installed, the services can be managed through the distribution's service manager. The services are named: `kea-dhcp4`, `kea-dhcp6`, `kea-dhcp-ddns`, and `kea-ctrl-agent`.

Note

The real service names on Debian and Ubuntu use slightly different package names, to maintain compatibility with some older scripts. A systemd service alias is used to allow users to refer to them with shorter names. Calling `systemctl enable` on these services requires the real service names, which are: `isc-kea-dhcp4-server`, `isc-kea-dhcp6-server`, `isc-kea-dhcp-ddns-server`, and `isc-kea-ctrl-agent`.

3.1.2 Caveats When Upgrading Kea Packages

To upgrade to a current version of Kea from version 2.3.2 or earlier on Debian and Ubuntu systems, run `apt dist-upgrade` instead of the usual `apt upgrade`. Once this upgrade has been completed, it is possible to upgrade to later versions normally using `apt upgrade` on Debian and Ubuntu systems.

Users may notice differences in the packages distributed in Kea versions prior to 2.3.2 and those distributed with 2.3.2 and later. As a result of an overhaul of our package design with that release, some packages were renamed or removed. To ensure that upgrades go as smoothly as possible, pay attention to which packages are being removed and installed by the upgrade transaction, and ensure that all required packages are reinstalled.

Specifically, there is a possibility for the following packages to be removed during the upgrade, depending on which packages were originally installed:

- `isc-kea-dhcp4`
- `isc-kea-dhcp6`
- `isc-kea-dhcp-ddns`
- `isc-kea-hooks`

To install the entire Kea software suite, simply run `apt install isc-kea` after upgrading, which will install all of the relevant subpackages that make up Kea.

This upgrade path issue does not apply to RPM and Alpine systems.

Customers with ISC support contracts who experience difficulties with upgrading are invited to open a ticket in their support queue. Other users are encouraged to describe their situation on the [kea-users mailing list](#) for best-effort support from other list members.

3.2 Installation Hierarchy

The following is the directory layout of the complete Kea installation. (All directory paths are relative to the installation directory.)

- `etc/kea/` — configuration files.
- `include/` — C++ development header files.
- `lib/` — libraries.
- `lib/kea/hooks` — additional hook libraries.
- `sbin/` — server software and commands used by the system administrator.
- `share/doc/kea/` — this guide, other supplementary documentation, and examples.
- `share/kea/` — API command examples and database schema scripts.
- `share/man/` — manual pages (online documentation).
- `var/lib/kea/` — server identification and lease database files.

- `var/log/kea/` - log files.
- `var/run/kea` - PID file and logger lock file.

3.3 Build Requirements

In addition to the runtime requirements (listed in *Required Software at Runtime*), building Kea from source code requires various development include headers and program development tools.

Note

Some operating systems have split their distribution packages into a runtime and a development package. The development package versions, which include header files and libraries, must be installed to build Kea from the source code.

Building from source code requires the following software installed on the system:

- The development tools `meson`, `ninja`, and `pkg-config`.
- Boost C++ libraries (<https://www.boost.org/>). The oldest Boost version used for testing is 1.67 (although Kea may also work with older versions e.g. 1.66). The Boost system library must also be installed. Installing a header-only version of Boost is not recommended.
- OpenSSL (at least version 1.0.2) or Botan (at least version 2). OpenSSL version 1.1.1 or later is strongly recommended.
- `log4cplus` (at least version 1.0.3) development include headers.
- A C++ compiler (with C++14 support) and standard development headers.
- The MySQL client and the client development libraries, when using the `-D mysql=enabled` configuration flag to build the Kea MySQL database backend. In this case, an instance of the MySQL server running locally or on a machine reachable over a network is required. Note that running the unit tests requires a local MySQL server.
- The PostgreSQL client and the client development libraries, when using the `-D postgresql=enabled` configuration flag to build the Kea PostgreSQL database backend. In this case an instance of the PostgreSQL server running locally or on a machine reachable over a network is required. Note that running the unit tests requires a local PostgreSQL server.
- `Sysrepo v1.4.140` and `libyang v1.0.240`, needed to connect to a Sysrepo datastore. Earlier versions are no longer supported. When compiling from sources, the configure switch that can be used is `-D netconf=enabled`. If these dependencies were installed in custom paths, point `PKG_CONFIG_PATH` to them.
- The MIT Kerberos 5 or Heimdal libraries are needed by Kea DDNS server to sign and verify DNS updates using GSS-TSIG. The configuration switch which enables this functionality is `-D krb5=enabled`. If this dependency is installed in a custom path, point `PKG_CONFIG_PATH` to the `.pc` file or `PATH` to the `krb5-config` tool.
- `googletest` (version 1.8 or later) is required when using the `-D tests=enabled` configuration option to build the unit tests.
- The documentation generation tools `Sphinx`, `texlive` with its extensions, and `Doxygen`, to create the documentation. Specifically, with Fedora, `python3-sphinx`, `python3-sphinx_rtd_theme`, `texlive`, and `texlive-collection-latexextra` are necessary. With Ubuntu, `python3-sphinx`, `python3-sphinx-rtd-theme`, `texlive`, and `texlive-latex-extra` are needed. If LaTeX packages are missing, Kea skips PDF generation and produces only HTML documents.

Visit ISC's Knowledgebase at <https://kb.isc.org/docs/installing-kea> for system-specific installation tips.

3.4 Installation From Source

Although Kea may be available in pre-compiled, ready-to-use packages from operating system vendors, it is open source software written in C++. As such, it is freely available in source code form from ISC as a downloadable tar file. The source code can also be obtained from the Kea GitLab repository at <https://gitlab.isc.org/isc-projects/kea>. This section describes how to build Kea from the source code.

3.4.1 Download Tar File

The Kea release tarballs may be downloaded from: <https://downloads.isc.org/isc/kea/>.

3.4.2 Verify The Tar File Signature

The tar file with the source code is distributed together with its GPG signature. The signature is a file with the same name as the tar file appended by the .asc extension. You can find the signature file on our download page, FTP, or CloudSmith.

The signature is created using the ISC code-signing key. The current set of ISC code-signing keys is available from the ISC website at <https://www.isc.org/pgpkey> (the [Current Set of ISC Code-Signing Keys](#) link).

The signature can be verified using the GnuPG `gpg` tool. The following commands import the code-signing keys (`isc-keyblock.asc`) and verify the signature:

```
$ gpg --import isc-keyblock.asc
$ gpg --verify kea-X.Y.Z.tar.xz.asc kea-X.Y.Z.tar.xz
```

The verification allows users to confirm that the tar file has not been tampered with and that it was created by ISC.

3.4.3 Retrieve From Git

The latest development code is available on GitLab (see <https://gitlab.isc.org/isc-projects/kea>). The Kea source is public and development is done in the “master” branch.

Downloading this “bleeding edge” code is recommended only for developers or advanced users. Using development code in a production environment is not recommended.

Note

When building from source code retrieved via git, additional software is required: `automake` (v1.11 or later), `libtoolize`, and `autoconf` (v2.69 or later). These may need to be installed.

The code can be checked out from <https://gitlab.isc.org/isc-projects/kea.git>:

```
$ git clone https://gitlab.isc.org/isc-projects/kea.git
```

The code checked out from the git repository does not include the build files. They can be created by running `meson setup build`.

Write access to the Kea repository is only granted to ISC staff. Developers planning to contribute to Kea should check our [Contributor's Guide](#). The [Kea Developer's Guide](#) contains more information about the process, and describes the requirements for contributed code to be accepted by ISC.

3.4.4 Set Up the Build

Kea uses Meson to discover build environment details. To generate the ninja file using the defaults, simply run:

```
$ meson setup build
```

Run `meson configure` to view the different build options. Some commonly used options are:

- `--prefix` Define the installation location (the default is `/usr/local`).
- `-D mysql=enabled` Build Kea with code to allow it to store leases and host reservations in a MySQL database.
- `-D postgresql=enabled` Build Kea with code to allow it to store leases and host reservations in a PostgreSQL database.
- `-D krb5=enabled` Build Kea with Kerberos5 with GSS-API support required by [libddns-gss-tsig.so](#).
- `-D netconf=enabled` Build Kea with libyang and sysrepo support required by [kea-netconf](#).
- `-D crypto=botan` Specify the name of the Botan pkg-config library e.g. `botan-2` to build with Botan for cryptographic functions. It is preferable to use OpenSSL (see below).
- `--D crypto=openssl` Use the OpenSSL cryptographic library instead of Botan. By default Meson searches for a valid Botan installation; if one is not found, Kea searches for OpenSSL. Normally this is not necessary.

Note

For instructions concerning the installation and configuration of database backends for Kea, see [DHCP Database Installation and Configuration](#).

If `meson setup build` fails, it may be due to missing or old dependencies.

When `meson setup build` succeeds, it displays a report with the parameters used to build the code. This report is saved into the file `build/config.report` and is also embedded into the executable binaries, e.g. [kea-dhcp6](#).

3.4.5 Build

After the setup step is complete, build the executables from the C++ code and prepare the Python scripts by running the command:

```
$ meson compile -C build
```

3.4.6 Install

To install the Kea executables, support files, and documentation, issue the command:

```
$ meson install -C build
```

Note

The install step may require superuser privileges.

It should not be required, but if shared libraries are not found at runtime, you can run `ldconfig` as root with `/usr/local/lib` (or with `${prefix}/lib` if set up with `--prefix`) in `/etc/ld.so.conf` (or the relevant linker cache configuration file for the OS):

```
$ ldconfig
```

If `ldconfig` is not run where required, users may see errors like the following:

```
program: error while loading shared libraries: libkea-something.so.1:
cannot open shared object file: No such file or directory
```

3.4.7 Cross-Building

It is possible to cross-build Kea, i.e. to create binaries in a separate system (the build system) from the one where Kea runs (the host system).

It is outside of the scope of common administrator operations and requires some developer skills, but the Developer Guide explains how to do that using an x86_64 Linux system to build Kea for a Raspberry Pi box running Raspbian: see this [Kea Cross-Compiling Example](#).

3.5 DHCP Database Installation and Configuration

Kea stores its leases in a lease database. The software has been written in a way that makes it possible to choose which database product should be used to store the lease information. Kea supports three database backends: MySQL, PostgreSQL, and memfile. To limit external dependencies, MySQL and PostgreSQL support are disabled by default and only memfile is available. Support for the optional external database backend must be explicitly included when Kea is built. This section covers the building of Kea with one of the optional backends and the creation of the lease database.

Note

When unit tests are built with Kea (i.e. the `-D tests=enabled` configuration option is specified), the databases must be manually pre-configured for the unit tests to run. The details of this configuration can be found in the [Kea Developer's Guide](#).

3.5.1 Building with MySQL Support

Install MySQL according to the instructions for the system. The client development libraries must be installed.

Build and install Kea as described in [Installation](#), with the following modification. To enable the MySQL database code, at the setup step (see [Set Up the Build](#)), the `-D mysql=enabled` switch should be specified:

```
$ meson setup build -D mysql=enabled
```

If MySQL was not installed in the default location, the location can be selected by setting `PKG_CONFIG_PATH=/path/to/mariadb.pc:${PKG_CONFIG_PATH}` or `PATH=/path/to/mariadb-config:${PATH}` prior to `meson setup`.

```
$ PKG_CONFIG_PATH=/opt/mariadb/lib/pkgconfig PATH=/opt/mariadb/bin meson setup build -D_
↪mysql=enabled
```

See [First-Time Creation of the MySQL Database](#) for details regarding MySQL database configuration.

3.5.2 Building with PostgreSQL Support

Install PostgreSQL according to the instructions for the system. The client development libraries must be installed. Client development libraries are often packaged as `libpq`.

Build and install Kea as described in [Installation](#), with the following modification. To enable the PostgreSQL database code, at the setup step (see [Set Up the Build](#)), the `-D postgresql=enabled` switch should be specified.

If PostgreSQL was not installed in the default location, the location can be selected by setting `PKG_CONFIG_PATH=/path/to/libpq.pc:${PKG_CONFIG_PATH}` or `PATH=/path/to/pg_config:${PATH}` prior to `meson setup`.

```
$ PKG_CONFIG_PATH=/opt/postgresql/lib/pkgconfig PATH=/opt/postgresql/bin meson setup.
↪ build -D postgresql=enabled
```

See [First-Time Creation of the PostgreSQL Database](#) for details regarding PostgreSQL database configuration.

3.6 Hammer Building Tool

Hammer is a Python 3 script that lets users automate tasks related to building Kea, such as setting up virtual machines, installing Kea dependencies, compiling Kea with various options, running unit-tests and more. This tool was created primarily for internal QA purposes at ISC and it is not included in the Kea distribution; however, it is available in the Kea git repository. This tool was developed primarily for internal purposes and ISC cannot guarantee its proper operation. Administrators who decide to use it should do so with care.

Note

Use of this tool is completely optional. Everything it does can be done manually.

The first-time user is strongly encouraged to look at Hammer's built-in help:

```
$ ./hammer.py --help
```

It will list available parameters.

Hammer is able to set up various operating systems running either in LXC or in VirtualBox. For a list of supported systems, use the `supported-systems` command:

```
$ ./hammer.py supported-systems
fedora:
- 37: lxc
- 38:
centos:
- 8: lxc, virtualbox
- 9:
rhel:
- 8: virtualbox
- 9:
ubuntu:
- 18.04: lxc, virtualbox
- 20.04: lxc
- 22.04: lxc
debian:
- 10: lxc, virtualbox
- 11: lxc
- 12: lxc
freebsd:
- 12.0: virtualbox
- 12.1:
```

(continues on next page)

(continued from previous page)

```
- 13.0: virtualbox
alpine:
- 3.15: lxc
- 3.16: lxc
- 3.17: lxc
```

It is also possible to run the build locally, in the current system (if the OS is supported).

First, the Hammer dependencies must be installed: Vagrant and either VirtualBox or LXC. Hammer can install Vagrant and the required Vagrant plugins using the command:

```
$ ./hammer.py ensure-hammer-deps
```

VirtualBox and LXC must be installed manually.

The basic functions provided by Hammer are to prepare the build environment and perform the actual build, and to run the unit tests locally in the current system. This can be achieved by running the command:

```
$ ./hammer.py build -p local
```

The scope of the process can be defined using the `--with (-w)` and `--without (-x)` options. By default, the build command builds Kea with documentation, installs it locally, and runs unit tests.

To exclude the installation and generation of docs, type:

```
$ ./hammer.py build -p local -x install docs
```

The basic scope can be extended by `mysql`, `pgsql`, `native-pkg`, `shell`, and `forge`. Please refer to `./hammer.py build --help` for more details.

Note

If building Kea locally, Hammer dependencies like Vagrant are not needed.

Hammer can be told to set up a new virtual machine with a specified operating system, without the build:

```
$ ./hammer.py prepare-system -p virtualbox -s freebsd -r 12.0
```

This way, a system can be prepared for our own use.

To prepare such a system using SSH, invoke:

```
$ ./hammer.py ssh -p virtualbox -s freebsd -r 12.0
```

It is possible to speed up subsequent Hammer builds via `ccache`. During compilation, `ccache` stores objects in a shared folder. In subsequent runs, instead of doing an actual compilation, `ccache` returns the stored earlier objects. The cache with these objects for reuse must be stored outside of VM or LXC. To indicate the folder, the `--ccache-dir` parameter for Hammer must be included. In the indicated folder, there are separate stored objects for each target operating system.

```
$ ./hammer.py build -p lxc -s ubuntu -r 18.04 --ccache-dir ~/kea-ccache
```

Note

`ccache` is currently only supported for LXC in Hammer; support for VirtualBox may be added later.

For more information check:

```
$ ./hammer.py --help
```

3.7 Running Kea From a Non-root Account on Linux

Both Kea DHCPv4 and DHCPv6 servers perform operations that in general require root access privileges. In particular, DHCPv4 opens raw sockets and both DHCPv4 and DHCPv6 open UDP sockets on privileged ports. However, with some extra system configuration, it is possible to run Kea from non-root accounts.

First, a regular user account must be created:

```
useradd admin
```

Then, change the binaries' ownership and group to the new user. Note that the specific path may be different. Please refer to the `--prefix` parameter passed to `meson setup`:

```
chown -R admin /opt/kea
chgrp -R admin /opt/kea
chown -R admin /var/log/kea/kea-dhcp4.log
chgrp -R admin /var/log/kea/kea-dhcp4.log
chown -R admin /var/log/kea/kea-dhcp6.log
chgrp -R admin /var/log/kea/kea-dhcp6.log
```

If using systemd, modify its service file (e.g. `/etc/systemd/system/kea-dhcp6.service`):

```
User=admin
Group=admin
```

The most important step is to set the capabilities of the binaries. Refer to the operating system man page for *capabilities* for more information.

```
setcap 'cap_net_bind_service,cap_net_raw=+ep' /opt/kea/sbin/kea-dhcp4
setcap 'cap_net_bind_service=+ep' /opt/kea/sbin/kea-dhcp6
```

If using systemd, also add this to the service file (e.g. `/etc/systemd/system/kea-dhcp6.service`):

```
ExecStartPre=setcap 'cap_net_bind_service=+ep' /opt/kea/sbin/kea-dhcp6
```

After this step is complete, the admin user should be able to run Kea. Note that the DHCPv4 server by default opens raw sockets. If the network is only using relayed traffic, Kea can be instructed to use regular UDP sockets (refer to `dhcp-socket-type` parameter in the [Interface Configuration](#) section) and the `cap_net_raw` capability can be skipped.

Note

It is possible to avoid running Kea with root privileges by instructing Kea to use non-privileged (greater than 1024) ports and redirecting traffic. This, however, only works for relayed traffic. This approach in general is considered experimental and has not been tested for deployment in production environments. Use with caution!

To use this approach, configure the server to listen on other non-privileged ports (e.g. 1547 and 1548) by running the process with the `-p` option in `/etc/systemd/system/kea-dhcp4.service`:


```
ExecStart=/opt/kea/sbin/kea-dhcp4 -d -c /etc/kea/kea-dhcp4.conf -p 2067
```

and /etc/systemd/system/kea-dhcp4.service:

```
ExecStart=/opt/kea/sbin/kea-dhcp6 -d -c /etc/kea/kea-dhcp6.conf -p 1547
```

Then configure port redirection with iptables and ip6tables for new ports (e.g. 1547 and 1548). Be sure to replace `ens4` with the specific interface name.

```
iptables -t nat -A PREROUTING -i ens4 -p udp --dport 67 -j REDIRECT --to-port 2067
iptables -t nat -A PREROUTING -i ens4 -p udp --dport 2068 -j REDIRECT --to-port 68
ip6tables -t nat -A PREROUTING -i ens4 -p udp --dport 547 -j REDIRECT --to-port 1547
ip6tables -t nat -A PREROUTING -i ens4 -p udp --dport 1548 -j REDIRECT --to-port 548
```

3.8 Deprecated Features

This section lists significant features that have been or will be removed. We try to deprecate features before removing them, to signal to current users to plan a migration. New users should not rely on deprecated features.

3.8.1 Sysrepo 0.x or 1.x

Kea 2.3.2 introduced support for Sysrepo 2.x. Unfortunately, Sysrepo continues to undergo major changes that are backward-incompatible, and current Kea versions do not support Sysrepo earlier than versions 2.x.

KEA DATABASE ADMINISTRATION

4.1 Databases and Schema Versions

Kea may be configured to use a database as storage for leases or as a source of servers' configurations and host reservations (i.e. static assignments of addresses, prefixes, options, etc.). As Kea is updated, new database schemas are introduced to facilitate new features and correct discovered issues with the existing schemas.

Each version of Kea expects a particular schema structure and checks for this by examining the version of the database it is using. Separate version numbers are maintained for the schemas, independent of the version of Kea itself. It is possible that the schema version will stay the same through several Kea revisions; similarly, it is possible that the version of the schema may go up several revisions during a single Kea version upgrade. Versions for each backend type are also independent, so an increment in the MySQL backend version does not imply an increment in that of PostgreSQL.

Schema versions are specified in a major.minor format. For the most recent versions, the minor version is always zero and only the major version is incremented.

Historically, the minor version used to be incremented when backward-compatible changes were introduced to the schema: for example - when a new index is added. This was opposed to incrementing the major version which implied an incompatible schema change: for example - changing the type of an existing column. If Kea attempts to run on a schema that is too old, as indicated by a mismatched schema version, it will fail; administrative action is required to upgrade the schema.

4.2 The `kea-admin` Tool

To manage the databases, Kea provides the `kea-admin` tool. It can initialize a new backend, check its version number, perform a backend upgrade, and dump lease data to a text file.

`kea-admin` takes two mandatory parameters: `command` and `backend`. Additional, non-mandatory options may be specified. The currently supported commands are:

- `db-init` — initializes a new database schema, which is useful during a new Kea installation. The new database is updated to match the Kea version being installed. `kea-admin` is automatically invoked with this command if a missing schema is detected during startup or reconfiguration of Kea DHCP servers.
- `db-version` — reports the database backend version number. This is not necessarily equal to the Kea version number, as each backend has its own versioning scheme.
- `db-upgrade` — conducts a database schema upgrade. This is useful when upgrading Kea.
- `lease-dump` — dumps the contents of the lease database (for MySQL or PostgreSQL backends) to a CSV (comma-separated values) text file.

The first line of the file contains the column names. This can be used as a way to switch from a database backend to a memfile backend. Alternatively, it can be used as a diagnostic tool, so it provides a portable form of the lease data.

- `lease-upload` — uploads leases from a CSV (comma-separated values) text file to a MySQL or a PostgreSQL lease database. The CSV file needs to be in memfile format.

backend specifies the type of backend database. The currently supported types are:

- `memfile` — lease information is stored on disk in a text file.
- `mysql` — information is stored in a MySQL relational database.
- `pgsql` — information is stored in a PostgreSQL relational database.

Additional parameters may be needed, depending on the setup and specific operation: username, password, and database name or the directory where specific files are located. See the appropriate manual page for details (man 8 `kea-admin`).

4.3 Supported Backends

The following table presents the capabilities of available backends. Please refer to the specific sections dedicated to each backend to better understand their capabilities and limitations. Choosing the right backend is essential for the success of the deployment.

Table 1: List of available backends

Feature	Memfile	MySQL	PostgreSQL
Status	Stable	Stable	Stable
Data format	CSV file	SQL RMDB	SQL RMDB
Leases	yes	yes	yes
Host reservations	no	yes	yes
Options defined on per host basis	no	yes	yes
Configuration backend	no	yes	yes

4.3.1 Memfile

The memfile backend is able to store lease information, but cannot store host reservation details; these must be stored in the configuration file. (There are no plans to add a host reservations storage capability to this backend.)

No special initialization steps are necessary for the memfile backend. During the first run, both `kea-dhcp4` and `kea-dhcp6` create an empty lease file if one is not present. Necessary disk-write permission is required.

4.3.1.1 Upgrading Memfile Lease Files From an Earlier Version of Kea

There are no special steps required to upgrade memfile lease files between versions of Kea. During startup, the servers check the schema version of the lease files against their own. If there is a mismatch, the servers automatically launch the LFC process to convert the files to the server's schema version. While this mechanism is primarily meant to ease the process of upgrading to newer versions of Kea, it can also be used for downgrading should the need arise. When upgrading, any values not present in the original lease files are assigned appropriate default values. When downgrading, any data present in the files but not in the server's schema are dropped. To convert the files manually prior to starting the servers, run the lease file cleanup (LFC) process. See *The LFC Process* for more information.

4.3.2 MySQL

MySQL is able to store leases, host reservations, options defined on a per-host basis, and a subset of the server configuration parameters (serving as a configuration backend).

4.3.2.1 MySQL 5.7 vs MySQL 8 vs MariaDB 10 and 11

In our Kea performance testing, MySQL 8 shows a 60-90% drop in speed in comparison with MySQL 5.7. Due to the upcoming MySQL 5.7 EOL, we recommend using MariaDB instead of MySQL 8.

MySQL 5.7, MySQL 8, MariaDB 10, and MariaDB 11 are fully compatible, interchangeable, and tested with Kea.

4.3.2.2 First-Time Creation of the MySQL Database

Before preparing any Kea-specific database and tables, the MySQL database must be configured to use the system timezone. It is recommended to use UTC as the timezone for both the system and the MySQL database.

To check the system timezone:

```
date +%Z
```

To check the MySQL timezone:

```
mysql> SELECT @@system_time_zone;
mysql> SELECT @@global.time_zone;
mysql> SELECT @@session.time_zone;
```

To configure the MySQL timezone for a specific server, please refer to the installed version documentation.

Usually the setting is configured in the [mysqld] section in /etc/mysql/my.cnf, /etc/mysql/mysql.cnf, /etc/mysql/mysql.cnf, or /etc/mysql/mysql.conf.d/mysqld.cnf.

```
[mysqld]
# using default-time-zone
default-time-zone='+00:00'

# or using timezone
timezone='UTC'
```

When setting up the MySQL database for the first time, the database area must be created within MySQL, and the MySQL user ID under which Kea will access the database must be set up. This needs to be done manually, rather than via *kea-admin*.

To create the database:

1. Log into MySQL as "root":

```
$ mysql -u root -p
Enter password:
mysql>
```

2. Create the MySQL database:

```
mysql> CREATE DATABASE database_name;
```

(database_name is the name chosen for the database.)

3. Create the user under which Kea will access the database (and give it a password), then grant it access to the database tables:

```
mysql> CREATE USER 'user-name'@'localhost' IDENTIFIED BY '1234';
mysql> GRANT ALL ON database-name.* TO 'user-name'@'localhost';
```

(user-name and password are the user ID and password used to allow Kea access to the MySQL instance. All apostrophes in the command lines above are required.)

4. Create the database.

Exit the MySQL client

```
mysql> quit
Bye
```

Then use the *kea-admin* tool to create the database.

```
$ kea-admin db-init mysql -u database-user -p database-password -n database-
↳name
```

While it is possible to create the database from within the MySQL client, we recommend using the *kea-admin* tool as it performs some necessary validations to ensure Kea can access the database at runtime. Among those checks is verification that the schema does not contain any pre-existing tables; any pre-existing tables must be removed manually. An additional check examines the user's ability to create functions and triggers. The following error indicates that the user does not have the necessary permissions to create functions or triggers:

```
ERROR 1419 (HY000) at line 1: You do not have the SUPER privilege and
↳binary logging is
enabled (you *might* want to use the less safe log_bin_trust_function_
↳creators variable)
ERROR/kea-admin: mysql_can_create cannot trigger, check user permissions,
↳mysql status = 1
mysql: [Warning] Using a password on the command line interface can be
↳insecure.
ERROR/kea-admin: Create failed, the user, keatest, has insufficient
↳privileges.
```

The simplest way around this is to set the global MySQL variable, `log_bin_trust_function_creators`, to 1 via the MySQL client. Note this must be done as a user with SUPER privileges:

```
mysql> set @@global.log_bin_trust_function_creators = 1;
Query OK, 0 rows affected (0.00 sec)
```

To create the database with MySQL directly, follow these steps:

```
mysql> CONNECT database-name;
mysql> SOURCE path-to-kea/share/kea/scripts/mysql/dhcpdb_create.mysql
```

(where path-to-kea is the location where Kea is installed.)

The database may also be dropped manually as follows:

```
mysql> CONNECT database-name;
mysql> SOURCE path-to-kea/share/kea/scripts/mysql/dhcpdb_drop.mysql
```

(where path-to-kea is the location where Kea is installed.)

Warning

Dropping the database results in the unrecoverable loss of any data it contains.

5. Exit MySQL:

```
mysql> quit
Bye
```

If the tables were not created in Step 4, run the *kea-admin* tool to create them now:

```
$ kea-admin db-init mysql -u database-user -p database-password -n database-name
```

Do not do this if the tables were created in Step 4. *kea-admin* implements rudimentary checks; it will refuse to initialize a database that contains any existing tables. To start from scratch, all data must be removed manually. (This process is a manual operation on purpose, to avoid accidentally irretrievable mistakes by *kea-admin*.)

4.3.2.3 Upgrading a MySQL Database From an Earlier Version of Kea

Sometimes a new Kea version uses a newer database schema, so the existing database needs to be upgraded. This can be done using the *kea-admin db-upgrade* command.

To check the current version of the database, use the following command:

```
$ kea-admin db-version mysql -u database-user -p database-password -n database-name
```

(See *Databases and Schema Versions* for a discussion about versioning.) If the version does not match the minimum required for the new version of Kea (as described in the release notes), the database needs to be upgraded.

Before upgrading, please make sure that the database is backed up. The upgrade process does not discard any data, but depending on the nature of the changes, it may be impossible to subsequently downgrade to an earlier version.

To perform an upgrade, issue the following command:

```
$ kea-admin db-upgrade mysql -u database-user -p database-password -n database-name
```

Note

To search host reservations by hostname, it is critical that the collation of the hostname column in the host table be case-insensitive. Fortunately, that is the default in MySQL, but it can be verified via this command:

```
mysql> SELECT COLLATION('');
+-----+
| COLLATION('') |
+-----+
| utf8_general_ci |
+-----+
```

According to mysql's naming convention, when the name ends in *_ci*, the collation is case-insensitive.

4.3.2.4 Improved Performance With MySQL

Changing the MySQL internal value `innodb_flush_log_at_trx_commit` from the default value of 1 to 2 can result in a huge gain in Kea performance. In some deployments, the gain was over 1000% (10 times faster when set to 2, compared to the default value of 1). It can be set per-session for testing:

```
mysql> SET GLOBAL innodb_flush_log_at_trx_commit=2;
mysql> SHOW SESSION VARIABLES LIKE 'innodb_flush_log%';
```

or permanently in `/etc/mysql/my.cnf`:

```
[mysqld]
innodb_flush_log_at_trx_commit=2
```

Be aware that changing this value can cause problems during data recovery after a crash, so we recommend checking the [MySQL documentation](#). With the default value of 1, MySQL writes changes to disk after every INSERT or UPDATE query (in Kea terms, every time a client gets a new lease or renews an existing lease). When `innodb_flush_log_at_trx_commit` is set to 2, MySQL writes the changes at intervals no longer than 1 second. Batching writes gives a substantial performance boost. The trade-off, however, is that in the worst-case scenario, all changes in the last second before crash could be lost. Given the fact that Kea is stable software and crashes very rarely, most deployments find it a beneficial trade-off.

4.3.3 PostgreSQL

PostgreSQL can store leases, host reservations, and options defined on a per-host basis.

4.3.3.1 First-Time Creation of the PostgreSQL Database

Before preparing any Kea-specific database and tables, the PostgreSQL database must be configured to use the system timezone. It is recommended to use UTC as the timezone for both the system and the PostgreSQL database.

To check the system timezone:

```
date +%Z
```

To check the PostgreSQL timezone:

```
postgres=# show timezone;
postgres=# SELECT * FROM pg_timezone_names WHERE name = current_setting(
↪ 'TIMEZONE');
```

To configure the PostgreSQL timezone for a specific server, please refer to the installed version documentation.

Usually the setting is configured in the `postgresql.conf` with the varying version path `/etc/postgresql/<version>/main/postgresql.conf`, but on some systems the files may be located in `/var/lib/pgsql/data`.

```
timezone = 'UTC'
```

The first task is to create both the database and the user under which the servers will access it. A number of steps are required:

1. Log into PostgreSQL as "postgres":

```
$ sudo -u postgres psql postgres
Enter password:
postgres=#
```

2. Create the database:


```
postgres=# CREATE DATABASE database-name;
CREATE DATABASE
postgres=#
```

(database-name is the name chosen for the database.)

3. Create the user under which Kea will access the database (and give it a password), then grant it access to the database:

```
postgres=# CREATE USER user-name WITH PASSWORD '1234';
CREATE ROLE
postgres=# GRANT ALL PRIVILEGES ON DATABASE database-name TO user-name;
GRANT
postgres=# \c database-name
You are now connected to database "database-name" as user "postgres".
postgres=# GRANT ALL PRIVILEGES ON SCHEMA public TO user-name;
GRANT
postgres=#
```

4. Exit PostgreSQL:

```
postgres=# \q
Bye
$
```

5. At this point, create the database tables either using the [kea-admin](#) tool, as explained in the next section (recommended), or manually. To create the tables manually, enter the following command. PostgreSQL will prompt the administrator to enter the new user's password that was specified in Step 3. When the command completes, Kea will return to the shell prompt. The output should be similar to the following:

```
$ psql -d database-name -U user-name -f path-to-kea/share/kea/scripts/pgsql/dhcpdb_
→create.pgsql
Password for user user-name:
CREATE TABLE
CREATE INDEX
CREATE INDEX
CREATE TABLE
CREATE INDEX
CREATE TABLE
START TRANSACTION
INSERT 0 1
INSERT 0 1
INSERT 0 1
COMMIT
CREATE TABLE
START TRANSACTION
INSERT 0 1
COMMIT
$
```

(path-to-kea is the location where Kea is installed.)

If instead an error is encountered, such as:

```
psql: FATAL: no pg_hba.conf entry for host "[local]", user "user-name", database
↳ "database-name", SSL off
```

... the PostgreSQL configuration will need to be altered. Kea uses password authentication when connecting to the database and must have the appropriate entries added to PostgreSQL's `pg_hba.conf` file. This file is normally located in the primary data directory for the PostgreSQL server. The precise path may vary depending on the operating system and version, but the default location for PostgreSQL is `/etc/postgresql/*/main/postgresql.conf`. However, on some systems, the file may reside in `/var/lib/pgsql/data`.

Assuming Kea is running on the same host as PostgreSQL, adding lines similar to the following should be sufficient to provide password-authenticated access to Kea's database:

local	database-name	user-name		password
host	database-name	user-name	127.0.0.1/32	password
host	database-name	user-name	::1/128	password

These edits are primarily intended as a starting point, and are not a definitive reference on PostgreSQL administration or database security. Please consult the PostgreSQL user manual before making these changes, as they may expose other databases that are running. It may be necessary to restart PostgreSQL for the changes to take effect.

4.3.3.2 Initialize the PostgreSQL Database Using `kea-admin`

If the tables were not created manually, do so now by running the `kea-admin` tool:

```
$ kea-admin db-init psql -u database-user -p database-password -n database-name
```

Do not do this if the tables were already created manually. `kea-admin` implements rudimentary checks; it will refuse to initialize a database that contains any existing tables. To start from scratch, all data must be removed manually. (This process is a manual operation on purpose, to avoid accidentally irretrievable mistakes by `kea-admin`.)

4.3.3.3 Upgrading a PostgreSQL Engine From an Earlier Version

If you upgraded your PostgreSQL from a version prior to 15.0, you need to grant additional privileges to the user:

First, log into PostgreSQL as "postgres":

```
$ sudo -u postgres psql -d database-name -U postgres
Enter password:
postgres=#
```

Next, grant the access to the public schema.

```
postgres=# GRANT ALL PRIVILEGES ON SCHEMA public TO user-name;
GRANT
postgres=#
```

Now, quit the PostgreSQL client:

```
postgres=# \q
Bye
$
```

4.3.3.4 Upgrading a PostgreSQL Database From an Earlier Version of Kea

The PostgreSQL database schema can be upgraded using the same tool and commands as described in *Upgrading a MySQL Database From an Earlier Version of Kea*, with the exception that the "pgsql" database backend type must be used in the commands.

Use the following command to check the current schema version:

```
$ kea-admin db-version pgsql -u database-user -p database-password -n database-name
```

Use the following command to perform an upgrade:

```
$ kea-admin db-upgrade pgsql -u database-user -p database-password -n database-name
```

4.3.3.5 Improved Performance With PostgreSQL

Changing the PostgreSQL internal value `synchronous_commit` from the default value of ON to OFF can result in significant gains in Kea performance; on slow systems, the gain can be over 1000%. It can be set per-session for testing:

```
postgres=# SET synchronous_commit = OFF;
```

or permanently via command (preferred method):

```
postgres=# ALTER SYSTEM SET synchronous_commit=OFF;
```

or permanently in `/etc/postgresql/[version]/main/postgresql.conf`:

```
synchronous_commit = off
```

Changing this value can cause problems during data recovery after a crash, so we recommend a careful read of the [PostgreSQL documentation](#). With the default value of ON, PostgreSQL writes changes to disk after every INSERT or UPDATE query (in Kea terms, every time a client gets a new lease or renews an existing lease). When `synchronous_commit` is set to OFF, PostgreSQL adds some delay before writing the changes. Batching writes gives a substantial performance boost, but in the worst-case scenario, all changes in the last moment before a crash could be lost. Since Kea is stable software and crashes very rarely, most deployments find the performance benefits outweigh the potential risks.

4.3.4 Using Read-Only Databases With Host Reservations

If a read-only database is used for storing host reservations, Kea must be explicitly configured to operate on the database in read-only mode. Sections *Using Read-Only Databases for Host Reservations With DHCPv4* and *Using Read-Only Databases for Host Reservations with DHCPv6* describe when such a configuration may be required, and how to configure Kea to operate in this way for both DHCPv4 and DHCPv6.

4.3.5 Limitations Related to the Use of SQL Databases

4.3.5.1 Year 2038 Issue

The lease expiration time in Kea is stored in the SQL database for each lease as a timestamp value. Kea developers have observed that the MySQL database does not accept timestamps beyond 2147483647 seconds (the maximum signed 32-bit number) from the beginning of the UNIX epoch (00:00:00 on 1 January 1970). Some versions of PostgreSQL do accept greater values, but the value is altered when it is read back. For this reason, the lease database backends put a restriction on the maximum timestamp to be stored in the database, which is equal to the maximum signed 32-bit number. This effectively means that the current Kea version cannot store leases whose expiration time is later than 2147483647 seconds since the beginning of the epoch (around the year 2038). This will be fixed when database support for longer timestamps is available.

KEA SECURITY

Kea was originally designed to be installed in a protected environment, in a network datacenter; it did not offer hardened security features. However, due to customer demand and evolving network requirements, support for basic HTTP authentication and Transport Layer Security (TLS) have been added to Kea.

5.1 TLS/HTTPS Support

Since Kea 1.9.6, TLS can be used to secure HTTP communication. There are three levels of protection possible:

- No TLS. The connection is plain-text, unencrypted HTTP. (This is the only option available in versions prior to Kea 1.9.6.)
- Encryption, which protects against passive attacks and eavesdropping. In this case, the server is authenticated but the client is not. This is the typical mode when securing a website, where clients and servers are not under the control of a common entity.
- Mutual authentication between the client and the server. This is the strictest security mode and is the default when TLS is enabled.

Note

TLS mutual authentication is for TLS entities only. When TLS and an HTTP authentication scheme are used together, there is no binding between the two security mechanisms, and therefore no proof that the TLS client and server are the same as the HTTP authentication client and server.

Note

It is recommend to use privileged ports for HTTP/HTTPS against local attacks (by users which are connected to the box where Kea servers/agents run). This measure also prevents against impersonation with HTTP, and Denial of Service in general.

Note

The server will issue an error when changing the socket type from HTTP to HTTPS or from HTTPS to HTTP using the same address and port. This action is not allowed as it might introduce a security issue accidentally caused by a user mistake. A different address or port must be specified when using the "config-set" command to switch from HTTP to HTTPS or from HTTPS to HTTP. The same applies when modifying the configuration file and then running "config-reload" command.

5.1.1 Building Kea with TLS/HTTPS Support

TLS/HTTPS support is available with either the OpenSSL or the Botan cryptographic library. There are some constraints on the Boost library that must be used:

- OpenSSL versions older than 1.0.2 are obsolete and should not be used. Kea TLS support has not been tested with and is not supported on these versions.
- OpenSSL version 1.0.2 has extended support, but only for OpenSSL premium customers. Kea TLS support has been tested but is not supported on this version.
- OpenSSL versions 1.1.x and later have been tested and are supported. Many recent operating system versions include TLS 1.3 support.
- OpenSSL 3.x has been released and Kea will build with it.
- LibreSSL 3.2.4 has been tested. LibreSSL shares the OpenSSL 1.0.2 API, so it should work, but is not supported.
- Botan 1.x versions are obsolete and must not be used.
- Botan versions 2.14.0 and later have been tested and are supported. Kea TLS support requires the four Asio header files which are included in Botan packages and which are installed only if Botan is configured with the `--with-boost` option.

Many packages provided by operating systems, such as Ubuntu 20.10, do not build Botan with Boost support, making those packages unusable for Kea with TLS.

It is still possible to take these files from the corresponding Botan distribution and install them manually in the Botan include directory, but this should be a last-resort procedure.

Without these header files, or with a Botan version prior to 2.14.0, Kea cannot build as the TLS/HTTPS support is considered essential for security.

- Very old Boost versions provide SSL support (based on OpenSSL) without offering a choice of the TLS version; Kea can still use them, but they are not recommended.
- Boost versions prior to 1.64 provide SSL support with a fixed choice of the TLS version; Kea enforces the use of TLS 1.2 with them.
- Boost versions 1.64 or newer provide SSL support with a generic TLS version; the best (highest) version available on both peers is selected.

5.1.2 TLS/HTTPS Configuration

The TLS configuration parameters are:

- `trust-anchor` - this string parameter specifies the name of a file or directory where the certification authority (CA) certificate of the other peer can be found. With OpenSSL, the directory must include hash symbolic links. With Botan, the directory is recursively searched for certificates.
- `cert-file` - this string parameter specifies the name of the file containing the end-entity certificate of the Kea instance being configured.
- `key-file` - this string parameter specifies the private key of the end-entity certificate of the Kea instance being configured. The file must not be encrypted; it is highly recommended to restrict its access.

The three string parameters must be either all unspecified (TLS disabled) or all specified (TLS enabled).

TLS is asymmetric: the authentication of the server by the client is mandatory but the authentication of the client by the server is optional. In TLS terms, this means the server may require the client certificate, or may not; there is a server-specific TLS parameter.

- `cert-required` - this boolean parameter allows a server to not require the client certificate. Its default value is `true`, which means the client certificate is required and the client must be authenticated. This flag has no meaning on the client side; the server always provides a certificate which is validated by the client.

Objects in files must be in the PEM format. Files can contain more than one certificate, but this has not been tested and is not supported.

Botan requires CA certificates to have the standard CA certificate attributes, verifies that end-entity certificates are version 3 certificates (as required by the TLS standard), and supports only PKCS 8 files for the private key.

Note

Some cryptographic libraries (e.g. Botan and recent OpenSSL) enforce minimal strength (i.e. key length), e.g. at least 2048 for RSA.

A sample set of certificates and associated objects is available at `src/lib/asiolink/testutils/ca` in the Kea sources, with a `doc.txt` file explaining how they were generated using the `openssl` command. These files are for testing purposes only. **Do not use them in production.**

TLS handshake, the phase where the cryptographic parameters are exchanged and authentication is verified, can fail in multiple ways. Error messages often do not help to pinpoint the source of the problem. Both OpenSSL and Botan provide a command-line tool with a `verify` command which can be used to understand and fix handshake issues.

5.1.3 OpenSSL Tuning

Kea accepts the default OpenSSL configuration parameters, but administrators can also fine-tune the OpenSSL settings. For example, it may be desirable to limit the TLS version.

The default OpenSSL configuration file is named `openssl.cnf`. It can be found in a system-dependent `etc` directory, and the location can be overridden using the `OPENSSL_CONF` environment variable. For OpenSSL versions greater than 1.0.2, the minimum acceptable protocol can be set via the `MinProtocol` variable.

For these examples, we assume that no variables are already set and no sections already exist; it is, of course, possible to reuse existing variables and sections.

In the default application, `openssl_conf`, the corresponding variable must be set to the name of the section that handles defaults: in this example, `default_conf`. If `openssl_conf` is not yet set, add this command at the beginning of the OpenSSL configuration file (before the first section):

```
openssl_conf = default_conf
```

In the `default_conf` section, the `ssl_conf` variable must be set to the name of the section that handles SSL/TLS defaults: in this example, `ssl_sect`.

```
[ default_conf ]
ssl_conf = ssl_sect
```

In the `ssl_sect` section, the `system_default` variable must be set to the name of the section that handles system defaults: in this example, `system_default_sect`.

```
[ ssl_sect ]
system_default = system_default_sect
```

In the `system_default_sect` section, the `MinProtocol` variable must be set to the desired minimal SSL/TLS version: in this example, `TLSv1.2`.


```
[ system_default_sect ]  
MinProtocol = TLSv1.2
```

The same steps can be used to enforce other crypto parameters if desired.

It is highly recommended to read the `openssl.cnf` manual page, normally called `config.5ssl` and displayed using `man config`.

5.2 Secure Kea Control Agent

The Kea Control Agent natively supports secure HTTP connections using TLS. This allows protection against users from the node where the agent runs, something that a reverse proxy cannot provide. More about TLS/HTTPS support in Kea can be found in [TLS/HTTPS Support](#).

TLS is configured using three string parameters with file names, and a boolean parameter:

- The `trust-anchor` specifies the Certification Authority file name or directory path.
- The `cert-file` specifies the server certificate file name.
- The `key-file` specifies the private key file name. The file must not be encrypted.
- The `cert-required` specifies whether client certificates are required or optional. The default is to require them and to perform mutual authentication.

The file format is PEM. Either all the string parameters are specified and HTTP over TLS (HTTPS) is used, or none is specified and plain HTTP is used. Configuring only one or two string parameters results in an error.

Note

When client certificates are not required, only the server side is authenticated, i.e. the communication is encrypted with an unknown client. This protects only against passive attacks; active attacks, such as "man-in-the-middle," are still possible.

Note

No standard HTTP authentication scheme cryptographically binds its end entity with TLS. This means that the TLS client and server can be mutually authenticated, but there is no proof they are the same as for the HTTP authentication.

The `kea-shell` tool also supports TLS.

5.3 Securing a Kea Deployment

Below is a list of considerations for administrators wishing to improve Kea's security. In many cases, there are trade-offs between convenience and security.

5.3.1 Component-Based Design

The Kea architecture is modular, with separate daemons for separate tasks. A Kea deployment may include DHCPv4, DHCPv6, and Dynamic DNS daemons; a Control Agent daemon run on each application server; the `kea-lfc` utility for doing periodic lease file cleanup; MySQL and or PostgreSQL databases, run either locally on the application servers or accessed over the internal network; a Netconf daemon to perform config and stats monitoring of Kea servers; and a

Stork monitoring system. This modular architecture allows the administrator to minimize the attack surface by minimizing the code that is loaded and running. For example, *kea-dhcp-ddns* should not be run unless DNS updates are required. Similarly, *kea-lfc* is never triggered (and can be safely removed or never installed) if memfile is not used. Potential Kea security issues can be minimized by running only those processes required in the local environment.

Note

As of Kea 2.7.9, *data-directory* (DHCPv6 only) is deprecated. The lease files (DHCPv4 and DHCPv6) and duid file (DHCPv6 only) may only be loaded from the directory determined at compilation: "[*kea-install-dir*]/var/lib/kea". This path may be overridden at startup by setting the environment variable *KEA_DHCP_DATA_DIR* to the desired path. If a path other than this value is used in *name* or *data-directory*, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. This restriction applies to writing lease file using *lease4-write* and *lease6-write* commands. If a path other than this value is used in *filename*, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in specifying a custom file name simply omit the path portion from *filename*. Same restriction also applies to writing cache file using *cache-write* command. For ease of use in specifying a custom file name simply omit the path portion from *filename*.

5.3.2 Limiting Application Permissions

The DHCPv4 and DHCPv6 protocols assume the server opens privileged UDP port 67 (DHCPv4) or 547 (DHCPv6), which requires root access under normal circumstances. However, via the capabilities mechanism on Linux systems, Kea can run from an unprivileged account. See *Running Kea From a Non-root Account on Linux* for details on how to run Kea without root access.

The Control Agent (CA) can accept incoming HTTP or HTTPS connections. The default port is 8000, which does not require privileged access.

5.3.3 Securing Kea Administrative Access

The three primary Kea daemons (*kea-dhcp4*, *kea-dhcp6* and *kea-dhcp-ddns*) all support a control channel, which is implemented as a UNIX socket. The control channel, which opens a UNIX socket, is disabled by default; however, many configuration examples have it enabled, as it is a very popular feature. To read from or write to this socket, root access is generally required, although if Kea is configured to run as non-root, the owner of the process can write to it. Access can be controlled using normal file-access control on POSIX systems (owner, group, others, read/write).

Note

As of Kea 2.7.9, control sockets may only reside in the directory determined during compilation as "[*kea-install-dir*]/var/run/kea", which must also have *0750* access rights. This path may be overridden at startup by setting the environment variable *KEA_CONTROL_SOCKET_DIR* to the desired path. If a path other than this value is used in *socket-name*, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in simply omit the path component from *socket-name*.

Since Kea version 2.7.2 DHCP servers support HTTP/HTTPS control channels so the Control Agent (CA) is no longer needed.

Since Kea-2.7.6 Kea supports multiple HTTP/HTTPS connections. Both IPv4 and IPv6 addresses can be used. Security can be enhanced if configuring HTTPS connections for all daemons.

Kea configuration is controlled by a JSON file on the Kea server. This file can be viewed or edited by anyone with file permissions (which are controlled by the operating system). Note that passwords are stored in clear text in the configuration file, so anyone with access to read the configuration file can find this information. As a practical matter,

anyone with permission to edit the configuration file has control over Kea. Limiting user permission to read or write the Kea configuration file is an important security step.

Note

As of Kea 2.7.9, the config file may only be written (using the `config-write` command) to the same directory as the config file used when starting Kea (passed as a `-c` argument).

5.3.4 Securing Database Connections

Kea can use an external MySQL or PostgreSQL database to store configuration, host reservations, or/and leases, or/and for forensic logging. The use of databases is a popular feature, but it is optional; it is also possible to store data in a flat file on disk.

When using a database, Kea stores and uses the following credentials to authenticate with the database: username, password, host, port, and database name. **These are stored in clear text in the configuration file.**

Depending on the database configuration, it is also possible to verify whether the system user matches the database username. Consult the MySQL or PostgreSQL manual for details.

Kea supports client TLS settings for MySQL database and it must be configured explicitly for all used connections (configuration, reservations, leases, forensic logging).

5.3.5 Information Leakage Through Logging

It is possible for Kea to log an entire configuration file, including passwords and secrets. Since Kea 1.9.7, this issue has been resolved by replacing the value of all entries ending in `password` or `secret` with asterisks, as was already done for database logs.

Logs are sent to stdout, stderr, files, or syslog; system file permissions system apply to stdout/stderr and files. Syslog may export the logs over the network, exposing them further to possible snooping.

Note

As of Kea 2.7.9, log files may only be written to the output directory determined during compilation as: "[kea-install-dir]/var/log/kea". This path may be overridden at startup by setting the environment variable `KEA_LOG_FILE_DIR` to the desired path. If a path other than this value is used in `output`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use simply omit the path component from `output` and specify only the file name.

Note

As of Kea 2.7.9, legal log files may only be written to the output directory determined during compilation as: "[kea-install-dir]/var/log/kea". This path may be overridden at startup by setting the environment variable `KEA_LEGAL_LOG_DIR` to the desired path. If a path other than this value is used in `path`, Kea will emit an error and refuse to start or if already running, log an unrecoverable error. For ease of use simply omit the `path` parameter.

5.3.6 Summary of Path Restrictions

Path restrictions mentioned through this section can be summarized according to the following table:

Restricted Element	Default Value	Environment Variable Override
Config Files (<code>config-write</code>)	Same Directory as Initial Config File	N/A
Host Cache Files (<code>cache-write</code>)	<code>var/lib/kea</code>	<code>KEA_DHCP_DATA_DIR</code>
Lease Files	<code>var/lib/kea</code>	<code>KEA_DHCP_DATA_DIR</code>
Log Files	<code>var/log/kea</code>	<code>KEA_LOG_FILE_DIR</code>
Forensic Log Files	<code>var/log/kea</code>	<code>KEA_LEGAL_LOG_DIR</code>
Unix Sockets	<code>var/run/kea</code>	<code>KEA_CONTROL_SOCKET_DIR</code>
Scripts used by hook libraries	<code>share/kea/scripts/</code>	<code>KEA_HOOK_SCRIPTS_PATH</code>

5.3.7 Cryptography Components

Kea supports the use of either of two cryptographic libraries: Botan or OpenSSL. The choice is made at compile time, and creates both compile and runtime dependencies between the Kea and the selected library. While OpenSSL is the most popular choice for deployments, Botan remains a fully supported alternative.

The primary use cases for the cryptographic libraries are:

- TLS support for the Control Agent (CA), introduced in Kea 1.9.6.
- TSIG signatures when sending DNS updates.
- calculating DHCID records when sending DNS updates.
- random number generation (but not for usage requiring a crypto grade generator).

For OpenSSL and Botan, only the low-level crypto interface is used (e.g. `libcrypto`). Kea does not link with `libssl`. Some dependent software systems, such as database client libraries, can also depend on a crypto library.

One way to limit exposure for potential OpenSSL or Botan vulnerabilities is not to use DDNS. The libraries would still be needed to build and run Kea, but the code would never be used, so any potential bugs in the libraries would not be exploitable.

5.3.8 TSIG Signatures

Kea supports the following algorithms when signing DNS updates with TSIG signatures:

- HMAC-MD5
- HMAC-SHA1
- HMAC-SHA224
- HMAC-SHA256
- HMAC-SHA384
- HMAC-SHA512

See [TSIG Key List](#) for an up-to-date list.

Kea uses SHA256 to calculate DHCID records. This is irrelevant from the cryptography perspective, as the DHCID record is only used to generate unique identifiers for two devices that may have been assigned the same IP address at different times.

5.3.9 Raw Socket Support

In principle, Kea DHCPv4 uses raw sockets to receive traffic from clients. The difficulty is with receiving packets from devices that do not yet have an IPv4 address. When dealing with direct traffic (where both client and server are connected to the same link, not separated by relays), the kernel normally drops the packet as the source IP address is 0.0.0.0. Therefore, Kea needs to open raw sockets to be able to receive this traffic.

However, this is not necessary if all the traffic is coming via relays, which is often the case in many networks. In that case normal UDP sockets can be used instead. There is a `dhcp-socket-type` parameter that controls this behavior.

The default is to permit raw socket usage, as it is more versatile.

When using raw sockets, Kea is able to receive raw layer 2 packets, bypassing most firewalls (including iptables). This effectively means that when raw sockets are used, the iptables cannot be used to block DHCP traffic. This is a design choice of the Linux kernel.

Kea can be switched to use UDP sockets. This is an option when all traffic is relayed. However, it does not work for directly connected devices. If Kea is limited to UDP sockets, iptables should work properly.

If raw sockets are not required, disabling this access can improve security.

5.3.10 Remote Administrative Access

Kea's Control Agent (CA) exposes a RESTful API over HTTP or HTTPS (HTTP over TLS). The CA is an optional feature that is disabled by default, but it is very popular. When enabled, it listens on the loopback address (127.0.0.1 or ::1) by default, unless configured otherwise. See [TLS/HTTPS Support](#) for information about protecting the TLS traffic. Limiting the incoming connections with a firewall, such as iptables, is generally a good idea.

Note that in High Availability (HA) deployments, DHCP partners connect to each other using a CA connection.

Since Kea version 2.7.2 DHCP and DDNS servers support HTTP/HTTPS control channels so the Control Agent (CA) is no longer needed.

Since Kea-2.7.6 Kea supports multiple HTTP/HTTPS connections. Both IPv4 and IPv6 addresses can be used. Security can be enhanced if configuring HTTPS connections for all daemons.

5.3.11 Authentication for Kea's RESTful API

Kea 1.9.0 added support for basic HTTP authentication ([RFC 7617](#)), to control access for incoming REST commands over HTTP. The credentials (username, password) are stored in a local Kea configuration file on disk. The username is logged with the API command, so it is possible to determine which authenticated user performed each command. The access control details are logged using a dedicated auth logger. Basic HTTP authentication is weak on its own as there are known dictionary attacks, but those attacks require a "man in the middle" to get access to the HTTP traffic. That can be eliminated by using basic HTTP authentication exclusively over TLS. In fact, if possible, using client certificates for TLS is better than using basic HTTP authentication.

Kea 1.9.2 introduced a new auth hook point. With this new hook point, it is possible to develop an external hook library to extend the access controls, integrate with another authentication authority, or add role-based access control to the Control Agent. This hookpoint was renamed as `http_auth` and is also supported by the DHCP and DDNS servers since Kea version 2.7.2.

5.4 Kea Security Processes

The following sections discuss how the Kea DHCP development team ensures code quality and handles vulnerabilities.

5.4.1 Vulnerability Handling

ISC is an experienced and active participant in the industry-standard vulnerability disclosure process and maintains accurate documentation on our process and vulnerabilities in ISC software. See <https://kb.isc.org/docs/aa-00861> for ISC's Software Defect and Security Vulnerability Disclosure Policy.

In case of a security vulnerability in Kea, ISC notifies support customers ahead of any public disclosure, and provides a patch and/or updated installer package to remediate the vulnerability.

When a security update is published, both the source tarballs and the ISC-maintained packages are published on the same day. This enables users of the native Linux update mechanisms (such as Debian's and Ubuntu's apt or RedHat's dnf) to update their systems promptly.

5.4.2 Code Quality and Testing

Kea undergoes extensive tests during its development. The following are some of the processes that are used to ensure adequate code quality:

- Each line of code goes through a formal review before it is accepted. The review process is documented and available publicly.
- Roughly 50% of the source code is dedicated to unit tests. As of May 2024, there were over 12000 unit tests and the number is increasing with time. Unit tests are required to commit any new feature.
- There are around 2000 system tests for Kea. These simulate both correct and invalid situations, covering network packets (mostly DHCP, but also DNS, HTTP, HTTPS and others), command-line usage, API calls, database interactions, scripts, and more.
- There are performance tests with over 80 scenarios that test Kea overall performance and resiliency to various levels of traffic, and measuring various metrics (latency, leases per seconds, packets per seconds, CPU usage, memory utilization, and others).
- Kea uses Continuous Integration (CI). This means that the great majority of tests (all unit and system tests, and in some cases also performance tests) are run for every commit. Many "lighter" tests are run on branches, before the code is even accepted.
- Many unit and system tests check for negative scenarios, such as incomplete, broken, or truncated packets, API commands, and configuration files, as well as incorrect sequences (such as sending packets in an invalid order) and more.
- The Kea development team uses many tools that perform automatic code quality checks, such as danger, as well as internally developed sanity checkers.
- The Kea team uses the following static code analyzers: Coverity Scan, cppcheck, clang-static-analyzer, shellcheck, flawfinder, semgrep and danger.
- The Kea team uses the following dynamic code analyzers: Valgrind, Thread Sanitizer (TSAN), Address Sanitizer (ASAN), Undefined Behavior Sanitizer (UBSAN).

5.4.3 Fuzz Testing

The Kea team has a process for running fuzz testing. Fuzzing is a software-testing technique whereby a program is presented with a variety of generated data as input and is monitored for abnormal conditions such as crashes or hangs.

There are two ways to fuzz Kea.

Option 1. With the libfuzzer harness function LLVMFuzzerTestOneInput.

Option 2. With the AFL (American Fuzzy Lop <https://github.com/google/AFL>) compiler.

Using the LLVMFuzzerTestOneInput Harness Function:

This mode of fuzzing works with virtually any compiler.

There are four types of fuzzers implemented with this mode:

- Config fuzzer
- HTTP endpoint fuzzer
- Packet fuzzer
- Unix socket fuzzer

There are two binaries under test:

- *kea-dhcp4*
- *kea-dhcp6*

5.4.4 Release Integrity

All ISC software releases are signed with PGP and distributed via the ISC website, which is itself DNSSEC-signed, so users can be confident the software has not been tampered with.

5.4.5 Bus Factor

According to the [Core Infrastructure project](#), a "bus factor" or "truck factor" is the minimum number of project members that have to suddenly disappear from a project ("be hit by a bus") before the project stalls due to lack of knowledgeable or competent personnel. It is hard to estimate precisely, but the bus factor for Kea is somewhere around five. As of 2021, there are six core developers and two quality assurance engineers, with many additional casual contributors (product manager, support team, IT, etc.). The team is geographically dispersed.

KEA CONFIGURATION

Kea uses JSON structures to represent server configurations. The following sections describe how the configuration structures are organized.

6.1 JSON Configuration

JSON is the notation used throughout the Kea project. The most obvious usage is for the configuration file, but JSON is also used for sending commands over the Management API (see *Management API*) and for communicating between DHCP servers and the DDNS update daemon.

Typical usage assumes that the servers are started from the command line, either directly or using a script, e.g. *keactrl*. The configuration file is specified upon startup using the `-c` parameter.

6.1.1 JSON Syntax

Configuration files for the DHCPv4, DHCPv6, DDNS, Control Agent, and NETCONF modules are defined in an extended JSON format. Basic JSON is defined in [RFC 7159](#) and [ECMA 404](#). In particular, the only boolean values allowed are true or false (all lowercase). The capitalized versions (True or False) are not accepted.

Even though the JSON standard (ECMA 404) does not require JSON objects (i.e. name/value maps) to have unique entries, Kea implements them using a C++ STL map with unique entries. Therefore, if there are multiple values for the same name in an object/map, the last value overwrites previous values. Since Kea 1.9.0, configuration file parsers raise a syntax error in such cases.

Kea components use extended JSON with additional features allowed:

- Shell comments: any text after the hash (#) character is ignored.
- C comments: any text after the double slashes (//) character is ignored.
- Multiline comments: any text between /* and */ is ignored. This comment can span multiple lines.
- File inclusion: JSON files can include other JSON files by using a statement of the form `<?include "file.json"?>`.
- Extra commas: to remove the inconvenience of errors caused by leftover commas after making changes to configuration. While parsing, a warning is printed with the location of the comma to give the user the ability to correct a potential mistake.

Warning

These features are meant to be used in a JSON configuration file. Their usage in any other way may result in errors.

The configuration file consists of a single object (often colloquially called a map) started with a curly bracket. It comprises only one of the "Dhcp4", "Dhcp6", "DhcpDdns", "Control-agent", or "Netconf" objects. It is possible to define additional elements but they will be ignored.

A very simple configuration for DHCPv4 could look like this:

```
# The whole configuration starts here.
{
    # DHCPv4 specific configuration starts here.
    "Dhcp4": {
        "interfaces-config": {
            "interfaces": [ "eth0" ],
            "dhcp-socket-type": "raw"
        },
        "valid-lifetime": 4000,
        "renew-timer": 1000,
        "rebind-timer": 2000,
        "subnet4": [{
            "pools": [ { "pool": "192.0.2.1-192.0.2.200" } ],
            "subnet": "192.0.2.0/24",
            "id": 1
        }],
        # Now loggers are inside the DHCPv4 object.
        "loggers": [{
            "name": "*",
            "severity": "DEBUG"
        }]
    }
}

# The whole configuration structure ends here.
}
```

More examples are available in the installed `share/doc/kea/examples` directory.

To avoid repetition of mostly similar structures, examples in the rest of this guide will showcase only the subset of parameters appropriate for a given context. For example, when discussing the IPv6 subnets configuration in DHCPv6, only `subnet6` parameters will be mentioned. It is implied that the remaining elements (the global map that holds `Dhcp6`) are present, but they are omitted for clarity. Usually, locations where extra parameters may appear are denoted by an ellipsis (...).

6.1.2 Comments and User Context

Shell, C, or C++ style comments are all permitted in the JSON configuration file if the file is used locally. This is convenient and works in simple cases where the configuration is kept statically using a local file. However, since comments are not part of JSON syntax, most JSON tools detect them as errors. Another problem with them is that once Kea loads its configuration, the shell, C, and C++ style comments are ignored. If commands such as `config-get` or `config-write` are used, those comments are lost. An example of such comments was presented in the previous section.

Historically, to address the problem, Kea code allowed the use of *comment* strings as valid JSON entities. This had the benefit of being retained through various operations (such as `config-get`), or allowing processing by JSON tools. An example JSON comment looks like this:

```
"Dhcp4": {
  "subnet4": [{
    "id": 1,
    "subnet": "192.0.2.0/24",
    "pools": [{ "pool": "192.0.2.10 - 192.0.2.20" }],
    "comment": "second floor"
  }]
}
```

However, the facts that the comment could only be a single line, and that it was not possible to add any other information in a more structured form, were frustrating. One specific example was a request to add floor levels and building numbers to subnets. This was one of the reasons why the concept of user context was introduced. It allows adding an arbitrary JSON structure to most Kea configuration structures.

This has a number of benefits compared to earlier approaches. First, it is fully compatible with JSON tools and Kea commands. Second, it allows storing simple comment strings, but it can also store much more complex data, such as multiple lines (as a string array), extra typed data (such as floor numbers being actual numbers), and more. Third, the data is exposed to hooks, so it is possible to develop third-party hooks that take advantage of that extra information. An example user context looks like this:

```
"Dhcp4": {
  "subnet4": [{
    "id": 1,
    "subnet": "192.0.2.0/24",
    "pools": [{ "pool": "192.0.2.10 - 192.0.2.20" }],
    "user-context": {
      "comment": "second floor",
      "floor": 2
    }
  }]
}
```

User contexts can store an arbitrary data file as long as it has valid JSON syntax and its top-level element is a map (i.e. the data must be enclosed in curly brackets). However, some hook libraries may expect specific formatting; please consult the specific hook library documentation for details.

The user-context mechanism has superseded the JSON comment capabilities; ISC encourages administrators to use user context instead of the older mechanisms. To promote this way of storing comments, Kea converts JSON comments to user context on the fly.

However, if the configuration uses the old JSON comment method, the `config-get` command returns a slightly modified configuration. It is not uncommon for a call for `config-set` followed by `config-get` to receive a slightly different structure. The best way to avoid this problem is simply to abandon JSON comments and use user context.

Kea supports user contexts at the following levels: global scope, interfaces configuration, multi-threading, shared networks, subnets, client classes, option data and definitions, pools, prefix delegation pools, host reservations, control socket, HTTP header, authentication, clients, queue control, DHCP-DDNS, loggers, leases, and server ID. These are supported in both DHCPv4 and DHCPv6, with the exception of prefix delegation pools and server ID, which are DHCPv6 only.

User context can be added and edited in structures supported by commands.

We encourage Kea users to utilize these functions to store information used by other systems and custom hooks.

For example, the `subnet4-update` command can be used to add user context data to an existing subnet.

```
{
  "subnet4": [ {
    "id": 1,
    "subnet": "10.20.30.0/24",
    "user-context": {
      "building": "Main",
      "floor": 1
    }
  } ]
}
```

The same can be done with many other commands, like *lease6-add*, etc.

Kea also uses user context to store non-standard data. Currently, only *Storing Extended Lease Information* uses this feature.

When enabled, it adds the ISC key in user-context to differentiate automatically added content.

Example of relay information stored in a lease:

```
{
  "arguments": {
    "client-id": "42:42:42:42:42:42:42:42",
    "cltt": 12345678,
    "fqdn-fwd": false,
    "fqdn-rev": true,
    "hostname": "myhost.example.com.",
    "hw-address": "08:08:08:08:08:08",
    "ip-address": "192.0.2.1",
    "state": 0,
    "subnet-id": 44,
    "valid-lft": 3600,
    "user-context": {
      "ISC": {
        "relays": [
          {
            "hop": 2,
            "link": "2001:db8::1",
            "peer": "2001:db8::2"
          },
          {
            "hop": 1,
            "link": "2001:db8::3",
            "options": "0x00C800080102030405060708",
            "peer": "2001:db8::4"
          }
        ]
      }
    }
  }
}
```

User context can store configuration for multiple hooks and comments at once.

For a discussion about user context used in hooks, see *User Contexts in Hooks*.

6.1.3 Simplified Notation

It is sometimes convenient to refer to a specific element in the configuration hierarchy. Each hierarchy level is separated by a slash. If there is an array, a specific instance within that array is referenced by a number in square brackets (with numbering starting at zero). For example, in the above configuration the valid-lifetime in the Dhcp4 component can be referred to as Dhcp4/valid-lifetime, and the pool in the first subnet defined in the DHCPv4 configuration as Dhcp4/subnet4[0]/pool.

6.1.4 Configuration Files Inclusion

The parser provides the ability to include files. The syntax was chosen to look similar to how Apache includes PHP scripts in HTML code. This particular syntax was chosen to emphasize that the include directive is an additional feature and not a part of JSON syntax.

The inclusion is implemented as a stack of files. You can use the include directive in nested includes. Up to ten nesting levels are supported. This arbitrarily chosen limit is protection against recursive inclusions.

The include directive has the form:

```
<?include "[PATH]"?>
```

The *[PATH]* pattern should be replaced with an absolute path or a path relative to the current working directory at the time the Kea process was launched.

To include one file from another, use the following syntax:

```
{
  "Dhcp6": {
    "interfaces-config": {
      "interfaces": [ "*" ],
      "preferred-lifetime": 3000,
      "rebind-timer": 2000,
      "renew-timer": 1000,
      <?include "subnets.json"?>
      "valid-lifetime": 4000
    }
  }
}
```

where the content of "subnets.json" may be:

```
{
  "subnet4": [
    {
      "id": 123,
      "subnet": "192.0.2.0/24"
    },
    {
      "id": 234,
      "subnet": "192.0.3.0/24"
    },
    {
      "id": 345,
      "subnet": "10.0.0.0/8"
    }
  ],
}
```

(continues on next page)

(continued from previous page)

```
...  
}
```

6.2 Kea Configuration Backend

6.2.1 Applicability

The Kea Configuration Backend (CB or config backend) gives Kea servers the ability to store almost all of their configuration in one or more databases.

Potential features and benefits include:

- Re-use of identical configuration sections across multiple Kea servers
- A "single source of truth" for all Kea servers
- All configuration done through an API with real-time logic checks
- Easier integration with third-party tools or in-house automation
- Database architecture provides concurrency, consistency, and atomicity
- The ability to mine the database for statistics and reporting
- Use of database replication for real-time fault-tolerance

Potential drawbacks include:

- Significant up-front effort to prepare integration and/or automation
- Supported scenarios require use of API for most configuration
- The API is only available to ISC customers with a paid support contract
- Incompatible with some uses, software, and scenarios

Note

Use of a database for storage of leases and/or host reservations is possible without using the CB. See the `host-databases` and `lease-database` config directives.

6.2.1.1 Example Scenario

Consider a large organization with many sites, each with a pair of Kea servers deployed in a high availability (HA) configuration. Typically such deployments use standardized configurations, leading to similar config parameters for every site. The members of each HA pair are often almost exactly the same, differing only by name. Often installations of this size will feature a high level of automation, including provisioning and management systems.

The CB is ideal for such deployments. While some effort is required to integrate with the provisioning automation, once accomplished, deployment of new Kea servers can be nearly automatic. A small "stub" JSON config, identical for every server, can be used, and all remaining configuration (shared networks, subnets, pools, etc.) loaded from the central CB database. Any updates to the central CB database are automatically propagated to all Kea instances. The CB becomes the "single source of truth" for DHCP throughout the organization.

6.2.2 Limitations and Warnings

6.2.2.1 Availability

libdhcp_cb_cmds.so is available only to ISC customers with a paid support contract. For more information on subscription options, please complete the form at <https://www.isc.org/contact>. While it is theoretically possible to use the CB without this hook, this is neither supported nor recommended.

6.2.2.2 Preparation is Required

The Configuration Backend is not a "plug-and-play" solution. Supported scenarios require use of the CB API commands. Configuration information must be loaded into the CB database using the API for the CB to have any effect. The general intent is for the CB to be integrated as part of a larger provisioning solution. Please do not define config-databases unless you have done the necessary preparation work.

6.2.2.3 Database Management

The only supported method for managing the contents of the CB database is through the *libdhcp_cb_cmds.so* hook, which provides API commands for config backends. As a practical matter, to use the CB, you must do almost all Kea configuration through the CB API.

While it is theoretically possible to use the CB without the API (using tools such as MySQL Workbench or the command-line MySQL client), these avenues are neither recommended nor supported.

6.2.2.4 Config Conflicts

We strongly recommend against mixing configuration information from the CB and JSON config file. In other words, do not use both JSON config declarations and CB to configure the same types of objects. Ideally, when using the CB, the Kea config files should contain the absolute bare minimum necessary, with everything else coming from the CB.

Using both CB and JSON as a source of configuration risks conflicting definitions. Such conflicts are confusing at best, and usually lead to undesired behavior or errors from Kea.

In the event of a conflict, configuration instructions from the CB database generally take precedence over instructions from a JSON file.

In certain carefully-controlled scenarios, it may be technically possible to use both. For example, defining one subnet in a JSON file and a second (different) subnet in the CB database would not conflict. However, other structures are replaced entirely. For example, if client classes are defined in the CB database, the DHCP server disregards any client classes defined in the JSON file.

In short, if you are using the CB, you should use only the CB.

6.2.2.5 Incompatible Software

API commands which modify Kea's configuration (other than the CB API) are contraindicated. This includes the *libdhcp_subnet_cmds.so* hook. Its APIs modify Kea's in-memory configuration, and the modifications can only be made persistent by using `config-write` to write a new JSON config file.

The *Stork* management suite does not currently support the CB. Stork makes all configuration changes through API avenues which expect to write changes into JSON config file. Support for the CB is planned for a future release of Stork.

In certain carefully-controlled scenarios, it may be possible to use these tools with the CB. Namely, if they are used in strictly "read-only" fashion, to retrieve Kea information, but never to modify it. However, no protection against accidental modification is provided, so this is not recommended.

6.2.2.6 Limitations

Currently, the Kea CB has the following limitations:

- It is only supported for MySQL and PostgreSQL databases.
- It is only supported for the DHCPv4 and DHCPv6 daemons; the Control Agent, D2 daemon, and the NETCONF daemon cannot be configured from the database,
- Only certain DHCP configuration parameters can be set in the database: global parameters, option definitions, global options, client classes, shared networks, and subnets. Other configuration parameters must be sourced from a JSON configuration file.

6.2.2.7 Custom Options

Using custom option formats requires creating definitions for these options. Suppose a user wishes to set option data in the configuration backend. In that case, we recommend specifying the definition for that option in the configuration backend as well. It is essential when multiple servers are managed via the configuration backend, and may differ in their configurations. The option data parser can search for an option definition appropriate for the server for which the option data is specified.

In a single-server deployment, or when all servers share the same configuration file information, it is possible to specify option definitions in the configuration files and option data in the configuration backend. The server receiving a command to set option data must have a valid definition in its configuration file, even when it sets option data for another server.

It is not supported to specify option definitions in the configuration backend and the corresponding option data in the server configuration files.

6.2.3 Components

The Kea Configuration Backend solution consists of the CB modules (hook libraries), the CB commands API (its own hook library), the external database software (MySQL or PostgreSQL), the database schema, and the Kea configuration information stored in the database.

In this documentation, the term "Configuration Backend" may also refer to the particular Kea module providing support for that database type. For example, the MySQL Configuration Backend, *libdhcp_mysql.so*, provides a complete set of functions to manage and fetch the configuration information from a MySQL database. The PostgreSQL Configuration Backend, *libdhcp_pgsql.so*, is the corresponding module for PostgreSQL. Similarly, the term "database" is used to refer to either a MySQL or PostgreSQL database.

The CB commands API provides a complete set of commands to manage Kea configuration information, as stored within the database. This API is implemented in its own hook library, *libdhcp_cb_cmds.so*. This library can be attached to both DHCPv4 and DHCPv6 server instances. It simplifies many typical operations, such as listing, adding, retrieving, and deleting global parameters, shared networks, subnets, pools, options, option definitions, and client classes. In addition, it provides essential business logic that ensures the logical integrity of the data. All CB API commands start with `remote-`. See *libdhcp_cb_cmds.so: Configuration Backend Commands* for more information.

Installation and maintenance of external database software is beyond the scope of this manual.

The database schema is typically installed via the *kea-admin* tool. See *Installation* for more information. The raw schema creation scripts are *dhcpdb_create.mysql* and *dhcpdb_create.pgsql*.

Use the CB commands API to populate the database with Kea configuration information.

Related design documents are available in our GitLab:

- [CB Design](#)
- [Client Classes in CB Design](#)

6.2.4 Installation

To use either Configuration Backend, the appropriate module library (*libdhcp_mysql.so* or *libdhcp_pgsql.so*) must be compiled during the Kea build. The `-D` switch specifies which database module to build, if any: `-D mysql=enabled` or `-D postgresql=enabled`. The appropriate database client libraries and header files must be installed prior to build. See *DHCP Database Installation and Configuration* for more information on building Kea with database support. ISC's Kea packaging, as well as some third-party distributions, provide separate packages for each database type.

The database server hosting the CB tables must be prepared with the Kea schema. When upgrading an existing Kea installation, the database schema may also need to be upgraded. The *kea-admin* tool can be used to more easily apply the schema, as described in *The kea-admin Tool*.

At runtime, the DHCP servers must be configured to load the module, in the `hooks-libraries` section. A `config-databases` directive must then be used to instruct Kea to load configuration using the database backend. The DHCPv4 and DHCPv6 server-specific configurations of the CB, as well as the list of supported configuration parameters, can be found in *Configuration Backend in DHCPv4* and *Configuration Backend in DHCPv6*, respectively.

Once installation is completed, the CB commands API can be used to populate the database with Kea configuration information.

6.2.5 Configuration Sharing and Server Tags

The configuration database is designed to store configuration information for multiple Kea servers. Depending on the use case, the entire configuration may be shared by all servers; parts of the configuration may be shared by multiple servers and the rest of the configuration may be different for these servers; or each server may have its own non-shared configuration.

The configuration elements in the database are associated with the servers by "server tags." The server tag is an arbitrary string holding the name of the Kea server instance. The tags of the DHCPv4 and DHCPv6 servers are independent in the database, i.e. the same server tag can be created for both the DHCPv4 and the DHCPv6 server. The value is configured using the `server-tag` parameter in the `Dhcp4` or `Dhcp6` scope. The current server tag can be checked with the *server-tag-get* command.

The server definition, which consists of the server tag and the server description, must be stored in the configuration database prior to creating the dedicated configuration for that server. In cases when all servers use the same configuration, e.g. a pair of servers running as High Availability peers, there is no need to configure the server tags for these servers in the database.

Commands which contain the logical server *all* are applied to all servers connecting to the database. The *all* server cannot be deleted or modified, and it is not returned among other servers as a result of the *remote-server4-get-all* and *remote-server6-get-all* commands.

In most cases, there are no server tags defined in the configuration database; all connecting servers get the same configuration regardless of the server tag they use. The server tag that a particular Kea instance presents to the database to fetch its configuration is specified in the Kea configuration file, using the *config-control* map (please refer to the *Enabling the Configuration Backend* and *Enabling the Configuration Backend* for details). All Kea instances presenting the same server tag to the configuration database are given the same configuration.

It is the administrator's choice whether multiple Kea instances use the same server tag or each Kea instance uses a different server tag. There is no requirement that the instances running on the same physical or virtual machine use the same server tag. It is even possible to configure the Kea server without assigning it a server tag. In such a case the server will be given the configuration specified for *all* servers.

To differentiate between different Kea server configurations, a list of the server tags used by the servers must be stored in the database. For the DHCPv4 and DHCPv6 servers, this can be done using the *remote-server4-set* and *remote-server6-set* commands. The server tags can then be used to associate the configuration information with the servers. However, it is important to note that some DHCP configuration elements may be associated with multiple server tags (known as "shareable" elements), while other configuration elements may be associated with only

one server tag ("non-shareable" elements). The *Configuration Backend in DHCPv4* and *Configuration Backend in DHCPv6* sections list the DHCP-specific shareable and non-shareable configuration elements; however, in this section we briefly explain the differences between them.

A shareable configuration element is one which has some unique property identifying it, and which may appear only once in the database. An example of a shareable DHCP element is a subnet instance: the subnet is a part of the network topology and we assume that any particular subnet may have only one definition within this network. Each subnet has two unique identifiers: the subnet identifier and the subnet prefix. The subnet identifier is used in Kea to uniquely identify the subnet within the network and to connect it with other configuration elements, e.g. in host reservations. Some commands provided by *libdhcp_cb_cmds.so* allow the subnet information to be accessed by either subnet identifier or prefix, and explicitly prohibit using the server tag to access the subnet. This is because, in general, the subnet definition is associated with multiple servers rather than a single server. In fact, it may even be associated with no servers (unassigned). Still, the unassigned subnet has an identifier and prefix which can be used to access the subnet.

A shareable configuration element may be associated with multiple servers, one server, or no servers. Deletion of the server which is associated with the shareable element does not cause the deletion of the shareable element. It merely deletes the association of the deleted server with the element.

Unlike a shareable element, a non-shareable element must not be explicitly associated with more than one server and must not exist after the server is deleted (must not remain unassigned). A non-shareable element only exists within the context of the server. An example of a non-shareable element in DHCP is a global parameter, e.g. *renew-timer*. The renew timer is the value to be used by a particular server and only this server. Other servers may have their respective renew timers set to the same or different values. The renew timer parameter has no unique identifier by which it could be accessed, modified, or otherwise used. Global parameters like the renew timer can be accessed by the parameter name and the tag of the server for which they are configured. For example, the *remote-global-parameter4-get* and *remote-global-parameter6-get* commands allow the value of the global parameter to be fetched by the parameter name and the server name. Getting the global parameter only by its name (without specifying the server tag) is not possible, because there may be many global parameters with a given name in the database.

When the server associated with a non-shareable configuration element is deleted, the configuration element is automatically deleted from the database along with the server because the non-shareable element must be always assigned to a server (or the logical server *all*).

The terms "shareable" and "non-shareable" only apply to associations with user-defined servers; all configuration elements associated with the logical server *all* are by definition shareable. For example: the *renew-timer* associated with *all* servers is used by all servers connecting to the database which do not have their specific renew timers defined. In a special case, when none of the configuration elements are associated with user-defined servers, the entire configuration in the database is shareable because all its pieces belong to *all* servers.

Note

Be very careful when associating configuration elements with different server tags. The configuration backend does not protect against some possible misconfigurations that may arise from the wrong server tags' assignments. For example: if a shared network is assigned to one server and the subnets belonging to this shared network to another server, the servers will fail upon trying to fetch and use this configuration. The server fetching the subnets will be aware that the subnets are associated with the shared network, but the shared network will not be found by this server since it doesn't belong to it. In such a case, both the shared network and the subnets should be assigned to the same set of servers.

MANAGING KEA WITH KEACTRL

7.1 Overview

keactrl is a shell script which controls the startup, shutdown, and reconfiguration of the Kea servers (*kea-dhcp4*, *kea-dhcp6*, *kea-dhcp-ddns*, *kea-ctrl-agent*, and *kea-netconf*). It also provides the means for checking the current status of the servers and determining the configuration files in use.

keactrl is available only when Kea is built from sources. When installing Kea using native packages, the native systemd scripts are provided. See *Native Packages and systemd* Section for details.

7.2 Command Line Options

keactrl is run as follows:

```
# keactrl <command> [-c keactrl-config-file] [-s server[,server,...]]
```

<command> is one of the commands described in *Commands*.

The optional `-c keactrl-config-file` switch allows specification of an alternate *keactrl* configuration file. (`--ctrl-config` is a synonym for `-c`.) In the absence of `-c`, *keactrl* uses the default configuration file `[kea-install-dir]/etc/kea/keactrl.conf`.

The optional `-s server[,server,...]` switch selects the servers to which the command is issued. (`--server` is a synonym for `-s`.) If absent, the command is sent to all servers enabled in the *keactrl* configuration file. If multiple servers are specified, they should be separated by commas with no intervening spaces.

7.3 The keactrl Configuration File

Depending on the administrator's requirements, it may not be necessary to run all of the available servers. The *keactrl* configuration file sets which servers are enabled and which are disabled. The default configuration file is `[kea-install-dir]/etc/kea/keactrl.conf`, but this can be overridden on a per-command basis using the `-c` switch.

The contents of `keactrl.conf` are:

```
# This is a configuration file for keactrl script which controls
# the startup, shutdown, reconfiguration and gathering the status
# of the Kea processes.

# prefix holds the location where the Kea is installed.
prefix=@prefix@
```

(continues on next page)

(continued from previous page)

```

# Location of Kea configuration file.
kea_dhcp4_config_file=@sysconfdir@/@PACKAGE@/kea-dhcp4.conf
kea_dhcp6_config_file=@sysconfdir@/@PACKAGE@/kea-dhcp6.conf
kea_dhcp_ddns_config_file=@sysconfdir@/@PACKAGE@/kea-dhcp-ddns.conf
kea_ctrl_agent_config_file=@sysconfdir@/@PACKAGE@/kea-ctrl-agent.conf
kea_netconf_config_file=@sysconfdir@/@PACKAGE@/kea-netconf.conf

# Location of Kea binaries.
exec_prefix=@exec_prefix@
dhcp4_srv=@sbindir@/kea-dhcp4
dhcp6_srv=@sbindir@/kea-dhcp6
dhcp_ddns_srv=@sbindir@/kea-dhcp-ddns
ctrl_agent_srv=@sbindir@/kea-ctrl-agent
netconf_srv=@sbindir@/kea-netconf

# Start DHCPv4 server?
dhcp4=yes

# Start DHCPv6 server?
dhcp6=yes

# Start DHCP DDNS server?
dhcp_ddns=no

# Start Control Agent?
ctrl_agent=yes

# Start Netconf?
netconf=no

# Be verbose?
kea_verbose=no

```

Note

In the example above, strings of the form @something@ are replaced by the appropriate values when Kea is installed.

Setting the `dhcp4`, `dhcp6`, `dhcp_ddns`, `ctrl_agent`, and `netconf` parameters set to "yes" configures *keactrl* to manage (start, reconfigure) all servers, i.e. *kea-dhcp4*, *kea-dhcp6*, *kea-dhcp-ddns*, *kea-ctrl-agent*, and *kea-netconf*. When any of these parameters is set to "no", *keactrl* ignores the corresponding server when starting or reconfiguring Kea. Some daemons (`dhcp_ddns` and `netconf`) are disabled by default.

By default, Kea servers managed by *keactrl* are located in `[kea-install-dir]/sbin`. This should work for most installations. If the default location needs to be altered, the paths specified with the `dhcp4_srv`, `dhcp6_srv`, `dhcp_ddns_srv`, `ctrl_agent_srv`, and `netconf_srv` parameters should be modified.

The `kea_verbose` parameter specifies the verbosity of the servers being started. When `kea_verbose` is set to `yes`, the logging level of the server is set to `DEBUG`. Modification of the logging severity in a configuration file, as described in *Logging*, will have no effect as long as `kea_verbose` is set to "yes." Setting it to "no" causes the server to use the logging levels specified in the Kea configuration file. If no logging configuration is specified, the default settings are used.

Note

The verbosity for the server is set when it is started. Once started, the verbosity can only be changed by stopping the server and starting it again with the new value of the `kea_verbosity` parameter.

7.4 Commands

The following commands are supported by `keactrl`:

- `start` - starts the selected servers.
- `stop` - stops all running servers.
- `reload` - triggers reconfiguration of the selected servers by sending the SIGHUP signal to them.
- `status` - returns the status of the servers (active or inactive) and the names of the configuration files in use.
- `version` - prints out the version of the `keactrl` tool itself, together with the versions of the Kea daemons.

Typical output from `keactrl` when starting the servers looks similar to the following:

```
$ keactrl start
INFO/keactrl: Starting kea-dhcp4 -c /usr/local/etc/kea/kea-dhcp4.conf -d
INFO/keactrl: Starting kea-dhcp6 -c /usr/local/etc/kea/kea-dhcp6.conf -d
INFO/keactrl: Starting kea-dhcp-ddns -c /usr/local/etc/kea/kea-dhcp-ddns.conf -d
INFO/keactrl: Starting kea-ctrl-agent -c /usr/local/etc/kea/kea-ctrl-agent.conf -d
INFO/keactrl: Starting kea-netconf -c /usr/local/etc/kea/kea-netconf.conf -d
```

Kea's servers create PID files upon startup. These files are used by `keactrl` to determine whether a given server is running. If one or more servers are running when the start command is issued, the output looks similar to the following:

```
$ keactrl start
INFO/keactrl: kea-dhcp4 appears to be running, see: PID 10918, PID file: /usr/local/var/
↳run/kea/kea.kea-dhcp4.pid.
INFO/keactrl: kea-dhcp6 appears to be running, see: PID 10924, PID file: /usr/local/var/
↳run/kea/kea.kea-dhcp6.pid.
INFO/keactrl: kea-dhcp-ddns appears to be running, see: PID 10930, PID file: /usr/local/
↳var/run/kea/kea.kea-dhcp-ddns.pid.
INFO/keactrl: kea-ctrl-agent appears to be running, see: PID 10931, PID file: /usr/local/
↳var/run/kea/kea.kea-ctrl-agent.pid.
INFO/keactrl: kea-netconf appears to be running, see: PID 10123, PID file: /usr/local/
↳var/run/kea/kea.kea-netconf.pid.
```

During normal shutdowns, these PID files are deleted; they may, however, be left over as remnants following a system crash. It is possible, though highly unlikely, that upon system restart the PIDs they contain may actually refer to processes unrelated to Kea. This condition will cause `keactrl` to decide that the servers are running, when in fact they are not. In such a case the PID files listed in the `keactrl` output must be manually deleted.

The following command stops all servers:

```
$ keactrl stop
INFO/keactrl: Stopping kea-dhcp4...
INFO/keactrl: Stopping kea-dhcp6...
INFO/keactrl: Stopping kea-dhcp-ddns...
INFO/keactrl: Stopping kea-ctrl-agent...
INFO/keactrl: Stopping kea-netconf...
```

Note that the `stop` command attempts to stop all servers regardless of whether they are "enabled" in `keactrl.conf`. If any of the servers are not running, an informational message is displayed as in the `stop` command output below.

```
$ keactrl stop
INFO/keactrl: kea-dhcp4 isn't running.
INFO/keactrl: kea-dhcp6 isn't running.
INFO/keactrl: kea-dhcp-ddns isn't running.
INFO/keactrl: kea-ctrl-agent isn't running.
INFO/keactrl: kea-netconf isn't running.
```

As already mentioned, the reconfiguration of each Kea server is triggered by the SIGHUP signal. The `reload` command sends the SIGHUP signal to any servers that are enabled in the `keactrl` configuration file and that are currently running. When a server receives the SIGHUP signal it rereads its configuration file and, if the new configuration is valid, uses the new configuration. If the new configuration proves to be invalid, the server retains its current configuration; however, in some cases a fatal error message is logged indicating that the server is no longer providing any service: a working configuration must be loaded as soon as possible.

A reload is executed as follows:

```
$ keactrl reload
INFO/keactrl: Reloading kea-dhcp4...
INFO/keactrl: Reloading kea-dhcp6...
INFO/keactrl: Reloading kea-dhcp-ddns...
INFO/keactrl: Reloading kea-ctrl-agent...
```

If any of the servers are not running, an informational message is displayed as in the `reload` command output below. `kea-netconf` does not support the SIGHUP signal. If its configuration has changed, please stop and restart it for the change to take effect.

```
$ keactrl stop
INFO/keactrl: kea-dhcp4 isn't running.
INFO/keactrl: kea-dhcp6 isn't running.
INFO/keactrl: kea-dhcp-ddns isn't running.
INFO/keactrl: kea-ctrl-agent isn't running.
INFO/keactrl: kea-netconf isn't running.
```

Note

NETCONF is an optional feature that is disabled by default and can be enabled during compilation. If Kea was compiled without NETCONF support, `keactrl` does not provide information about it. The NETCONF entries are still present in the `keactrl.conf` file, but NETCONF status is not shown and other commands ignore it.

Note

Currently `keactrl` does not report configuration failures when the server is started or reconfigured. To check if the server's configuration succeeded, the Kea log must be examined for errors. By default, the log is written to the `syslog` file.

Sometimes it is useful to check which servers are running. The `status` command reports this, with typical output that looks like:


```
$ keactrl status
DHCPv4 server: active
DHCPv6 server: inactive
DHCP DDNS: active
Control Agent: active
Netconf agent: inactive
Kea configuration file: /usr/local/etc/kea/kea.conf
Kea DHCPv4 configuration file: /usr/local/etc/kea/kea-dhcp4.conf
Kea DHCPv6 configuration file: /usr/local/etc/kea/kea-dhcp6.conf
Kea DHCP DDNS configuration file: /usr/local/etc/kea/kea-dhcp-ddns.conf
Kea Control Agent configuration file: /usr/local/etc/kea/kea-ctrl-agent.conf
Kea Netconf configuration file: /usr/local/etc/kea/kea-netconf.conf
keactrl configuration file: /usr/local/etc/kea/keactrl.conf
```

`keactrl status` offers basic reporting capabilities. For more extensive insight into Kea's health and status, consider deploying Stork. For details, see [Monitoring Kea With Stork](#).

7.5 Overriding the Server Selection

The optional `-s` switch allows the selection of the server(s) to which the `keactrl` command is issued. For example, the following instructs `keactrl` to stop the `kea-dhcp4` and `kea-dhcp6` servers and leave the `kea-dhcp-ddns` and `kea-ctrl-agent` running:

```
$ keactrl stop -s dhcp4,dhcp6
```

Similarly, the following starts only the `kea-dhcp4` and `kea-dhcp-ddns` servers, but not `kea-dhcp6` or `kea-ctrl-agent`.

```
$ keactrl start -s dhcp4,dhcp_ddns
```

Note that the behavior of the `-s` switch with the `start` and `reload` commands is different from its behavior with the `stop` command. On `start` and `reload`, `keactrl` checks whether the servers given as parameters to the `-s` switch are enabled in the `keactrl` configuration file; if not, the server is ignored. For `stop`, however, this check is not made; the command is applied to all listed servers, regardless of whether they have been enabled in the file.

The following keywords can be used with the `-s` command-line option:

- `dhcp4` for `kea-dhcp4`.
- `dhcp6` for `kea-dhcp6`.
- `dhcp_ddns` for `kea-dhcp-ddns`.
- `ctrl_agent` for `kea-ctrl-agent`.
- `netconf` for `kea-netconf`.
- `all` for all servers (default).

7.6 Native Packages and systemd

`keactrl` is a script that was developed to assist in managing Kea processes. However, all modern operating systems have their own process-management scripts, such as `systemd`. In general, these native scripts should be used, as they have several advantages. `systemd` scripts handle processes in a uniform way, so Kea is handled in a similar fashion to HTTP or a mail server. Second and more importantly, `systemd` allows dependencies to be defined between services. For example, it is easy to specify that the Kea server should not start until the network interfaces are operational. Using

native scripts also has other benefits, such as the ability to enable or disable services using commands, and the ability to temporarily start a disabled service.

Thus, it is recommended to use `systemctl` commands if they are available. Native Kea packages do not provide `keactrl`; `systemctl` service definitions are provided instead. Consult the system documentation for details.

Briefly, here are example commands to check status, start, stop, and restart various Kea daemons:

```
# systemctl status kea-ctrl-agent
# systemctl start kea-dhcp4
# systemctl stop kea-dhcp6
# systemctl restart kea-dhcp-ddns
```

Note that the service names may be slightly different between Linux distributions; in general, we have followed the naming conventions in third-party packages. In particular, some systems may not have the *isc*- prefix.

THE KEA CONTROL AGENT

8.1 Overview of the Kea Control Agent

The Kea Control Agent (CA) is a daemon which exposes a RESTful control interface for managing Kea servers. The daemon can receive control commands over HTTP and either forward these commands to the respective Kea servers or handle these commands on its own. The determination whether the command should be handled by the CA or forwarded is made by checking the value of the `service` parameter, which may be included in the command from the controlling client. The details of the supported commands, as well as their structures, are provided in *Management API*.

The CA can use hook libraries to provide support for additional commands or to program custom behavior of existing commands. Such hook libraries must implement callouts for the `control_command_receive` hook point. Details about creating new hook libraries and supported hook points can be found in the *Kea Developer's Guide*.

The CA processes received commands according to the following algorithm:

- Pass command into any installed hooks (regardless of service value(s)). If the command is handled by a hook, return the response.
- If the service specifies one or more services, forward the command to the specified services and return the accumulated responses.
- If the service is not specified or is an empty list, handle the command if the CA supports it.

Note

The `kea-ctrl-agent` is deprecated. As of Kea 2.7.2 the Kea servers: `kea-dhcp4`, `kea-dhcp6`, and `kea-dhcp-ddns` directly support command API connections over HTTPS/HTTP. (*Migration from the Control Agent*).

8.2 Configuration

The following example demonstrates the basic CA configuration.

```
{
  "Control-agent": {
    "http-host": "10.20.30.40",
    "http-port": 8000,
    "http-headers": [
      {
        "name": "Strict-Transport-Security",
        "value": "max-age=31536000"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "trust-anchor": "/path/to/the/ca-cert.pem",
    "cert-file": "/path/to/the/agent-cert.pem",
    "key-file": "/path/to/the/agent-key.pem",
    "cert-required": true,
    "authentication": {
        "type": "basic",
        "realm": "kea-control-agent",
        "clients": [
            {
                "user": "admin",
                "password": "1234"
            }
        ]
    },
    "control-sockets": {
        "dhcp4": {
            "comment": "main server",
            "socket-type": "unix",
            "socket-name": "/path/to/the/unix/socket-v4"
        },
        "dhcp6": {
            "socket-type": "unix",
            "socket-name": "/path/to/the/unix/socket-v6",
            "user-context": { "version": 3 }
        },
        "d2": {
            "socket-type": "unix",
            "socket-name": "/path/to/the/unix/socket-d2"
        }
    },
    "hooks-libraries": [
        {
            "library": "/opt/local/custom_hooks_example.so",
            "parameters": {
                "param1": "foo"
            }
        }
    ],
    "loggers": [ {
        "name": "kea-ctrl-agent",
        "severity": "INFO"
    } ]
}

```

The `http-host` and `http-port` parameters specify an IP address and port to which HTTP service will be bound. In the example configuration provided above, the RESTful service will be available at the URL `https://10.20.30.40:8000/`. If these parameters are not specified, the default URL is `http://127.0.0.1:8000/`.

When using Kea's HA hook library with multi-threading, the address:port combination used for CA must be different from the HA peer URLs, which are strictly for internal HA traffic between the peers. User commands should still be

sent via the CA.

Since Kea 2.7.5 the `http-headers` parameter specifies a list of extra HTTP headers to add to HTTP responses.

The `trust-anchor`, `cert-file`, `key-file`, and `cert-required` parameters specify the TLS setup for HTTP, i.e. HTTPS. If these parameters are not specified, HTTP is used. The TLS/HTTPS support in Kea is described in [TLS/HTTPS Support](#).

As mentioned in [Overview of the Kea Control Agent](#), the CA can forward received commands to the Kea servers for processing. For example, `config-get` is sent to retrieve the configuration of one of the Kea services. When the CA receives this command, including a `service` parameter indicating that the client wishes to retrieve the configuration of the DHCPv4 server, the CA forwards the command to that server and passes the received response back to the client. More about the `service` parameter and the general structure of commands can be found in [Management API](#).

The CA uses UNIX domain sockets to forward control commands and receive responses from other Kea services. The `dhcp4`, `dhcp6`, and `d2` maps specify the files to which UNIX domain sockets are bound. In the configuration above, the CA connects to the DHCPv4 server via `/path/to/the/unix/socket-v4` to forward the commands to it. Obviously, the DHCPv4 server must be configured to listen to connections via this same socket. In other words, the command-socket configuration for the DHCPv4 server and the CA (for that server) must match. Consult [UNIX Control Socket](#), [UNIX Control Socket](#), and [UNIX Control Socket](#) to learn how the UNIX socket configuration is specified for the DHCPv4, DHCPv6, and D2 services.

Note

As of Kea 2.7.9, control sockets may only reside in the directory determined during compilation as "[`kea-install-dir`]/var/run/kea", which must also have `0750` access rights. This path may be overridden at startup by setting the environment variable `KEA_CONTROL_SOCKET_DIR` to the desired path. If a path other than this value is used in `socket-name`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use simply omit the path component from `socket-name`.

User contexts can store arbitrary data as long as they are in valid JSON syntax and their top-level element is a map (i.e. the data must be enclosed in curly brackets). Some hook libraries may expect specific formatting; please consult the relevant hook library documentation for details.

User contexts can be specified on either global scope, control socket, basic authentication, or loggers. One other useful feature is the ability to store comments or descriptions; the parser translates a "comment" entry into a user context with the entry, which allows a comment to be attached within the configuration itself.

Basic HTTP authentication protects against unauthorized uses of the control agent by local users. For protection against remote attackers, HTTPS and reverse proxy of [Secure Connections](#) provide stronger security.

The authentication is described in the `authentication` block with the mandatory `type` parameter, which selects the authentication. Currently only the basic HTTP authentication (type `basic`) is supported.

The `realm` authentication parameter is used for error messages when the basic HTTP authentication is required but the client is not authorized.

When the `clients` authentication list is configured and not empty, basic HTTP authentication is required. Each element of the list specifies a user ID and a password. The user ID is mandatory, must not be empty, and must not contain the colon (:) character. The password is optional; when it is not specified an empty password is used.

Note

The basic HTTP authentication user ID and password are encoded in UTF-8, but the current Kea JSON syntax only supports the Latin-1 (i.e. 0x00..0xff) Unicode subset.

To avoid exposing the user ID and/or the associated password, these values can be read from files. The syntax is extended by:

- The `directory` authentication parameter, which handles the common part of file paths. The default value is the empty string.
- The `password-file` client parameter, which, alongside the `directory` parameter, specifies the path of a file that can contain the password, or when no user ID is given, the whole basic HTTP authentication secret.
- The `user-file` client parameter, which, with the `directory` parameter, specifies the path of a file where the user ID can be read.

When files are used, they are read when the configuration is loaded, to detect configuration errors as soon as possible.

Hook libraries can be loaded by `kea-ctrl-agent` in the same way as they are loaded by `kea-dhcp4` and `kea-dhcp6`. The CA currently supports one hook point - `control_command_receive` - which makes it possible to delegate the processing of some commands to the hook library. The `hooks-libraries` list contains the list of hook libraries that should be loaded by `kea-ctrl-agent`, along with their configuration information specified with `parameters`.

Please consult [Logging](#) for the details on how to configure logging. The CA's root logger's name is `kea-ctrl-agent`, as given in the example above.

8.3 Secure Connections

Configuration options related to Kea Control Agent security can be found in the [Secure Kea Control Agent](#) section.

8.4 Starting and Stopping the Control Agent

`kea-ctrl-agent` accepts the following command-line switches:

- `-c file` - specifies the configuration file.
- `-d` - specifies whether the agent logging should be switched to debug/verbose mode. In verbose mode, the logging severity and debuglevel specified in the configuration file are ignored and "debug" severity and the maximum debuglevel (99) are assumed. The flag is convenient for temporarily switching the server into maximum verbosity, e.g. when debugging.
- `-t file` - specifies the configuration file to be tested. `kea-netconf` attempts to load it and conducts sanity checks; certain checks are possible only while running the actual server. The actual status is reported with exit code (0 = configuration appears valid, 1 = error encountered). Kea prints out log messages to standard output and error to standard error when testing the configuration.
- `-v` - displays the version of `kea-ctrl-agent` and exits.
- `-V` - displays the extended version information for `kea-ctrl-agent` and exits. The listing includes the versions of the libraries dynamically linked to Kea.
- `-W` - displays the Kea configuration report and exits. The report is a copy of the `config.report` file produced by `meson setup`; it is embedded in the executable binary.

The contents of the `config.report` file may also be accessed by examining certain libraries in the installation tree or in the source tree.

```
# from installation using libkea-process.so
$ strings ${prefix}/lib/libkea-process.so | sed -n 's/;;;; //p'

# from sources using libkea-process.so
$ strings src/lib/process/.libs/libkea-process.so | sed -n 's/;;;; //p'
```

(continues on next page)

(continued from previous page)

```
# from sources using libkea-process.a
$ strings src/lib/process/.libs/libkea-process.a | sed -n 's/;;; //p'

# from sources using libcfgrpt.a
$ strings src/lib/process/cfgrpt/.libs/libcfgrpt.a | sed -n 's/;;; //p'
```

- **-X** - As of Kea 3.0, disables path and permissions restrictions. The server will emit a warning at startup that security restrictions have been disabled. Do not use this mode of operation without careful consideration and taking any necessary precautions. Failure to do so can expose deployments to security vulnerabilities. For more information please read section *Securing a Kea Deployment*.

The CA is started by running its binary and specifying the configuration file it should use. For example:

```
$ ./kea-ctrl-agent -c /usr/local/etc/kea/kea-ctrl-agent.conf
```

It can be started by *keactrl* as well (see *Managing Kea with keactrl*).

8.5 Connecting to the Control Agent

For an example of a tool that can take advantage of the RESTful API, see *The Kea Shell*.

THE DHCPV4 SERVER

9.1 Starting and Stopping the DHCPv4 Server

It is recommended that the Kea DHCPv4 server be started and stopped using *keactrl* (described in *Managing Kea with keactrl*); however, it is also possible to run the server directly via the *kea-dhcp4* command, which accepts the following command-line switches:

- **-c file** - specifies the configuration file. This is the only mandatory switch.
- **-d** - specifies whether the server logging should be switched to debug/verbose mode. In verbose mode, the logging severity and debuglevel specified in the configuration file are ignored; "debug" severity and the maximum debuglevel (99) are assumed. The flag is convenient for temporarily switching the server into maximum verbosity, e.g. when debugging.
- **-p server-port** - specifies the local UDP port on which the server listens. This is only useful during testing, as a DHCPv4 server listening on ports other than the standard ones is not able to handle regular DHCPv4 queries.
- **-P client-port** - specifies the remote UDP port to which the server sends all responses. This is only useful during testing, as a DHCPv4 server sending responses to ports other than the standard ones is not able to handle regular DHCPv4 queries.
- **-t file** - specifies a configuration file to be tested. *kea-dhcp4* loads it, checks it, and exits. During the test, log messages are printed to standard output and error messages to standard error. The result of the test is reported through the exit code (0 = configuration looks OK, 1 = error encountered). The check is not comprehensive; certain checks are possible only when running the server.
- **-T file** - specifies a configuration file to be tested. *kea-dhcp4* loads it, checks it, and exits. It performs extra checks beyond what **-t** offers, such as establishing database connections (for the lease backend, host reservations backend, configuration backend, and forensic logging backend), loading hook libraries, parsing hook-library configurations, etc. It does not open UNIX or TCP/UDP sockets, nor does it open or rotate files, as any of these actions could interfere with a running process on the same machine.
- **-v** - displays the Kea version and exits.
- **-V** - displays the Kea extended version with additional parameters and exits. The listing includes the versions of the libraries dynamically linked to Kea.
- **-W** - displays the Kea configuration report and exits. The report is a copy of the `config.report` file produced by `meson setup`; it is embedded in the executable binary.

The contents of the `config.report` file may also be accessed by examining certain libraries in the installation tree or in the source tree.

```
# from installation using libkea-process.so
$ strings ${prefix}/lib/libkea-process.so | sed -n 's/;;; //p'
```

(continues on next page)

(continued from previous page)

```
# from sources using libkea-process.so
$ strings src/lib/process/.libs/libkea-process.so | sed -n 's/;;;; //p'

# from sources using libkea-process.a
$ strings src/lib/process/.libs/libkea-process.a | sed -n 's/;;;; //p'

# from sources using libcfgrpt.a
$ strings src/lib/process/cfgrpt/.libs/libcfgrpt.a | sed -n 's/;;;; //p'
```

- **-X** - As of Kea 3.0, disables path and permissions restrictions. The server will emit a warning at startup that security restrictions have been disabled. Do not use this mode of operation without careful consideration and taking any necessary precautions. Failure to do so can expose deployments to security vulnerabilities. For more information please read section *Securing a Kea Deployment*.

On startup, the server detects available network interfaces and attempts to open UDP sockets on all interfaces listed in the configuration file. Since the DHCPv4 server opens privileged ports, it requires root access; this daemon must be run as root.

During startup, the server attempts to create a PID file of the form: `[pidfile_dir]/[conf name].kea-dhcp4.pid` where:

- **pidfile_dir** - is `[prefix]/[localstatedir]/run/kea` where `prefix` and `localstatedir` are the values passed into meson setup using `--prefix` and `--localstatedir` which default to `/usr/local` and `var` respectively. So the whole `pidfile_dir` defaults to `/usr/local/var`. Note that this value may be overridden at runtime by setting the environment variable `KEA_PIDFILE_DIR` intended primarily for testing purposes.
- **conf name**: The configuration file name used to start the server, minus all preceding paths and the file extension. For example, given a pathname of `/usr/local/etc/kea/myconf.txt`, the portion used would be `myconf`.

If the file already exists and contains the PID of a live process, the server issues a `DHCP4_ALREADY_RUNNING` log message and exits. It is possible, though unlikely, that the file is a remnant of a system crash and the process to which the PID belongs is unrelated to Kea. In such a case, it would be necessary to manually delete the PID file.

The server can be stopped using the `kill` command. When running in a console, the server can also be shut down by pressing `Ctrl-c`. Kea detects the key combination and shuts down gracefully.

The reconfiguration of each Kea server is triggered by the `SIGHUP` signal. When a server receives the `SIGHUP` signal it rereads its configuration file and, if the new configuration is valid, uses the new configuration. If the new configuration proves to be invalid, the server retains its current configuration; however, in some cases a fatal error message is logged indicating that the server is no longer providing any service: a working configuration must be loaded as soon as possible.

9.2 DHCPv4 Server Configuration

9.2.1 Introduction

This section explains how to configure the Kea DHCPv4 server using a configuration file.

Before DHCPv4 is started, its configuration file must be created. The basic configuration is as follows:

```
{
# DHCPv4 configuration starts on the next line
"Dhcp4": {

# First we set up global values
  "valid-lifetime": 4000,
  "renew-timer": 1000,
```

(continues on next page)

(continued from previous page)

```

    "rebind-timer": 2000,

# Next we set up the interfaces to be used by the server.
    "interfaces-config": {
        "interfaces": [ "eth0" ]
    },

# And we specify the type of lease database
    "lease-database": {
        "type": "memfile",
        "persist": true,
        "name": "/var/lib/kea/dhcp4.leases"
    },

# Finally, we list the subnets from which we will be leasing addresses.
    "subnet4": [
        {
            "id": 1,
            "subnet": "192.0.2.0/24",
            "pools": [
                {
                    "pool": "192.0.2.1 - 192.0.2.200"
                }
            ]
        }
    ]
# DHCPv4 configuration ends with the next line
}
}

```

The following paragraphs provide a brief overview of the parameters in the above example, along with their format. Subsequent sections of this chapter go into much greater detail for these and other parameters.

The lines starting with a hash (#) are comments and are ignored by the server; they do not impact its operation in any way.

The configuration starts in the first line with the initial opening curly bracket (or brace). Each configuration must contain an object specifying the configuration of the Kea module using it. In the example above, this object is called Dhc4.

The Dhc4 configuration starts with the "Dhc4": { line and ends with the corresponding closing brace (in the above example, the brace after the last comment). Everything defined between those lines is considered to be the Dhc4 configuration.

In general, the order in which those parameters appear does not matter, but there are two caveats. The first one is that the configuration file must be well-formed JSON, meaning that the parameters for any given scope must be separated by a comma, and there must not be a comma after the last parameter. When reordering a configuration file, moving a parameter to or from the last position in a given scope may also require moving the comma. The second caveat is that it is uncommon — although legal JSON — to repeat the same parameter multiple times. If that happens, the last occurrence of a given parameter in a given scope is used, while all previous instances are ignored. This is unlikely to cause any confusion as there are no real-life reasons to keep multiple copies of the same parameter in the configuration file.

The first few DHCPv4 configuration elements define some global parameters. `valid-lifetime` defines how long the

addresses (leases) given out by the server are valid; the default is for a client to be allowed to use a given address for 4000 seconds. (Note that integer numbers are specified as is, without any quotes around them.) `renew-timer` and `rebind-timer` are values (also in seconds) that define the T1 and T2 timers that govern when the client begins the renewal and rebind processes.

Note

The lease valid lifetime is expressed as a triplet with minimum, default, and maximum values using configuration entries `min-valid-lifetime`, `valid-lifetime`, and `max-valid-lifetime`. Since Kea 1.9.5, these values may be specified in client classes. The procedure the server uses to select which lifetime value to use is as follows:

If the client query is a BOOTP query, the server always uses the infinite lease time (e.g. 0xffffffff). Otherwise, the server must determine which configured triplet to use by first searching all classes assigned to the query, and then the subnet selected for the query.

Classes are searched in the order they were assigned to the query; the server uses the triplet from the first class that specifies it. If no classes specify the triplet, the server uses the triplet specified by the subnet selected for the client. If the subnet does not explicitly specify it, the server next looks at the subnet's shared-network (if one exists), then for a global specification, and finally the global default.

If the client requested a lifetime value via DHCP option 51, then the lifetime value used is the requested value bounded by the configured triplet. In other words, if the requested lifetime is less than the configured minimum, the configured minimum is used; if it is more than the configured maximum, the configured maximum is used. If the client did not provide a requested value, the lifetime value used is the triplet default value.

Note

Both `renew-timer` and `rebind-timer` are optional. The server only sends `rebind-timer` to the client, via DHCPv4 option code 59, if it is less than `valid-lifetime`; and it only sends `renew-timer`, via DHCPv4 option code 58, if it is less than `rebind-timer` (or `valid-lifetime` if `rebind-timer` was not specified). In their absence, the client should select values for T1 and T2 timers according to [RFC 2131](#). See section *Sending T1 (Option 58) and T2 (Option 59)* for more details on generating T1 and T2.

The `interfaces-config` map specifies the network interfaces on which the server should listen to DHCP messages. The `interfaces` parameter specifies a list of network interfaces on which the server should listen. Lists are opened and closed with square brackets, with elements separated by commas. To listen on two interfaces, the `interfaces-config` element should look like this:

```
{
"interfaces-config": {
  "interfaces": [ "eth0", "eth1" ]
},
...
}
```

The next lines define the lease database, the place where the server stores its lease information. This particular example tells the server to use memfile, which is the simplest and fastest database backend. It uses an in-memory database and stores leases on disk in a CSV (comma-separated values) file. This is a very simple configuration example; usually the lease database configuration is more extensive and contains additional parameters. Note that `lease-database` is an object and opens up a new scope, using an opening brace. Its parameters (just one in this example: `type`) follow. If there were more than one, they would be separated by commas. This scope is closed with a closing brace. As more parameters for the `Dhcp4` definition follow, a trailing comma is present.

Finally, we need to define a list of IPv4 subnets. This is the most important DHCPv4 configuration structure, as the

server uses that information to process clients' requests. It defines all subnets from which the server is expected to receive DHCP requests. The subnets are specified with the `subnet4` parameter. It is a list, so it starts and ends with square brackets. Each subnet definition in the list has several attributes associated with it, so it is a structure and is opened and closed with braces. At a minimum, a subnet definition must have at least two parameters: `subnet`, which defines the whole subnet; and `pools`, which is a list of dynamically allocated pools that are governed by the DHCP server.

The example contains a single subnet. If more than one were defined, additional elements in the `subnet4` parameter would be specified and separated by commas. For example, to define three subnets, the following syntax would be used:

```
{
"subnet4": [
  {
    "id": 1,
    "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
    "subnet": "192.0.2.0/24"
  },
  {
    "id": 2,
    "pools": [ { "pool": "192.0.3.100 - 192.0.3.200" } ],
    "subnet": "192.0.3.0/24"
  },
  {
    "id": 3,
    "pools": [ { "pool": "192.0.4.1 - 192.0.4.254" } ],
    "subnet": "192.0.4.0/24"
  }
],
...
}
```

Note that indentation is optional and is used for aesthetic purposes only. In some cases it may be preferable to use more compact notation.

After all the parameters have been specified, there are two contexts open: `global` and `Dhcp4`; thus, two closing curly brackets must be used to close them.

9.2.2 Lease Storage

All leases issued by the server are stored in the lease database. There are three database backends available: `memfile` (the default), `MySQL`, `PostgreSQL`.

9.2.2.1 Memfile - Basic Storage for Leases

The server is able to store lease data in different repositories. Larger deployments may elect to store leases in a database; *Lease Database Configuration* describes this option. In typical smaller deployments, though, the server stores lease information in a CSV file rather than a database. As well as requiring less administration, an advantage of using a file for storage is that it eliminates a dependency on third-party database software.

The configuration of the `memfile` backend is controlled through the `Dhcp4/lease-database` parameters. The `type` parameter is mandatory and specifies which storage for leases the server should use, through the `"memfile"` value. The following list gives additional optional parameters that can be used to configure the `memfile` backend.

- **`persist`**: controls whether the new leases and updates to existing leases are written to the file. It is strongly recommended that the value of this parameter be set to `true` at all times during the server's normal operation. Not writing leases to disk means that if a server is restarted (e.g. after a power failure), it will not know which

addresses have been assigned. As a result, it may assign new clients addresses that are already in use. The value of `false` is mostly useful for performance-testing purposes. The default value of the `persist` parameter is `true`, which enables writing lease updates to the lease file.

- **name:** specifies the lease file in which new leases and lease updates are recorded. The default value for this parameter is `"[kea-install-dir]/var/lib/kea/kea-leases4.csv"`.

Note

As of Kea 2.7.9, lease files may only be loaded from the data directory determined during compilation: `"[kea-install-dir]/var/lib/kea"`. This path may be overridden at startup by setting the environment variable `KEA_DHCP_DATA_DIR` to the desired path. If a path other than this value is used in **name**, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in specifying a custom file name simply omit the path component from **name**.

- **lfc-interval:** specifies the interval, in seconds, at which the server will perform a lease file cleanup (LFC). This removes redundant (historical) information from the lease file and effectively reduces the lease file size. The cleanup process is described in more detail later in this section. The default value of the `lfc-interval` is `3600`. A value of `0` disables the LFC.
- **max-row-errors:** specifies the number of row errors before the server stops attempting to load a lease file. When the server loads a lease file, it is processed row by row, each row containing a single lease. If a row is flawed and cannot be processed correctly the server logs it, discards the row, and goes on to the next row. This parameter can be used to set a limit on the number of such discards that can occur, after which the server abandons the effort and exits. The default value of `0` disables the limit and allows the server to process the entire file, regardless of how many rows are discarded.

An example configuration of the memfile backend is presented below:

```
"Dhcp4": {
  "lease-database": {
    "type": "memfile",
    "persist": true,
    "name": "kea-leases4.csv",
    "lfc-interval": 1800,
    "max-row-errors": 100
  }
}
```

This configuration selects `kea-leases4.csv` as the storage for lease information and enables persistence (writing lease updates to this file). It also configures the backend to perform a periodic cleanup of the lease file every 1800 seconds (30 minutes) and sets the maximum number of row errors to 100.

9.2.2.2 Why Is Lease File Cleanup Necessary?

It is important to know how the lease file contents are organized to understand why the periodic lease file cleanup is needed. Every time the server updates a lease or creates a new lease for a client, the new lease information must be recorded in the lease file. For performance reasons, the server does not update the existing client's lease in the file, as this would potentially require rewriting the entire file. Instead, it simply appends the new lease information to the end of the file; the previous lease entries for the client are not removed. When the server loads leases from the lease file, e.g. at server startup, it assumes that the latest lease entry for the client is the valid one. Previous entries are discarded, meaning that the server can reconstruct accurate information about the leases even though there may be many lease entries for each client. However, storing many entries for each client results in a bloated lease file and impairs the performance of the server's startup and reconfiguration, as it needs to process a larger number of lease entries.

Lease file cleanup (LFC) removes all previous entries for each client and leaves only the latest ones. The interval at

which the cleanup is performed is configurable, and it should be selected according to the frequency of lease renewals initiated by the clients. The more frequent the renewals, the smaller the value of `lfc-interval` should be. Note, however, that the LFC takes time and thus it is possible (although unlikely) that, if the `lfc-interval` is too short, a new cleanup may be started while the previous one is still running. The server would recover from this by skipping the new cleanup when it detected that the previous cleanup was still in progress, but it implies that the actual cleanups will be triggered more rarely than the configured interval. Moreover, triggering a new cleanup adds overhead to the server, which is not able to respond to new requests for a short period of time when the new cleanup process is spawned. Therefore, it is recommended that the `lfc-interval` value be selected in a way that allows the LFC to complete the cleanup before a new cleanup is triggered.

Lease file cleanup is performed by a separate process (in the background) to avoid a performance impact on the server process. To avoid conflicts between two processes using the same lease files, the LFC process starts with Kea opening a new lease file; the actual LFC process operates on the lease file that is no longer used by the server. There are also other files created as a side effect of the lease file cleanup. The detailed description of the LFC process is located later in this Kea Administrator's Reference Manual: *The LFC Process*.

9.2.2.3 Lease Database Configuration

Note

Lease database access information must be configured for the DHCPv4 server, even if it has already been configured for the DHCPv6 server. The servers store their information independently, so each server can use a separate database or both servers can use the same database.

Note

Kea requires the database timezone to match the system timezone. For more details, see *First-Time Creation of the MySQL Database* and *First-Time Creation of the PostgreSQL Database*.

Lease database configuration is controlled through the `Dhcp4/lease-database` parameters. The database type must be set to `memfile`, `mysql` or `postgresql`, e.g.:

```
"Dhcp4": { "lease-database": { "type": "mysql", ... }, ... }
```

Next, the name of the database to hold the leases must be set; this is the name used when the database was created (see *First-Time Creation of the MySQL Database* or *First-Time Creation of the PostgreSQL Database*).

For MySQL or PostgreSQL:

```
"Dhcp4": { "lease-database": { "name": "database-name", ... }, ... }
```

If the database is located on a different system from the DHCPv4 server, the database host name must also be specified:

```
"Dhcp4": { "lease-database": { "host": "remote-host-name", ... }, ... }
```

Normally, the database is on the same machine as the DHCPv4 server. In this case, set the value to the empty string:

```
"Dhcp4": { "lease-database": { "host": "", ... }, ... }
```

Should the database use a port other than the default, it may be specified as well:

```
"Dhcp4": { "lease-database": { "port": 12345, ... }, ... }
```

Should the database be located on a different system, the administrator may need to specify a longer interval for the connection timeout:

```
"Dhcp4": { "lease-database": { "connect-timeout" : timeout-in-seconds, ... }, ... }
```

The default value of five seconds should be more than adequate for local connections. If a timeout is given, though, it should be an integer greater than zero.

The maximum number of times the server automatically attempts to reconnect to the lease database after connectivity has been lost may be specified:

```
"Dhcp4": { "lease-database": { "max-reconnect-tries" : number-of-tries, ... }, ... }
```

If the server is unable to reconnect to the database after making the maximum number of attempts, the server will exit. A value of 0 (the default) disables automatic recovery and the server will exit immediately upon detecting a loss of connectivity (MySQL and PostgreSQL only).

The number of milliseconds the server waits between attempts to reconnect to the lease database after connectivity has been lost may also be specified:

```
"Dhcp4": { "lease-database": { "reconnect-wait-time" : number-of-milliseconds, ... }, ... }
↪ }
```

The default value for MySQL and PostgreSQL is 0, which disables automatic recovery and causes the server to exit immediately upon detecting the loss of connectivity.

```
"Dhcp4": { "lease-database": { "on-fail" : "stop-retry-exit", ... }, ... }
```

The possible values are:

- **stop-retry-exit** - disables the DHCP service while trying to automatically recover lost connections, and shuts down the server on failure after exhausting `max-reconnect-tries`. This is the default value for the lease backend, the host backend, and the configuration backend.
- **serve-retry-exit** - continues the DHCP service while trying to automatically recover lost connections, and shuts down the server on failure after exhausting `max-reconnect-tries`.
- **serve-retry-continue** - continues the DHCP service and does not shut down the server even if the recovery fails. This is the default value for forensic logging.

Note

Automatic reconnection to database backends is configured individually per backend; this allows users to tailor the recovery parameters to each backend they use. We suggest that users enable it either for all backends or none, so behavior is consistent.

Losing connectivity to a backend for which reconnection is disabled results (if configured) in the server shutting itself down. This includes cases when the lease database backend and the hosts database backend are connected to the same database instance.

It is highly recommended not to change the `stop-retry-exit` default setting for the lease manager, as it is critical for the connection to be active while processing DHCP traffic. Change this only if the server is used exclusively as a configuration tool.

```
"Dhcp4": { "lease-database": { "retry-on-startup" : true, ... }, ... }
```

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server starts reconnection attempts even at server startup or on reconfigure events, and honors the action specified in the `on-fail` parameter.

The `host` parameter is used by the MySQL and PostgreSQL backends.

Finally, the credentials of the account under which the server will access the database should be set:

```
"Dhcp4": {
  "lease-database": {
    "user": "user-name",
    "password": "1234",
    ...
  },
  ...
}
```

If there is no password to the account, set the password to the empty string `""`. (This is the default.)

9.2.2.4 Tuning Database Timeouts

In rare cases, reading or writing to the database may hang. This can be caused by a temporary network issue, or by mis-configuration of the proxy server switching the connection between different database instances. These situations are rare, but users have reported that Kea sometimes hangs while performing database IO operations. Setting appropriate timeout values can mitigate such issues.

MySQL exposes two distinct connection options to configure the read and write timeouts. Kea's corresponding `read-timeout` and `write-timeout` configuration parameters specify the timeouts in seconds. For example:

```
"Dhcp4": { "lease-database": { "read-timeout" : 10, "write-timeout": 20, ... }, ... }
```

Setting these parameters to 0 is equivalent to not specifying them, and causes the Kea server to establish a connection to the database with the MySQL defaults. In this case, Kea waits indefinitely for the completion of the read and write operations.

MySQL versions earlier than 5.6 do not support setting timeouts for read and write operations. Moreover, the `read-timeout` and `write-timeout` parameters can only be specified for the MySQL backend; setting them for any other backend database type causes a configuration error.

To set a timeout in seconds for PostgreSQL, use the `tcp-user-timeout` parameter. For example:

```
"Dhcp4": { "lease-database": { "tcp-user-timeout" : 10, ... }, ... }
```

Specifying this parameter for other backend types causes a configuration error.

Note

The timeouts described here are only effective for TCP connections. Please note that the MySQL client library used by the Kea servers typically connects to the database via a UNIX domain socket when the `host` parameter is `localhost`, but establishes a TCP connection for `127.0.0.1`.

Since Kea.2.7.4, the `libdhcp_mysql.so` hook library must be loaded in order to store leases in the MySQL Lease Database Backend. Specify the lease backend hook library location:

```
"Dhcp4": { "hooks-libraries": [
  {
    // the MySQL lease backend hook library required for lease storage.
    "library": "/opt/lib/kea/hooks/libdhcp_mysql.so"
  }, ... ], ... }
```

Since Kea.2.7.4, the libdhcp_pgsql.so hook library must be loaded in order to store leases in the PostgreSQL Lease Database Backend. Specify the lease backend hook library location.

```
"Dhcp4": { "hooks-libraries": [
  {
    // the PostgreSQL lease backend hook library required for lease storage.
    "library": "/opt/lib/kea/hooks/libdhcp_pgsql.so"
  }, ... ], ... }
```

9.2.3 Hosts Storage

Kea is also able to store information about host reservations in the database. The hosts database configuration uses the same syntax as the lease database. In fact, the Kea server opens independent connections for each purpose, be it lease or hosts information, which gives the most flexibility. Kea can keep leases and host reservations separately, but can also point to the same database. Currently the supported hosts database types are MySQL and PostgreSQL.

The following configuration can be used to configure a connection to MySQL:

```
"Dhcp4": {
  "hosts-database": {
    "type": "mysql",
    "name": "kea",
    "user": "kea",
    "password": "1234",
    "host": "localhost",
    "port": 3306
  }
}
```

Depending on the database configuration, many of the parameters may be optional.

Please note that usage of hosts storage is optional. A user can define all host reservations in the configuration file, and that is the recommended way if the number of reservations is small. However, when the number of reservations grows, it is more convenient to use host storage. Please note that both storage methods (the configuration file and one of the supported databases) can be used together. If hosts are defined in both places, the definitions from the configuration file are checked first and external storage is checked later, if necessary.

Host information can be placed in multiple stores. Operations are performed on the stores in the order they are defined in the configuration file, although this leads to a restriction in ordering in the case of a host reservation addition; read-only stores must be configured after a (required) read-write store, or the addition will fail.

Note

Kea requires the database timezone to match the system timezone. For more details, see *First-Time Creation of the MySQL Database* and *First-Time Creation of the PostgreSQL Database*.

9.2.3.1 DHCPv4 Hosts Database Configuration

Hosts database configuration is controlled through the `Dhcp4/hosts-database` parameters. If enabled, the type of database must be set to `mysql` or `postgresql`.

```
"Dhcp4": { "hosts-database": { "type": "mysql", ... }, ... }
```

Next, the name of the database to hold the reservations must be set; this is the name used when the lease database was created (see *Supported Backends* for instructions on how to set up the desired database type):

```
"Dhcp4": { "hosts-database": { "name": "database-name" , ... }, ... }
```

If the database is located on a different system than the DHCPv4 server, the database host name must also be specified:

```
"Dhcp4": { "hosts-database": { "host": remote-host-name, ... }, ... }
```

Normally, the database is on the same machine as the DHCPv4 server. In this case, set the value to the empty string:

```
"Dhcp4": { "hosts-database": { "host" : "", ... }, ... }
```

Should the database use a port different than the default, it may be specified as well:

```
"Dhcp4": { "hosts-database": { "port" : 12345, ... }, ... }
```

The maximum number of times the server automatically attempts to reconnect to the host database after connectivity has been lost may be specified:

```
"Dhcp4": { "hosts-database": { "max-reconnect-tries" : number-of-tries, ... }, ... }
```

If the server is unable to reconnect to the database after making the maximum number of attempts, the server will exit. A value of 0 (the default) disables automatic recovery and the server will exit immediately upon detecting a loss of connectivity (MySQL and PostgreSQL only).

The number of milliseconds the server waits between attempts to reconnect to the host database after connectivity has been lost may also be specified:

```
"Dhcp4": { "hosts-database": { "reconnect-wait-time" : number-of-milliseconds, ... }, ... }
↪ }
```

The default value for MySQL and PostgreSQL is 0, which disables automatic recovery and causes the server to exit immediately upon detecting the loss of connectivity.

```
"Dhcp4": { "hosts-database": { "on-fail" : "stop-retry-exit", ... }, ... }
```

The possible values are:

- `stop-retry-exit` - disables the DHCP service while trying to automatically recover lost connections. Shuts down the server on failure after exhausting `max-reconnect-tries`. This is the default value for MySQL and PostgreSQL.
- `serve-retry-exit` - continues the DHCP service while trying to automatically recover lost connections. Shuts down the server on failure after exhausting `max-reconnect-tries`.
- `serve-retry-continue` - continues the DHCP service and does not shut down the server even if the recovery fails.

Note

Automatic reconnection to database backends is configured individually per backend. This allows users to tailor the recovery parameters to each backend they use. We suggest that users enable it either for all backends or none, so behavior is consistent.

Losing connectivity to a backend for which reconnection is disabled results (if configured) in the server shutting itself down. This includes cases when the lease database backend and the hosts database backend are connected to the same database instance.

```
"Dhcp4": { "hosts-database": { "retry-on-startup" : true, ... }, ... }
```

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server starts reconnection attempts even at server startup or on reconfigure events, and honors the action specified in the `on-fail` parameter.

Finally, the credentials of the account under which the server will access the database should be set:

```
"Dhcp4": {
  "hosts-database": {
    "user": "user-name",
    "password": "1234",
    ...
  },
  ...
}
```

If there is no password to the account, set the password to the empty string `""`. (This is the default.)

The multiple-storage extension uses a similar syntax; a configuration is placed into a `hosts-databases` list instead of into a `hosts-database` entry, as in:

```
"Dhcp4": { "hosts-databases": [ { "type": "mysql", ... }, ... ], ... }
```

If the same host is configured both in-file and in-database, Kea does not issue a warning, as it would if both were specified in the same data source. Instead, the host configured in-file has priority over the one configured in-database.

9.2.3.2 Using Read-Only Databases for Host Reservations With DHCPv4

In some deployments, the user whose name is specified in the database backend configuration may not have write privileges to the database. This is often required by the policy within a given network to secure the data from being unintentionally modified. In many cases administrators have deployed inventory databases, which contain substantially more information about the hosts than just the static reservations assigned to them. The inventory database can be used to create a view of a Kea hosts database and such a view is often read-only.

Kea host-database backends operate with an implicit configuration to both read from and write to the database. If the user does not have write access to the host database, the backend will fail to start and the server will refuse to start (or reconfigure). However, if access to a read-only host database is required for retrieving reservations for clients and/or assigning specific addresses and options, it is possible to explicitly configure Kea to start in "read-only" mode. This is controlled by the `readonly` boolean parameter as follows:

```
"Dhcp4": { "hosts-database": { "readonly": true, ... }, ... }
```

Setting this parameter to `false` configures the database backend to operate in "read-write" mode, which is also the default configuration if the parameter is not specified.

Note

The `readonly` parameter is only supported for MySQL and PostgreSQL databases.

Since Kea.2.7.4, the `libdhcp_mysql.so` hook library must be loaded in order to store host reservations in the MySQL Host Database Backend. Specify the lease backend hook library location:

```
"Dhcp4": { "hooks-libraries": [
  {
    // the MySQL host backend hook library required for host storage.
    "library": "/opt/lib/kea/hooks/libdhcp_mysql.so"
  }, ... ], ... }
```

Since Kea.2.7.4, the `libdhcp_pgsql.so` hook library must be loaded in order to store host reservations in the PostgreSQL Host Database Backend. Specify the lease backend hook library location.

```
"Dhcp4": { "hooks-libraries": [
  {
    // the PostgreSQL host backend hook library required for host storage.
    "library": "/opt/lib/kea/hooks/libdhcp_pgsql.so"
  }, ... ], ... }
```

9.2.3.3 Tuning Database Timeouts for Hosts Storage

See *Tuning Database Timeouts*.

9.2.4 Interface Configuration

The DHCPv4 server must be configured to listen on specific network interfaces. The simplest network interface configuration tells the server to listen on all available interfaces:

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "*" ]
  },
  ...
}
```

The asterisk plays the role of a wildcard and means "listen on all interfaces." However, it is usually a good idea to explicitly specify interface names:

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ]
  },
  ...
}
```

It is possible to use an interface wildcard (*) concurrently with explicit interface names:


```
"Dhcp4": {  
  "interfaces-config": {  
    "interfaces": [ "eth1", "eth3", "*" ]  
  },  
  ...  
}
```

This format should only be used when it is desired to temporarily override a list of interface names and listen on all interfaces.

Some deployments of DHCP servers require that the servers listen on interfaces with multiple IPv4 addresses configured. In these situations, the address to use can be selected by appending an IPv4 address to the interface name in the following manner:

```
"Dhcp4": {  
  "interfaces-config": {  
    "interfaces": [ "eth1/10.0.0.1", "eth3/192.0.2.3" ]  
  },  
  ...  
}
```

Should the server be required to listen on multiple IPv4 addresses assigned to the same interface, multiple addresses can be specified for an interface as in the example below:

```
"Dhcp4": {  
  "interfaces-config": {  
    "interfaces": [ "eth1/10.0.0.1", "eth1/10.0.0.2" ]  
  },  
  ...  
}
```

Alternatively, if the server should listen on all addresses for the particular interface, an interface name without any address should be specified.

Kea supports responding to directly connected clients which do not have an address configured. This requires the server to inject the hardware address of the destination into the data-link layer of the packet being sent to the client. The DHCPv4 server uses raw sockets to achieve this, and builds the entire IP/UDP stack for the outgoing packets. The downside of raw socket use, however, is that incoming and outgoing packets bypass the firewalls (e.g. iptables).

Handling traffic on multiple IPv4 addresses assigned to the same interface can be a challenge, as raw sockets are bound to the interface. When the DHCP server is configured to use the raw socket on an interface to receive DHCP traffic, advanced packet filtering techniques (e.g. the BPF) must be used to receive unicast traffic on the desired addresses assigned to the interface. Whether clients use the raw socket or the UDP socket depends on whether they are directly connected (raw socket) or relayed (either raw or UDP socket).

Therefore, in deployments where the server does not need to provision the directly connected clients and only receives the unicast packets from the relay agents, the Kea server should be configured to use UDP sockets instead of raw sockets. The following configuration demonstrates how this can be achieved:

```
"Dhcp4": {  
  "interfaces-config": {  
    "interfaces": [ "eth1", "eth3" ],  
    "dhcp-socket-type": "udp"  
  },  
  ...  
}
```

The `dhcp-socket-type` parameter specifies that the IP/UDP sockets will be opened on all interfaces on which the server listens, i.e. "eth1" and "eth3" in this example. If `dhcp-socket-type` is set to `raw`, it configures the server to use raw sockets instead. If the `dhcp-socket-type` value is not specified, the default value `raw` is used.

Using UDP sockets automatically disables the reception of broadcast packets from directly connected clients. This effectively means that UDP sockets can be used for relayed traffic only. When using raw sockets, both the traffic from the directly connected clients and the relayed traffic are handled.

Caution should be taken when configuring the server to open multiple raw sockets on the interface with several IPv4 addresses assigned. If the directly connected client sends the message to the broadcast address, all sockets on this link will receive this message and multiple responses will be sent to the client. Therefore, the configuration with multiple IPv4 addresses assigned to the interface should not be used when the directly connected clients are operating on that link. To use a single address on such an interface, the "interface-name/address" notation should be used.

Note

Specifying the value `raw` as the socket type does not guarantee that raw sockets will be used! The use of raw sockets to handle traffic from the directly connected clients is currently supported on Linux and BSD systems only. If raw sockets are not supported on the particular OS in use, the server issues a warning and fall back to using IP/UDP sockets.

In a typical environment, the DHCP server is expected to send back a response on the same network interface on which the query was received. This is the default behavior. However, in some deployments it is desired that the outbound (response) packets be sent as regular traffic and the outbound interface be determined by the routing tables. This kind of asymmetric traffic is uncommon, but valid. Kea supports a parameter called `outbound-interface` that controls this behavior. It supports two values: the first one, `same-as-inbound`, tells Kea to send back the response on the same interface where the query packet was received. This is the default behavior. The second parameter, `use-routing`, tells Kea to send regular UDP packets and let the kernel's routing table determine the most appropriate interface. This only works when `dhcp-socket-type` is set to `udp`. An example configuration looks as follows:

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "dhcp-socket-type": "udp",
    "outbound-interface": "use-routing"
  },
  ...
}
```

Interfaces are re-detected at each reconfiguration. This behavior can be disabled by setting the `re-detect` value to `false`, for instance:

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "re-detect": false
  },
  ...
}
```

Note that interfaces are not re-detected during `config-test`.

Usually loopback interfaces (e.g. the `lo` or `lo0` interface) are not configured, but if a loopback interface is explicitly configured and IP/UDP sockets are specified, the loopback interface is accepted.

For example, this setup can be used to run Kea in a FreeBSD jail having only a loopback interface, to service a relayed DHCP request:

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "lo0" ],
    "dhcp-socket-type": "udp"
  },
  ...
}
```

Kea binds the service sockets for each interface on startup. If another process is already using a port, then Kea logs the message and suppresses an error. DHCP service runs, but it is unavailable on some interfaces.

The "service-sockets-require-all" option makes Kea require all sockets to be successfully bound. If any opening fails, Kea interrupts the initialization and exits with a non-zero status. (Default is false).

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "service-sockets-require-all": true
  },
  ...
}
```

Sometimes, immediate interruption isn't a good choice. The port can be unavailable only temporary. In this case, retrying the opening may resolve the problem. Kea provides two options to specify the retrying: `service-sockets-max-retries` and `service-sockets-retry-wait-time`.

The first defines a maximal number of retries that Kea makes to open a socket. The zero value (default) means that the Kea doesn't retry the process.

The second defines a wait time (in milliseconds) between attempts. The default value is 5000 (5 seconds).

```
"Dhcp4": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "service-sockets-max-retries": 5,
    "service-sockets-retry-wait-time": 5000
  },
  ...
}
```

If "service-sockets-max-retries" is non-zero and "service-sockets-require-all" is false, then Kea retries the opening (if needed) but does not fail if any socket is still not opened.

9.2.5 Issues With Unicast Responses to DHCPINFORM

The use of UDP sockets has certain benefits in deployments where the server receives only relayed traffic; these benefits are mentioned in *Interface Configuration*. From the administrator's perspective it is often desirable to configure the system's firewall to filter out unwanted traffic, and the use of UDP sockets facilitates this. However, the administrator must also be aware of the implications related to filtering certain types of traffic, as it may impair the DHCP server's operation.

In this section we focus on the case when the server receives the DHCPINFORM message from the client via a relay. According to [RFC 2131](#), the server should unicast the DHCPACK response to the address carried in the `ciaddr` field. When the UDP socket is in use, the DHCP server relies on the low-level functions of an operating system to build the

data link, IP, and UDP layers of the outgoing message. Typically, the OS first uses ARP to obtain the client's link-layer address to be inserted into the frame's header, if the address is not cached from a previous transaction that the client had with the server. When the ARP exchange is successful, the DHCP message can be unicast to the client, using the obtained address.

Some system administrators block ARP messages in their network, which causes issues for the server when it responds to the DHCPINFORM messages because the server is unable to send the DHCPACK if the preceding ARP communication fails. Since the OS is entirely responsible for the ARP communication and then sending the DHCP packet over the wire, the DHCP server has no means to determine that the ARP exchange failed and the DHCP response message was dropped. Thus, the server does not log any error messages when the outgoing DHCP response is dropped. At the same time, all hooks pertaining to the packet-sending operation will be called, even though the message never reaches its destination.

Note that the issue described in this section is not observed when raw sockets are in use, because, in this case, the DHCP server builds all the layers of the outgoing message on its own and does not use ARP. Instead, it inserts the value carried in the `chaddr` field of the DHCPINFORM message into the link layer.

Server administrators willing to support DHCPINFORM messages via relays should not block ARP traffic in their networks, or should use raw sockets instead of UDP sockets.

9.2.6 IPv4 Subnet Identifier

The subnet identifier (subnet ID) is a unique number associated with a particular subnet. In principle, it is used to associate clients' leases with their respective subnets. The server configuration must contain unique and stable identifiers for all subnets.

Note

Subnet IDs must be greater than zero and less than 4294967295.

The following configuration assigns the specified subnet identifier to a newly configured subnet:

```
"Dhcp4": {
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "id": 1024,
      ...
    }
  ]
}
```

9.2.7 IPv4 Subnet Prefix

The subnet prefix is the second way to identify a subnet. Kea can accept non-canonical subnet addresses; for instance, this configuration is accepted:

```
"Dhcp4": {
  "subnet4": [
    {
      "subnet": "192.0.2.1/24",
      ...
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
]
}
```

This works even if there is another subnet with the "192.0.2.0/24" prefix; only the textual form of subnets are compared to avoid duplicates.

Note

Abuse of this feature can lead to incorrect subnet selection (see *How the DHCPv4 Server Selects a Subnet for the Client*).

9.2.8 Configuration of IPv4 Address Pools

The main role of a DHCPv4 server is address assignment. For this, the server must be configured with at least one subnet and one pool of dynamic addresses to be managed. For example, assume that the server is connected to a network segment that uses the 192.0.2.0/24 prefix. The administrator of that network decides that addresses from the range 192.0.2.10 to 192.0.2.20 are going to be managed by the DHCPv4 server. Such a configuration can be achieved in the following way:

```
"Dhcp4": {
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "pools": [
        { "pool": "192.0.2.10 - 192.0.2.20" }
      ],
      ...
    }
  ]
}
```

Note that `subnet` is defined as a simple string, but the `pools` parameter is actually a list of pools; for this reason, the pool definition is enclosed in square brackets, even though only one range of addresses is specified.

Each `pool` is a structure that contains the parameters that describe a single pool. Currently there is only one parameter, `pool`, which gives the range of addresses in the pool.

It is possible to define more than one pool in a subnet; continuing the previous example, further assume that 192.0.2.64/26 should also be managed by the server. It could be written as 192.0.2.64 to 192.0.2.127, or it can be expressed more simply as 192.0.2.64/26. Both formats are supported by `Dhcp4` and can be mixed in the pool list. For example, the following pools could be defined:

```
"Dhcp4": {
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "pools": [
        { "pool": "192.0.2.10-192.0.2.20" },
        { "pool": "192.0.2.64/26" }
      ],
      ...
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    ...
}

```

White space in pool definitions is ignored, so spaces before and after the hyphen are optional. They can be used to improve readability.

The number of pools is not limited, but for performance reasons it is recommended to use as few as possible.

The server may be configured to serve more than one subnet. To add a second subnet, use a command similar to the following:

```

"Dhcp4": {
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
      ...
    },
    {
      "subnet": "192.0.3.0/24",
      "pools": [ { "pool": "192.0.3.100 - 192.0.3.200" } ],
      ...
    },
    {
      "subnet": "192.0.4.0/24",
      "pools": [ { "pool": "192.0.4.1 - 192.0.4.254" } ],
      ...
    }
  ]
}

```

When configuring a DHCPv4 server using prefix/length notation, please pay attention to the boundary values. When specifying that the server can use a given pool, it is also able to allocate the first (typically a network address) and the last (typically a broadcast address) address from that pool. In the aforementioned example of pool 192.0.3.0/24, both the 192.0.3.0 and 192.0.3.255 addresses may be assigned as well. This may be invalid in some network configurations. To avoid this, use the `min-max` notation.

In a subnet whose prefix length is less than 24, users may wish to exclude all addresses ending in .0 and .255 from being dynamically allocated. For instance, in the subnet 10.0.0.0/8, an administrator may wish to exclude 10.x.y.0 and 10.x.y.255 for all values of x and y, even though only 10.0.0.0 and 10.255.255.255 must be excluded according to RFC standards. The `exclude-first-last-24` configuration compatibility flag (*Kea DHCPv4 Compatibility Configuration Parameters*) does this automatically, rather than requiring explicit configuration of many pools or reservations for fake hosts. When `true`, it applies only to subnets of prefix length 24 or smaller, i.e. larger address space; the default is `false`.

In this case, "exclude" means to skip these addresses in the free address pickup routine of the allocation engine; if a client explicitly requests or has a host reservation for an address in .0 or .255, it will get it.

Note

Here are some liberties and limits to the values that subnets and pools can take in unusual Kea configurations:

Kea configuration case	Allow	Comment
Overlapping subnets	Yes	Administrator should consider how clients are matched to these subnets.
Overlapping pools in one subnet	No	Startup error: DHCP4_PARSER_FAIL
Overlapping address pools in different subnets	Yes	Specifying the same address pool in different subnets can be used as an equivalent of the global address pool. In that case, the server can assign addresses from the same range regardless of the client's subnet. If an address from such a pool is assigned to a client in one subnet, the same address will be renewed for this client if it moves to another subnet. Another client in a different subnet will not be assigned an address already assigned to the client in any of the subnets.
Pools not matching the subnet prefix	No	Startup error: DHCP4_PARSER_FAIL

9.2.9 Sending T1 (Option 58) and T2 (Option 59)

According to [RFC 2131](#), servers should send values for T1 and T2 that are 50% and 87.5% of the lease lifetime, respectively. By default, *kea-dhcp4* does not send either value; it can be configured to send values that are either specified explicitly or that are calculated as percentages of the lease time. The server's behavior is governed by a combination of configuration parameters, two of which have already been mentioned. To send specific, fixed values use the following two parameters:

- **renew-timer** - specifies the value of T1 in seconds.
- **rebind-timer** - specifies the value of T2 in seconds.

The server only sends T2 if it is less than the valid lease time. T1 is only sent if T2 is being sent and T1 is less than T2; or T2 is not being sent and T1 is less than the valid lease time.

Calculating the values is controlled by the following three parameters.

- **calculate-tee-times** - when true, T1 and T2 are calculated as percentages of the valid lease time. It defaults to false.
- **t1-percent** - the percentage of the valid lease time to use for T1. It is expressed as a real number between 0.0 and 1.0 and must be less than **t2-percent**. The default value is 0.50, per RFC 2131.
- **t2-percent** - the percentage of the valid lease time to use for T2. It is expressed as a real number between 0.0 and 1.0 and must be greater than **t1-percent**. The default value is .875, per RFC 2131.

Note

In the event that both explicit values are specified and **calculate-tee-times** is true, the server will use the

explicit values. Administrators with a setup where some subnets or shared-networks use explicit values and some use calculated values must not define the explicit values at any level higher than where they will be used. Inheriting them from too high a scope, such as global, will cause them to have explicit values at every level underneath (shared-networks and subnets), effectively disabling calculated values.

9.2.10 Standard DHCPv4 Options

One of the major features of the DHCPv4 server is the ability to provide configuration options to clients. Most of the options are sent by the server only if the client explicitly requests them using the Parameter Request List option. Those that do not require inclusion in the Parameter Request List option are commonly used options, e.g. "Domain Server", and options which require special behavior, e.g. "Client FQDN", which is returned to the client if the client has included this option in its message to the server.

List of standard DHCPv4 options configurable by an administrator comprises the list of the standard DHCPv4 options whose values can be configured using the configuration structures described in this section. This table excludes the options which require special processing and thus cannot be configured with fixed values. The last column of the table indicates which options can be sent by the server even when they are not requested in the Parameter Request List option, and those which are sent only when explicitly requested.

The following example shows how to configure the addresses of DNS servers, which is one of the most frequently used options. Options specified in this way are considered global and apply to all configured subnets.

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "code": 6,
      "space": "dhcp4",
      "csv-format": true,
      "data": "192.0.2.1, 192.0.2.2"
    },
    ...
  ]
}
```

Note that either name or code is required; there is no need to specify both. space has a default value of dhcp4, so this can be skipped as well if a regular (not encapsulated) DHCPv4 option is defined. Finally, csv-format defaults to true, so it too can be skipped, unless the option value is specified as a hexadecimal string. Therefore, the above example can be simplified to:

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "data": "192.0.2.1, 192.0.2.2"
    },
    ...
  ]
}
```

Defined options are added to the response when the client requests them, with a few exceptions which are always added. To enforce the addition of a particular option, set the always-send flag to true as in:

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "data": "192.0.2.1, 192.0.2.2",
      "always-send": true
    },
    ...
  ]
}
```

The effect is the same as if the client added the option code in the Parameter Request List option (or its equivalent for vendor options):

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "data": "192.0.2.1, 192.0.2.2",
      "always-send": true
    },
    ...
  ],
  "subnet4": [
    {
      "subnet": "192.0.3.0/24",
      "option-data": [
        {
          "name": "domain-name-servers",
          "data": "192.0.3.1, 192.0.3.2"
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}
```

In the example above, the `domain-name-servers` option respects the global `always-send` flag and is always added to responses, but for subnet `192.0.3.0/24`, the value is taken from the subnet-level option data specification.

Contrary to `always-send`, if the `never-send` flag is set to `true` for a particular option, the server does not add it to the response. The effect is the same as if the client removed the option code in the Parameter Request List option (or its equivalent for vendor options):

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "data": "192.0.2.1, 192.0.2.2"
    },
    ...
  ]
}
```

(continues on next page)

(continued from previous page)

```

],
"subnet4": [
  {
    "subnet": "192.0.3.0/24",
    "option-data": [
      {
        "name": "domain-name-servers",
        "never-send": true
      },
      ...
    ],
    ...
  },
  ...
],
...
}

```

In the example above, the `domain-name-servers` option is never added to responses on subnet `192.0.3.0/24`. `never-send` has precedence over `always-send`, so if both are `true` the option is not added.

Note

The `always-send` and `never-send` flags are sticky, meaning they do not follow the usual configuration inheritance rules. Instead, if they are enabled at least once along the configuration inheritance chain, they are applied - even if they are disabled in other places which would normally receive a higher priority. For instance, if one of the flags is enabled in the global scope, but disabled at the subnet level, it is enabled, disregarding the subnet-level setting.

Note

The `never-send` flag is less powerful than *[libdhcp_flex_option.so](#)*; for instance, it has no effect on options managed by the server itself. Both `always-send` and `never-send` have no effect on options which cannot be requested, for instance from a custom space.

Note

Beginning with Kea 2.7.4, option inclusion can also be controlled through option class-tagging, see *[Option Class-Tagging](#)*

The `name` parameter specifies the option name. For a list of currently supported names, see *[List of standard DHCPv4 options configurable by an administrator](#)* below. The `code` parameter specifies the option code, which must match one of the values from that list. The next line specifies the option space, which must always be set to `dhcp4` as these are standard DHCPv4 options. For other option spaces, including custom option spaces, see *[Nested DHCPv4 Options \(Custom Option Spaces\)](#)*. The next line specifies the format in which the data will be entered; use of CSV (comma-separated values) is recommended. The sixth line gives the actual value to be sent to clients. The `data` parameter is specified as normal text, with values separated by commas if more than one value is allowed.

Options can also be configured as hexadecimal values. If `csv-format` is set to `false`, option data must be specified as a hexadecimal string. The following commands configure the `domain-name-servers` option for all subnets with

the following addresses: 192.0.3.1 and 192.0.3.2. Note that `csv-format` is set to `false`.

```
"Dhcp4": {
  "option-data": [
    {
      "name": "domain-name-servers",
      "code": 6,
      "space": "dhcp4",
      "csv-format": false,
      "data": "C0 00 03 01 C0 00 03 02"
    },
    ...
  ],
  ...
}
```

Kea supports the following formats when specifying hexadecimal data:

- **Delimited octets** - one or more octets separated by either colons or spaces (":" or " "). While each octet may contain one or two digits, we strongly recommend always using two digits. Valid examples are "ab:cd:ef" and "ab cd ef".
- **String of digits** - a continuous string of hexadecimal digits with or without a "0x" prefix. Valid examples are "0xabcdef" and "abcdef".

Care should be taken to use proper encoding when using hexadecimal format; Kea's ability to validate data correctness in hexadecimal is limited.

It is also possible to specify data for binary options as a single-quoted text string within double quotes as shown (note that `csv-format` must be set to `false`):

```
"Dhcp4": {
  "option-data": [
    {
      "name": "user-class",
      "code": 77,
      "space": "dhcp4",
      "csv-format": false,
      "data": "'convert this text to binary'"
    },
    ...
  ],
  ...
}
```

Most of the parameters in the `option-data` structure are optional and can be omitted in some circumstances, as discussed in *Unspecified Parameters for DHCPv4 Option Configuration*.

It is possible to specify or override options on a per-subnet basis. If clients connected to most subnets are expected to get the same values of a given option, administrators should use global options. On the other hand, if different values are used in each subnet, it does not make sense to specify global option values; rather, only subnet-specific ones should be set.

The following commands override the global DNS servers option for a particular subnet, setting a single DNS server with address 192.0.2.3:

```

"Dhcp4": {
  "subnet4": [
    {
      "option-data": [
        {
          "name": "domain-name-servers",
          "code": 6,
          "space": "dhcp4",
          "csv-format": true,
          "data": "192.0.2.3"
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}

```

In some cases it is useful to associate some options with an address pool from which a client is assigned a lease. Pool-specific option values override subnet-specific and global option values; it is not possible to prioritize assignment of pool-specific options via the order of pool declarations in the server configuration.

The following configuration snippet demonstrates how to specify the DNS servers option, which is assigned to a client only if the client obtains an address from the given pool:

```

"Dhcp4": {
  "subnet4": [
    {
      "pools": [
        {
          "pool": "192.0.2.1 - 192.0.2.200",
          "option-data": [
            {
              "name": "domain-name-servers",
              "data": "192.0.2.3"
            },
            ...
          ],
          ...
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}

```

Options can also be specified in class or host-reservation scope. The current Kea options precedence order is (from most important to least): host reservation, pool, subnet, shared network, class, global.

When a data field is a string and that string contains the comma (,; U+002C) character, the comma must be escaped

with two backslashes (`\\`; U+005C). This double escape is required because both the routine splitting of CSV data into fields and JSON use the same escape character; a single escape (`\`,) would make the JSON invalid. For example, the string "foo,bar" must be represented as:

```
"Dhcp4": {
  "subnet4": [
    {
      "pools": [
        {
          "option-data": [
            {
              "name": "boot-file-name",
              "data": "foo\\,bar"
            }
          ]
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}
```

Some options are designated as arrays, which means that more than one value is allowed. For example, the option `time-servers` allows the specification of more than one IPv4 address, enabling clients to obtain the addresses of multiple NTP servers.

Custom DHCPv4 Options describes the configuration syntax to create custom option definitions (formats). Creation of custom definitions for standard options is generally not permitted, even if the definition being created matches the actual option format defined in the RFCs. However, there is an exception to this rule for standard options for which Kea currently does not provide a definition. To use such options, a server administrator must create a definition as described in *Custom DHCPv4 Options* in the `dhcp4` option space. This definition should match the option format described in the relevant RFC, but the configuration mechanism allows any option format as there is currently no way to validate it.

The currently supported standard DHCPv4 options are listed in the table below. "Name" and "Code" are the values that should be used as a name/code in the option-data structures. "Type" designates the format of the data; the meanings of the various types are given in *List of standard DHCP option types*.

Table 1: List of standard DHCPv4 options configurable by an administrator

Name	Code	Type	Array?	Returned if not requested
time-offset	2	int32	false	false
routers	3	ipv4-address	true	true
time-servers	4	ipv4-address	true	false
name-servers	5	ipv4-address	true	false
domain-name-servers	6	ipv4-address	true	true
log-servers	7	ipv4-address	true	false
cookie-servers	8	ipv4-address	true	false
lpr-servers	9	ipv4-address	true	false
impress-servers	10	ipv4-address	true	false
resource-location-servers	11	ipv4-address	true	false

continues on next

Table 1 – continued from previous page

Name	Code	Type	Array?	Returned if not requested
boot-size	13	uint16	false	false
merit-dump	14	string	false	false
domain-name	15	fqdn	false	true
swap-server	16	ipv4-address	false	false
root-path	17	string	false	false
extensions-path	18	string	false	false
ip-forwarding	19	boolean	false	false
non-local-source-routing	20	boolean	false	false
policy-filter	21	ipv4-address	true	false
max-dgram-reassembly	22	uint16	false	false
default-ip-ttl	23	uint8	false	false
path-mtu-aging-timeout	24	uint32	false	false
path-mtu-plateau-table	25	uint16	true	false
interface-mtu	26	uint16	false	false
all-subnets-local	27	boolean	false	false
broadcast-address	28	ipv4-address	false	false
perform-mask-discovery	29	boolean	false	false
mask-supplier	30	boolean	false	false
router-discovery	31	boolean	false	false
router-solicitation-address	32	ipv4-address	false	false
static-routes	33	ipv4-address	true	false
trailer-encapsulation	34	boolean	false	false
arp-cache-timeout	35	uint32	false	false
ieee802-3-encapsulation	36	boolean	false	false
default-tcp-ttl	37	uint8	false	false
tcp-keepalive-interval	38	uint32	false	false
tcp-keepalive-garbage	39	boolean	false	false
nis-domain	40	string	false	false
nis-servers	41	ipv4-address	true	false
ntp-servers	42	ipv4-address	true	false
vendor-encapsulated-options	43	empty	false	false
netbios-name-servers	44	ipv4-address	true	false
netbios-dd-server	45	ipv4-address	true	false
netbios-node-type	46	uint8	false	false
netbios-scope	47	string	false	false
font-servers	48	ipv4-address	true	false
x-display-manager	49	ipv4-address	true	false
dhcp-option-overload	52	uint8	false	false
dhcp-server-identifier	54	ipv4-address	false	true
dhcp-message	56	string	false	false
dhcp-max-message-size	57	uint16	false	false
vendor-class-identifier	60	string	false	false
nwip-domain-name	62	string	false	false
nwip-suboptions	63	binary	false	false
nisplus-domain-name	64	string	false	false
nisplus-servers	65	ipv4-address	true	false
tftp-server-name	66	string	false	false
boot-file-name	67	string	false	false
mobile-ip-home-agent	68	ipv4-address	true	false
smtp-server	69	ipv4-address	true	false

continues on next

Table 1 – continued from previous page

Name	Code	Type	Array?	Returned if not requested
pop-server	70	ipv4-address	true	false
nntp-server	71	ipv4-address	true	false
www-server	72	ipv4-address	true	false
finger-server	73	ipv4-address	true	false
irc-server	74	ipv4-address	true	false
streettalk-server	75	ipv4-address	true	false
streettalk-directory-assistance-server	76	ipv4-address	true	false
user-class	77	binary	false	false
slp-directory-agent	78	record (boolean, ipv4-address)	true	false
slp-service-scope	79	record (boolean, string)	false	false
nds-servers	85	ipv4-address	true	false
nds-tree-name	86	string	false	false
nds-context	87	string	false	false
bcms-controller-names	88	fqdn	true	false
bcms-controller-address	89	ipv4-address	true	false
client-system	93	uint16	true	false
client-ndi	94	record (uint8, uint8, uint8)	false	false
uuid-guid	97	record (uint8, binary)	false	false
uap-servers	98	string	false	false
geoconf-civic	99	binary	false	false
pcode	100	string	false	false
tcode	101	string	false	false
v6-only-preferred	108	uint32	false	false
netinfo-server-address	112	ipv4-address	true	false
netinfo-server-tag	113	string	false	false
v4-captive-portal	114	string	false	false
auto-config	116	uint8	false	false
name-service-search	117	uint16	true	false
domain-search	119	fqdn	true	false
classless-static-route	121	internal	false	false
cablelabs-client-conf	122	empty	false	false
vivco-suboptions	124	record (uint32, binary)	false	false
vivso-suboptions	125	uint32	false	false
pana-agent	136	ipv4-address	true	false
v4-lost	137	fqdn	false	false
capwap-ac-v4	138	ipv4-address	true	false
sip-ua-cs-domains	141	fqdn	true	false
v4-sztp-redirect	143	tuple	true	false
rdnss-selection	146	record (uint8, ipv4-address, ipv4-address, fqdn)	true	false
v4-portparams	159	record (uint8, psid)	false	false
v4-dnr	162	record (uint16, uint16, uint8, fqdn, binary)	false	false
option-6rd	212	record (uint8, uint8, ipv6-address, ipv4-address)	true	false
v4-access-domain	213	fqdn	false	false

Note

The `default-url` option was replaced with `v4-captive-portal` in Kea 2.1.2, as introduced by [RFC 8910](#). The new option has exactly the same format as the old one. The general perception is that `default-url` was seldom used. Migrating users should replace `default-url` with `v4-captive-portal` in their configurations.

Kea also supports other options than those listed above; the following options are returned by the Kea engine itself and in general should not be configured manually.

Table 2: List of standard DHCPv4 options managed by Kea on its own and not directly configurable by an administrator

Name	Code	Type	Description
subnet-mask	1	ipv4-address	calculated automatically, based on subnet definition.
host-name	12	string	sent by client, generally governed by the DNS configuration.
dhcp-requested-address	50	ipv4-address	may be sent by the client and the server should not set it.
dhcp-lease-time	51	uint32	set automatically based on the <code>valid-lifetime</code> parameter.
dhcp-message-type	53	string	sent by clients and servers. Set by the Kea engine depending on the situation and should never be configured explicitly.
dhcp-parameter-request-list	55	uint8 array	sent by clients and should never be sent by the server.
dhcp-renewal-time	58	uint32	governed by <code>renew-timer</code> parameter.
dhcp-rebinding-time	59	uint32	governed by <code>rebind-timer</code> parameter.
dhcp-client-identifier	61	binary	sent by client, echoed back with the value sent by the client.
fqdn	81	record (uint8, uint8, uint8, fqdn)	part of the DDNS and D2 configuration.
dhcp-agent-options	82	empty	sent by the relay agent. This is an empty container option; see RAI option detail later in this section.
authenticate	90	binary	sent by client, Kea does not yet validate it.
client-last-transaction-time	91	uint32	sent by client, server does not set it.
associated-ip	92	ipv4-address array	sent by client, server responds with list of addresses.
subnet-selection	118	ipv4-address	if present in client's messages, will be used in the subnet selection process.
status-code	151	record (uint8, string)	set by server in bulk leasequery responses.
base-time	152	uint32	set by server in bulk leasequery responses.
start-time-of-state	153	uint32	set by server in bulk leasequery responses.
query-start-time	154	uint32	set by requestor during bulk leasequery.
query-end-time	155	uint32	set by requestor during bulk leasequery.
dhcp-state	156	uint8	set by server in bulk leasequery responses.
data-source	157	uint8	set by server in bulk leasequery responses.

The following table lists all option types used in the previous two tables with a description of what values are accepted for them.

Table 3: List of standard DHCP option types

Nam	Meaning
bi-nary	An arbitrary string of bytes, specified as a set of hexadecimal digits.
boole	A boolean value with allowed values true or false.
empt	No value; data is carried in sub-options.
fqdn	Fully qualified domain name (e.g. www.example.com).
ipv4-addr	IPv4 address in the usual dotted-decimal notation (e.g. 192.0.2.1).
ipv6-addr	IPv6 address in the usual colon notation (e.g. 2001:db8::1).
ipv6-prefix	IPv6 prefix and prefix length specified using CIDR notation, e.g. 2001:db8:1::/64. This data type is used to represent an 8-bit field conveying a prefix length and the variable length prefix value.
psid	PSID and PSID length separated by a slash, e.g. 3/4 specifies PSID=3 and PSID length=4. In the wire format it is represented by an 8-bit field carrying PSID length (in this case equal to 4) and the 16-bits-long PSID value field (in this case equal to "0011000000000000b" using binary notation). Allowed values for a PSID length are 0 to 16. See RFC 7597 for details about the PSID wire representation.
recor	Structured data that may be comprised of any types (except "record" and "empty"). The array flag applies to the last field only.
string	Any text. Please note that Kea silently discards any terminating/trailing nulls from the end of "string" options when unpacking received packets. This is in keeping with RFC 2132, Section 2 .
tu-ple	A length field encoded as an 8-bit or 16-bit unsigned integer followed by a string of this length. Typically, for DHCPv4, the length is 8-bit, and for DHCPv6, it is 16-bit. However, there are exceptions to that rule. E.g. for the DHCPv4 SZTP Redirect Option, bootstrap-server-list is encoded as a list of tuples where the URI-length in each tuple is a 16-bit unsigned integer.
uint8	An 8-bit unsigned integer with allowed values 0 to 255.
uint1	A 16-bit unsigned integer with allowed values 0 to 65535.
uint3	A 32-bit unsigned integer with allowed values 0 to 4294967295.
int8	An 8-bit signed integer with allowed values -128 to 127.
int16	A 16-bit signed integer with allowed values -32768 to 32767.
int32	A 32-bit signed integer with allowed values -2147483648 to 2147483647.

Kea also supports the Relay Agent Information (RAI, defined in [RFC 3046](#)) option, sometimes referred to as the relay option, agent option, or simply option 82. The option itself is just a container and does not convey any information on its own. The following table contains a list of RAI sub-options that Kea can understand. The RAI and its sub-options are inserted by the relay agent and received by Kea; there is no need for Kea to be configured with those options. Kea's classification and flexible ID features in host reservations can be used to process those and other options not listed in the table below.

Table 4: List of RAI sub-options that Kea can understand

Name	Code	Comment
circuit-id	1	Used when host-reservation-identifiers is set to <i>circuit-id</i> .
remote-id	2	Can be used with flex-id to identify hosts.
link-selection	5	If present, used to select the appropriate subnet.
subscriber-id	6	Can be used with flex-id to identify hosts.
server-id-override	11	If sent by the relay, Kea accepts it as the <i>server-id</i> .
relay-id	12	Identifies the relay
relay-port	19	If sent by the relay, Kea sends back its responses to this port.

All other RAI sub-options (including those not listed here) can be used in client classification to classify incoming pack-

ets to specific classes, and/or by *libdhcp_flex_id.so* to construct a unique device identifier. For more information about expressions used in client classification and flexible identifiers, see *Client Classification*. The RAI sub-options can be referenced using `relay4[option-code].hex`; for example, to classify packets based on the `remote-id` (sub-option code 2), one would use `relay4[2].hex`. An example client class that includes all packets with a specific `remote-id` value would look as follows:

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "remote-id-1020304",
      "test": "relay4[2].hex == 0x01020304",
      ...
    }
  ],
  ...
}
```

Classes may be used to segregate traffic into a relatively small number of groups, which then can be used to select specific subnets, pools and extra options, and more. If per-host behavior is necessary, using host reservations with flexible identifiers is strongly recommended.

9.2.11 CableLabs Client Conf Suboptions

CableLabs client conf option ("cablelabs-client-conf" code 122) is a container of suboptions in the "cablelabs-client-conf" space listed in the table below.

Table 5: List of CableLabs Client Conf sub-options that Kea can understand

Name	Code	Format
tsp-primary-server	1	IPv4 address
tsp-secondary-server	2	IPv4 address
tsp-as-parameters	4	record of 3 uint32's
tsp-ap-parameters	5	record of 3 uint32's
tsp-realm	6	Fully Qualified Domain Name
tsp-use-tgt	7	boolean
tsp-provisioning-timer	8	uint8
tsp-sct	9	uint16
kdc-server	10	array of IPv4 addresses

These suboptions are defined in [RFC 3495](#) including its errata which clarifies the realm format, [RFC 3594](#) and [RFC 3634](#).

Note

The suboption 3 carries the TSP provisioning server address as an either IPv4 address or a FQDN. This can't be defined in Kea so no standard suboption is defined for the code 3 leaving the choice to configure its content as a binary value, or if it is used only as an IPv4 address or a FQDN to define it as a record of an `uint8` (set to 0) and `ipv4-address`, or a record of an `uint8` (set to 1) and a `fqdn`, allowing the use of CSV formatted data.

9.2.12 Custom DHCPv4 Options

Kea supports custom (non-standard) DHCPv4 options. Let's say that we want to define a new DHCPv4 option called `foo`, which will have code 222 and will convey a single, unsigned, 32-bit integer value. Such an option can be defined by putting the following entry in the configuration file:

```
"Dhcp4": {
  "option-def": [
    {
      "name": "foo",
      "code": 222,
      "type": "uint32",
      "array": false,
      "record-types": "",
      "space": "dhcp4",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

The `false` value of the `array` parameter determines that the option does NOT comprise an array of `uint32` values but is, instead, a single value. Two other parameters have been left blank: `record-types` and `encapsulate`. The former specifies the comma-separated list of option data fields, if the option comprises a record of data fields. The `record-types` value should be non-empty if `type` is set to "record"; otherwise it must be left blank. The latter parameter specifies the name of the option space being encapsulated by the particular option. If the particular option does not encapsulate any option space, the parameter should be left blank. Note that the `option-def` configuration statement only defines the format of an option and does not set its value(s).

The `name`, `code`, and `type` parameters are required; all others are optional. The `array` parameter's default value is `false`. The `record-types` and `encapsulate` parameters' default values are blank (`""`). The default `space` is `dhcp4`.

Once the new option format is defined, its value is set in the same way as for a standard option. For example, the following commands set a global value that applies to all subnets.

```
"Dhcp4": {
  "option-data": [
    {
      "name": "foo",
      "code": 222,
      "space": "dhcp4",
      "csv-format": true,
      "data": "12345"
    },
    ...
  ],
  ...
}
```

New options can take more complex forms than the simple use of primitives (`uint8`, `string`, `ipv4-address`, etc.); it is possible to define an option comprising a number of existing primitives.

For example, say we want to define a new option that will consist of an IPv4 address, followed by an unsigned 16-bit integer, followed by a boolean value, followed by a text string. Such an option could be defined in the following way:

```
"Dhcp4": {
  "option-def": [
    {
      "name": "bar",
      "code": 223,
      "space": "dhcp4",
      "type": "record",
      "array": false,
      "record-types": "ipv4-address, uint16, boolean, string",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

The type parameter is set to "record" to indicate that the option contains multiple values of different types. These types are given as a comma-separated list in the record-types field and should be ones from those listed in *List of standard DHCP option types*.

The option's values are set in an option-data statement as follows:

```
"Dhcp4": {
  "option-data": [
    {
      "name": "bar",
      "space": "dhcp4",
      "code": 223,
      "csv-format": true,
      "data": "192.0.2.100, 123, true, Hello World"
    }
  ],
  ...
}
```

The csv-format parameter is set to true to indicate that the data field comprises a comma-separated list of values. The values in data must correspond to the types set in the record-types field of the option definition.

When array is set to true and type is set to "record", the last field is an array, i.e. it can contain more than one value, as in:

```
"Dhcp4": {
  "option-def": [
    {
      "name": "bar",
      "code": 223,
      "space": "dhcp4",
      "type": "record",
      "array": true,
      "record-types": "ipv4-address, uint16",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

(continues on next page)

(continued from previous page)

```

    ...
}

```

The new option content is one IPv4 address followed by one or more 16-bit unsigned integers.

Note

In general, boolean values are specified as `true` or `false`, without quotes. Some specific boolean parameters may also accept `"true"`, `"false"`, `0`, `1`, `"0"`, and `"1"`.

Note

Numbers can be specified in decimal or hexadecimal format. The hexadecimal format can be either plain (e.g. `abcd`) or prefixed with `0x` (e.g. `0xabcd`).

9.2.13 DHCPv4 Private Options

Options with a code between 224 and 254 are reserved for private use. They can be defined at the global scope or at the client-class local scope; this allows option definitions to be used depending on context, and option data to be set accordingly. For instance, to configure an old PXEClient vendor:

```

"Dhcp4": {
  "client-classes": [
    {
      "name": "pxeclient",
      "test": "option[vendor-class-identifier].text == 'PXEClient'",
      "option-def": [
        {
          "name": "configfile",
          "code": 209,
          "type": "string"
        }
      ],
      ...
    },
    ...
  ],
  ...
}

```

As the Vendor-Specific Information (VSI) option (code 43) has a vendor-specific format, i.e. can carry either raw binary value or sub-options, this mechanism is also available for this option.

In the following example taken from a real configuration, two vendor classes use option 43 for different and incompatible purposes:

```

"Dhcp4": {
  "option-def": [
    {
      "name": "cookie",

```

(continues on next page)

(continued from previous page)

```

        "code": 1,
        "type": "string",
        "space": "APC"
    },
    {
        "name": "mtftp-ip",
        "code": 1,
        "type": "ipv4-address",
        "space": "PXE"
    },
    ...
],
"client-classes": [
    {
        "name": "APC",
        "test": "option[vendor-class-identifier].text == 'APC'",
        "option-def": [
            {
                "name": "vendor-encapsulated-options",
                "type": "empty",
                "encapsulate": "APC"
            }
        ],
        "option-data": [
            {
                "name": "cookie",
                "space": "APC",
                "data": "1APC"
            },
            {
                "name": "vendor-encapsulated-options"
            },
            ...
        ],
        ...
    },
    {
        "name": "PXE",
        "test": "option[vendor-class-identifier].text == 'PXE'",
        "option-def": [
            {
                "name": "vendor-encapsulated-options",
                "type": "empty",
                "encapsulate": "PXE"
            }
        ],
        "option-data": [
            {
                "name": "mtftp-ip",
                "space": "PXE",
                "data": "0.0.0.0"
            },
            ...
        ],
        ...
    }
],
...

```

(continues on next page)

(continued from previous page)

```

        {
            "name": "vendor-encapsulated-options"
        },
        ...
    ],
    ...
},
...
],
...
}

```

The definition used to decode a VSI option is:

1. The local definition of a client class the incoming packet belongs to;
2. If none, the global definition;
3. If none, the last-resort definition described in the next section, *DHCPv4 Vendor-Specific Options* (backward-compatible with previous Kea versions).

Note

This last-resort definition for the Vendor-Specific Information option (code 43) is not compatible with a raw binary value. When there are known cases where a raw binary value will be used, a client class must be defined with both a classification expression matching these cases and an option definition for the VSI option with a binary type and no encapsulation.

Note

By default, in the Vendor-Specific Information option (code 43), sub-option code 0 and 255 mean PAD and END respectively, according to [RFC 2132](#). In other words, the sub-option code values of 0 and 255 are reserved. Kea does, however, allow users to define sub-option codes from 0 to 255. If sub-options with codes 0 and/or 255 are defined, bytes with that value are no longer treated as a PAD or an END, but as the sub-option code when parsing a VSI option in an incoming query.

Option 43 input processing (also called unpacking) is deferred so that it happens after classification. This means clients cannot be classified using option 43 sub-options. The definition used to unpack option 43 is determined as follows:

- If defined at the global scope, this definition is used.
- If defined at client class scope and the packet belongs to this class, the client class definition is used.
- If not defined at global scope nor in a client class to which the packet belongs, the built-in last resort definition is used. This definition only says the sub-option space is "vendor-encapsulated-options-space".

The output definition selection is a bit simpler:

- If the packet belongs to a client class which defines the option 43, use this definition.
- If defined at the global scope, use this definition.
- Otherwise, use the built-in last-resort definition.

Since they use a specific/per vendor option space, sub-options are defined at the global scope.

Note

Option definitions in client classes are allowed only for this limited option set (codes 43 and from 224 to 254), and only for DHCPv4.

9.2.14 DHCPv4 Vendor-Specific Options

Currently there are two option spaces defined for *kea-dhcp4*: *dhcp4* (for the top-level DHCPv4 options) and *"vendor-encapsulated-options-space"*, which is empty by default but in which options can be defined. Those options are carried in the Vendor-Specific Information option (code 43). The following examples show how to define an option *foo* with code 1 that comprises an IPv4 address, an unsigned 16-bit integer, and a string. The *foo* option is conveyed in a Vendor-Specific Information option.

The first step is to define the format of the option:

```
"Dhcp4": {
  "option-def": [
    {
      "name": "foo",
      "code": 1,
      "space": "vendor-encapsulated-options-space",
      "type": "record",
      "array": false,
      "record-types": "ipv4-address, uint16, string",
      "encapsulate": ""
    }
  ],
  ...
}
```

Note that the option space is set to *"vendor-encapsulated-options-space"*. Once the option format is defined, the next step is to define actual values for that option:

```
"Dhcp4": {
  "option-data": [
    {
      "name": "foo",
      "space": "vendor-encapsulated-options-space",
      "code": 1,
      "csv-format": true,
      "data": "192.0.2.3, 123, Hello World"
    }
  ],
  ...
}
```

In this example, we also include the Vendor-Specific Information option, which conveys our sub-option *foo*. This is required; otherwise, the option will not be included in messages sent to the client.

```
"Dhcp4": {
  "option-data": [
    {
      "name": "vendor-encapsulated-options"
```

(continues on next page)

(continued from previous page)

```

    }
  ],
  ...
}

```

Alternatively, the option can be specified using its code.

```

"Dhcp4": {
  "option-data": [
    {
      "code": 43
    }
  ],
  ...
}

```

Another popular option that is often somewhat imprecisely called the "vendor option" is option 125. Its proper name is the "vendor-independent vendor-specific information option" or "vivso". The idea behind vivso options is that each vendor has its own unique set of options with their own custom formats. The vendor is identified by a 32-bit unsigned integer called `enterprise-number` or `vendor-id`.

The standard spaces defined in Kea and their options are:

- `vendor-4491`: Cable Television Laboratories, Inc. for DOCSIS3 options:

option code	option name	option description
1	oro	ORO (or Option Request Option), used by clients to request a list of options they are interested in.
2	tftp-servers	a list of IPv4 addresses of TFTP servers to be used by the cable modem

In Kea, each vendor is represented by its own vendor space. Since there are hundreds of vendors and they sometimes use different option definitions for different hardware, it is impossible for Kea to support them all natively. Fortunately, it is easy to define support for new vendor options. As an example, the Genexis home gateway device requires the vivso 125 option to be sent with a sub-option 2 that contains a string with the TFTP server URL. To support such a device, three steps are needed: first, establish option definitions that explain how the option is supposed to be formed; second, define option values; and third, tell Kea when to send those specific options, via client classification.

An example snippet of a configuration could look similar to the following:

```

"Dhcp4": {
  // First, we need to define that the sub-option 2 in vivso option for
  // vendor-id 25167 has a specific format (it's a plain string in this example).
  // After this definition, we can specify values for option tftp.
  "option-def": [
    {
      // We define a short name, so the option can be referenced by name.
      // The option has code 2 and resides within vendor space 25167.
      // Its data is a plain string.
      "name": "tftp",
      "code": 2,
      "space": "vendor-25167",
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

        "type": "string"
    }
],
"client-classes": [
    {
        // We now need to tell Kea how to recognize when to use vendor space 25167.
        // Usually we can use a simple expression, such as checking if the device
        // sent a vivso option with specific vendor-id, e.g. "vendor[4491].exists".
        // Unfortunately, Genexis is a bit unusual in this aspect, because it
        // doesn't send vivso. In this case we need to look into the vendor class
        // (option code 60) and see if there's a specific string that identifies
        // the device. Alternatively, one can make use of the automated `VENDOR_
↪CLASS_`
        // client class and replace "name" and "test" with `"name": "VENDOR_CLASS_
↪HMC1000"`
        // and no test expression.
        "name": "cpe_genexis",
        "test": "substring(option[60].hex,0,7) == 'HMC1000'",

        // Once the device is recognized, we want to send two options:
        // the vivso option with vendor-id set to 25167, and a sub-option 2.
        "option-data": [
            {
                "name": "vivso-suboptions",
                "data": "25167"
            },

            // The sub-option 2 value is defined as any other option. However,
            // we want to send this sub-option 2, even when the client didn't
            // explicitly request it (often there is no way to do that for
            // vendor options). Therefore we use always-send to force Kea
            // to always send this option when 25167 vendor space is involved.
            {
                "name": "tftp",
                "space": "vendor-25167",
                "data": "tftp://192.0.2.1/genexis/HMC1000.v1.3.0-R.img",
                "always-send": true
            }
        ]
    }
]
}

```

By default, Kea sends back only those options that are requested by a client, unless there are protocol rules that tell the DHCP server to always send an option. This approach works nicely in most cases and avoids problems with clients refusing responses with options they do not understand. However, the situation with vendor options is more complex, as they are not requested the same way as other options, are not well-documented in official RFCs, or vary by vendor.

Some vendors (such as DOCSIS, identified by vendor option 4491) have a mechanism to request specific vendor options and Kea is able to honor those (sub-option 1). Unfortunately, for many other vendors, such as Genexis (25167, discussed above), Kea does not have such a mechanism, so it cannot send any sub-options on its own. To solve this issue, we devised the concept of persistent options. Kea can be told to always send options, even if the client did not request

them. This can be achieved by adding `"always-send": true` to the option data entry. Note that in this particular case an option is defined in vendor space 25167. With `always-send` enabled, the option is sent every time there is a need to deal with vendor space 25167.

This is also how `kea-dhcp4` can be configured to send multiple vendor options from different vendors, along with each of their specific vendor IDs. If these options need to be sent by the server regardless of whether the client specified any enterprise number, `"always-send": true` must be configured for the suboptions that will be included in the `vivso-suboptions` option (code 125).

```
"Dhcp4": {
  "option-data": [
    # Typically DHCPv4 clients will send a Parameter Request List option (code 55).
    ↪for
    # vivso-suboptions (code 125), and that is enough for Kea to understand that it
    ↪needs to
    # send the option. These options still need to be defined in the configuration,
    ↪one per
    # each vendor, but they don't need "always-send" enabled in that case. For
    ↪misbehaving
    # clients that do not explicitly request it, one may alternatively set "always-
    ↪send"
    # to true for them as well. This is referring to the following two entries in
    ↪option-data.
    {
      "name": "vivso-suboptions",
      "space": "dhcp4",
      "data": "2234"
    },
    {
      "name": "vivso-suboptions",
      "space": "dhcp4",
      "data": "3561"
    },
    {
      "always-send": true,
      "data": "tagged",
      "name": "tag",
      "space": "vendor-2234"
    },
    {
      "always-send": true,
      "data": "https://example.com:1234/path",
      "name": "url",
      "space": "vendor-3561"
    }
  ],
  "option-def": [
    {
      "code": 22,
      "name": "tag",
      "space": "vendor-2234",
      "type": "string"
    },
    {
```

(continues on next page)

(continued from previous page)

```

        "code": 11,
        "name": "url",
        "space": "vendor-3561",
        "type": "string"
    }
]
}

```

Another possibility is to redefine the option; see *DHCPv4 Private Options*.

Kea comes with several example configuration files. Some of them showcase how to configure options 60 and 43. See `doc/examples/kea4/vendor-specific.json` and `doc/examples/kea4/vivso.json` in the Kea sources.

Note

kea-dhcp4 is able to recognize multiple Vendor Class Identifier options (code 60) with different vendor IDs in the client requests, and to send multiple vivso options (code 125) in the responses, one for each vendor.

kea-dhcp4 honors DOCSIS sub-option 1 (ORO) and adds only requested options if this sub-option is present in the client request.

Currently only one vendor is supported for the vivco-suboptions (code 124) option. Specifying multiple enterprise numbers within a single option instance or multiple options with different enterprise numbers is not supported.

9.2.15 Nested DHCPv4 Options (Custom Option Spaces)

It is sometimes useful to define a completely new option space, such as when a user creates a new option in the standard option space (dhcp4) and wants this option to convey sub-options. Since they are in a separate space, sub-option codes have a separate numbering scheme and may overlap with the codes of standard options.

Note that the creation of a new option space is not required when defining sub-options for a standard option, because one is created by default if the standard option is meant to convey any sub-options (see *DHCPv4 Vendor-Specific Options*).

If we want a DHCPv4 option called `container` with code 222, that conveys two sub-options with codes 1 and 2, we first need to define the new sub-options:

```

"Dhcp4": {
  "option-def": [
    {
      "name": "subopt1",
      "code": 1,
      "space": "isc",
      "type": "ipv4-address",
      "record-types": "",
      "array": false,
      "encapsulate": ""
    },
    {
      "name": "subopt2",
      "code": 2,
      "space": "isc",
      "type": "string",
      "record-types": "",
      "array": false,

```

(continues on next page)

(continued from previous page)

```

        "encapsulate": ""
    }
],
...
}

```

Note that we have defined the options to belong to a new option space (in this case, "isc").

The next step is to define a regular DHCPv4 option with the desired code and specify that it should include options from the new option space:

```

"Dhcp4": {
    "option-def": [
        {
            "name": "container",
            "code": 222,
            "space": "dhcp4",
            "type": "empty",
            "array": false,
            "record-types": "",
            "encapsulate": "isc"
        },
        ...
    ],
    ...
}

```

The name of the option space in which the sub-options are defined is set in the `encapsulate` field. The `type` field is set to "empty", to indicate that this option does not carry any data other than sub-options.

Finally, we can set values for the new options:

```

{
    "Dhcp4": {
        "option-data": [
            {
                "name": "subopt1",
                "code": 1,
                "space": "isc",
                "data": "192.0.2.3"
            },
            {
                "name": "subopt2",
                "code": 2,
                "space": "isc",
                "data": "Hello world"
            },
            {
                "name": "container",
                "code": 222,
                "space": "dhcp4"
            }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

It is possible to create an option which carries some data in addition to the sub-options defined in the encapsulated option space. For example, if the `container` option from the previous example were required to carry a `uint16` value as well as the sub-options, the `type` value would have to be set to `"uint16"` in the option definition. (Such an option would then have the following data structure: DHCP header, `uint16` value, sub-options.) The value specified with the `data` parameter — which should be a valid integer enclosed in quotes, e.g. `"123"` — would then be assigned to the `uint16` field in the `container` option.

9.2.16 Unspecified Parameters for DHCPv4 Option Configuration

In many cases it is not required to specify all parameters for an option configuration, and the default values can be used. However, it is important to understand the implications of not specifying some of them, as it may result in configuration errors. The list below explains the behavior of the server when a particular parameter is not explicitly specified:

- **name** - the server requires either an option name or an option code to identify an option. If this parameter is unspecified, the option code must be specified.
- **code** - the server requires either an option name or an option code to identify an option; this parameter may be left unspecified if the `name` parameter is specified. However, this also requires that the particular option have a definition (either as a standard option or an administrator-created definition for the option using an `option-def` structure), as the option definition associates an option with a particular name. It is possible to configure an option for which there is no definition (unspecified option format). Configuration of such options requires the use of the option code.
- **space** - if the option space is unspecified it defaults to `dhcp4`, which is an option space holding standard DHCPv4 options.
- **data** - if the option data is unspecified it defaults to an empty value. The empty value is mostly used for the options which have no payload (boolean options), but it is legal to specify empty values for some options which carry variable-length data and for which the specification allows a length of 0. For such options, the `data` parameter may be omitted in the configuration.
- **csv-format** - if this value is not specified, the server assumes that the option data is specified as a list of comma-separated values to be assigned to individual fields of the DHCP option.

9.2.17 Support for Long Options

The `kea-dhcp4` server partially supports long options (RFC 3396). Since Kea 2.1.6, the server accepts configuring long options and sub-options (longer than 255 bytes). The options and sub-options are stored internally in their unwrapped form and they can be processed as usual using the parser language. On send, the server splits long options and sub-options into multiple options and sub-options, using the respective option code.

```
{
"option-def": [
  {
    "array": false,
    "code": 240,
    "encapsulate": "",
    "name": "my-option",
    "space": "dhcp4",
    "type": "string"
  }
],
```

(continues on next page)

(continued from previous page)

```

"subnet4": [
  {
    "id": 1,
    "subnet": "192.0.2.0/24",
    "reservations": [
      {
        "hw-address": "aa:bb:cc:dd:ee:ff",
        "option-data": [
          {
            "always-send": false,
            "code": 240,
            "name": "my-option",
            "csv-format": true,
            "data": "data \
-00010203040506070809-00010203040506070809-
↪00010203040506070809-00010203040506070809 \
-00010203040506070809-00010203040506070809-
↪00010203040506070809-00010203040506070809 \
-00010203040506070809-00010203040506070809-
↪00010203040506070809-00010203040506070809 \
-data",
            "space": "dhcp4"
          }
        ]
      }
    ]
  }
],
...
}

```

Note

In the example above, the data has been wrapped into several lines for clarity, but Kea does not support wrapping in the configuration file.

This example illustrates configuring a custom long option (exceeding 255 octets) in a reservation. When sending a response, the server splits this option into two options, each with the code 240.

Note

Currently the server does not support storing long options in databases, either via host reservations or the configuration backend.

The server is also able to receive packets with split options (options using the same option code) and to fuse the data chunks into one option. This is also supported for sub-options if each sub-option data chunk also contains the sub-option code and sub-option length.

9.2.18 Support for IPv6-Only Preferred Option

The v6-only-preferred (code 108) option is handled in a specific way described in [RFC 8925](#) by *kea-dhcp4* when it is configured in a subnet or a shared network: when the client requests the option (i.e. puts the 108 code in the DHCP parameter request list option) and the subnet or shared network is selected the 0.0.0.0 address is offered and the option returned in the response.

9.2.19 Stateless Configuration of DHCPv4 Clients

The DHCPv4 server supports stateless client configuration, whereby the client has an IP address configured (e.g. using manual configuration) and only contacts the server to obtain other configuration parameters, such as addresses of DNS servers. To obtain the stateless configuration parameters, the client sends the DHCPINFORM message to the server with the `ciaddr` set to the address that the client is currently using. The server unicasts the DHCPACK message to the client that includes the stateless configuration ("`yiaddr`" not set).

The server responds to the DHCPINFORM when the client is associated with a subnet defined in the server's configuration. An example subnet configuration looks like this:

```
"Dhcp4": {
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "option-data": [
        {
          "name": "domain-name-servers",
          "code": 6,
          "data": "192.0.2.200,192.0.2.201",
          "csv-format": true,
          "space": "dhcp4"
        }
      ]
    }
  ]
}
```

This subnet specifies the single option which will be included in the DHCPACK message to the client in response to DHCPINFORM. The subnet definition does not require the address pool configuration if it will be used solely for stateless configuration.

This server will associate the subnet with the client if one of the following conditions is met:

- The DHCPINFORM is relayed and the `giaddr` matches the configured subnet.
- The DHCPINFORM is unicast from the client and the `ciaddr` matches the configured subnet.
- The DHCPINFORM is unicast from the client and the `ciaddr` is not set, but the source address of the IP packet matches the configured subnet.
- The DHCPINFORM is not relayed and the IP address on the interface on which the message is received matches the configured subnet.

9.2.20 Client Classification in DHCPv4

The DHCPv4 server includes support for client classification. For a deeper discussion of the classification process, see *Client Classification*.

In certain cases it is useful to configure the server to differentiate between DHCP client types and treat them accordingly. Client classification can be used to modify the behavior of almost any part of DHCP message processing. Kea

currently offers client classification via private options and option 43 deferred unpacking; subnet selection; pool selection; assignment of different options; and, for cable modems, specific options for use with the TFTP server address and the boot file field.

Kea can be instructed to limit access to given subnets based on class information. This is particularly useful for cases where two types of devices share the same link and are expected to be served from two different subnets. The primary use case for such a scenario is cable networks, where there are two classes of devices: the cable modem itself, which should be handed a lease from subnet A; and all other devices behind the modem, which should get leases from subnet B. That segregation is essential to prevent overly curious end-users from playing with their cable modems. For details on how to set up class restrictions on subnets, see *Configuring Subnets With Class Information*.

When subnets belong to a shared network, the classification applies to subnet selection but not to pools; that is, a pool in a subnet limited to a particular class can still be used by clients which do not belong to the class, if the pool they are expected to use is exhausted. The limit on access based on class information is also available at the pool level within a subnet: see *Configuring Pools With Class Information*. This is useful when segregating clients belonging to the same subnet into different address ranges.

In a similar way, a pool can be constrained to serve only known clients, i.e. clients which have a reservation, using the built-in KNOWN or UNKNOWN classes. Addresses can be assigned to registered clients without giving a different address per reservation: for instance, when there are not enough available addresses. The determination whether there is a reservation for a given client is made after a subnet is selected, so it is not possible to use KNOWN/UNKNOWN classes to select a shared network or a subnet.

The process of classification is conducted in five steps. The first step is to assess an incoming packet and assign it to zero or more classes. The second step is to choose a subnet, possibly based on the class information. When the incoming packet is in the special class DROP, it is dropped and a debug message logged. The next step is to evaluate class expressions depending on the built-in KNOWN/UNKNOWN classes after host reservation lookup, using them for pool selection and assigning classes from host reservations. The list of required classes is then built and each class of the list has its expression evaluated; when it returns `true`, the packet is added as a member of the class. The last step is to assign options, again possibly based on the class information. More complete and detailed information is available in *Client Classification*.

There are two main methods of classification. The first is automatic and relies on examining the values in the vendor class options or the existence of a host reservation. Information from these options is extracted, and a class name is constructed from it and added to the class list for the packet. The second method specifies an expression that is evaluated for each packet. If the result is `true`, the packet is a member of the class.

Note

The new `early-global-reservations-lookup` global parameter flag enables a lookup for global reservations before the subnet selection phase. This lookup is similar to the general lookup described above with two differences:

- the lookup is limited to global host reservations
- the UNKNOWN class is never set

Note

Care should be taken with client classification, as it is easy for clients that do not meet class criteria to be denied all service.

9.2.20.1 Setting Fixed Fields in Classification

It is possible to specify that clients belonging to a particular class should receive packets with specific values in certain fixed fields. In particular, three fixed fields are supported: `next-server` (conveys an IPv4 address, which is set in the `siaddr` field), `server-hostname` (conveys a server hostname, can be up to 64 bytes long, and is sent in the `sname` field) and `boot-file-name` (conveys the configuration file, can be up to 128 bytes long, and is sent using the `file` field).

Obviously, there are many ways to assign clients to specific classes, but for PXE clients the client architecture type option (code 93) seems to be particularly suited to make the distinction. The following example checks whether the client identifies itself as a PXE device with architecture EFI x86-64, and sets several fields if it does. See [Section 2.1 of RFC 4578](#)) or the client documentation for specific values.

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "ipxe_efi_x64",
      "test": "option[93].hex == 0x0009",
      "next-server": "192.0.2.254",
      "server-hostname": "hal9000",
      "boot-file-name": "/dev/null"
    },
    ...
  ],
  ...
}
```

If an incoming packet is matched to multiple classes, then the value used for each field will come from the first class that specifies the field, in the order the classes are assigned to the packet.

Note

The classes are ordered as specified in the configuration.

9.2.20.2 Using Vendor Class Information in Classification

The server checks whether an incoming packet includes the vendor class identifier option (60). If it does, the content of that option is prepended with `VENDOR_CLASS_`, and it is interpreted as a class. For example, modern cable modems send this option with value `docsis3.0`, so the packet belongs to the class `VENDOR_CLASS_docsis3.0`.

Note

Certain special actions for clients in `VENDOR_CLASS_docsis3.0` can be achieved by defining `VENDOR_CLASS_docsis3.0` and setting its `next-server` and `boot-file-name` values appropriately.

This example shows a configuration using an automatically generated `VENDOR_CLASS_` class. The administrator of the network has decided that addresses from the range 192.0.2.10 to 192.0.2.20 are going to be managed by the DhcP4 server and only clients belonging to the DOCSIS 3.0 client class are allowed to use that pool.

```
"Dhcp4": {
  "subnet4": [
    {
      "id": 1,
```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.2.0/24",
        "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
        "client-classes": [ "VENDOR_CLASS_docsis3.0" ]
    }
],
...
}

```

9.2.20.3 Defining and Using Custom Classes

The following example shows how to configure a class using an expression and a subnet using that class. This configuration defines the class named `Client_foo`. It is comprised of all clients whose client IDs (option 61) start with the string `foo`. Members of this class will be given addresses from 192.0.2.10 to 192.0.2.20 and the addresses of their DNS servers set to 192.0.2.1 and 192.0.2.2.

```

"Dhcp4": {
    "client-classes": [
        {
            "name": "Client_foo",
            "test": "substring(option[61].hex,0,3) == 'foo'",
            "option-data": [
                {
                    "name": "domain-name-servers",
                    "code": 6,
                    "space": "dhcp4",
                    "csv-format": true,
                    "data": "192.0.2.1, 192.0.2.2"
                }
            ]
        }
    ],
    ...
],
"subnet4": [
    {
        "id": 1,
        "subnet": "192.0.2.0/24",
        "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
        "client-classes": [ "Client_foo" ]
    },
    ...
],
...
}

```

9.2.20.4 Additional Classification

In some cases it is useful to limit the scope of a class to a pool, subnet, or shared network. There are two parameters which are used to limit the scope of the class by instructing the server to evaluate test expressions when required.

The `evaluate-additional-classes`, which takes a list of class names and is valid in pool, subnet, and shared network scope. Classes in these lists are marked as additional and evaluated after selection of this specific pool/subnet/shared network and before output-option processing.

The second one is the `per-class only-in-additional-list` flag, which is `false` by default. When it is set to `true`,

the test expression of the class is not evaluated at the reception of the incoming packet but later, and only if the class is present in an `evaluate-additional-classes` list.

In this example, a class is assigned to the incoming packet when the specified subnet is used:

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "Client_foo",
      "test": "member('ALL')",
      "only-in-additional-list": true
    },
    ...
  ],
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
      "evaluate-additional-classes": [ "Client_foo" ],
      ...
    },
    ...
  ],
  ...
}
```

Additional evaluation can be used to express complex dependencies like subnet membership. It can also be used to reverse the precedence; if `option-data` is set in a subnet, it takes precedence over `option-data` in a class. If `option-data` is moved to a required class and required in the subnet, a class evaluated earlier may take precedence.

Additional evaluation is also available at shared network and pool levels. The order in which additional classes are considered is: pool, subnet, and shared network, i.e. in the same order from the way in which `option-data` is processed.

Since Kea version 2.7.4 additional classes configured without a test expression are unconditionally added, i.e. they are considered to always be evaluated to `true`.

Note

Because additional evaluation occurs after lease assignment, parameters that would otherwise impact lease life times (e.g. `valid-lifetime`, `offer-lifetime`) will have no effect when specified in a class that also sets `only-in-additional-list` `true`.

Note

As of Kea version 2.7.4, `only-if-required` and `require-client-classes` have been renamed to `only-in-additional-list` and `evaluate-additional-classes` respectively. The original names will still be accepted as input to allow users to migrate but will eventually be rejected.

Note

Vendor-Identifying Vendor Options are a special case: for all other options an option is identified by its code point, but `vivco-suboptions` (124) and `vivso-suboptions` (125) are identified by the pair of code point and vendor

identifier. This has no visible effect for `vivso-suboptions`, whose value is the vendor identifier, but it is different for `vivco-suboptions`, where the value is a record with the vendor identifier and a binary value. For instance, in:

```
"Dhcp4": {
  "option-data": [
    {
      "name": "vivco-suboptions",
      "always-send": true,
      "data": "1234, 03666f6f"
    },
    {
      "name": "vivco-suboptions",
      "always-send": true,
      "data": "5678, 03626172"
    },
    ...
  ],
  ...
}
```

The first `option-data` entry does not hide the second one, because the vendor identifiers (1234 and 5678) are different: the responses will carry two instances of the `vivco-suboptions` option, each for a different vendor.

9.2.21 DDNS for DHCPv4

As mentioned earlier, *kea-dhcp4* can be configured to generate requests to the DHCP-DDNS server, *kea-dhcp-ddns*, (referred to herein as "D2") to update DNS entries. These requests are known as NameChangeRequests or NCRs. Each NCR contains the following information:

1. Whether it is a request to add (update) or remove DNS entries.
2. Whether the change requests forward DNS updates (A records), reverse DNS updates (PTR records), or both.
3. The Fully Qualified Domain Name (FQDN), lease address, and DHCID (information identifying the client associated with the FQDN).

DDNS-related parameters are split into two groups:

1. Connectivity Parameters

These are parameters which specify where and how *kea-dhcp4* connects to and communicates with D2. These parameters can only be specified within the top-level `dhcp-ddns` section in the *kea-dhcp4* configuration. The connectivity parameters are listed below:

- `enable-updates`
- `server-ip`
- `server-port`
- `sender-ip`
- `sender-port`
- `max-queue-size`
- `ncr-protocol`
- `ncr-format`

2. Behavioral Parameters

These parameters influence behavior such as how client host names and FQDN options are handled. They have been moved out of the `dhcp-ddns` section so that they may be specified at the global, shared-network, subnet, and/or pool levels. Furthermore, they are inherited downward from global to shared-network to subnet to pool. In other words, if a parameter is not specified at a given level, the value for that level comes from the level above it. The behavioral parameters are as follows:

- `ddns-send-updates`
- `ddns-override-no-update`
- `ddns-override-client-update`
- `ddns-replace-client-name`
- `ddns-generated-prefix`
- `ddns-qualifying-suffix`
- `ddns-update-on-renew`
- `ddns-conflict-resolution-mode`
- `ddns-ttl`
- `ddns-ttl-percent`
- `ddns-ttl-min`
- `ddns-ttl-max`
- `hostname-char-set`
- `hostname-char-replacement`

Note

Kea 2.7.6 added support for these parameters to the pool level.

Note

Behavioral parameters that affect the FQDN are in effect even if both `enable-updates` and `ddns-send-updates` are `false`, to support environments in which clients are responsible for their own DNS updates. This applies to `ddns-replace-client-name`, `ddns-generated-prefix`, `ddns-qualifying-suffix`, `hostname-char-set`, and `hostname-char-replacement`.

The default configuration and values would appear as follows:

```
"Dhcp4": {
  "dhcp-ddns": {
    // Connectivity parameters
    "enable-updates": false,
    "server-ip": "127.0.0.1",
    "server-port": 53001,
    "sender-ip": "",
    "sender-port": 0,
    "max-queue-size": 1024,
```

(continues on next page)

(continued from previous page)

```

        "ncr-protocol": "UDP",
        "ncr-format": "JSON"
    },

    // Behavioral parameters (global)
    "ddns-send-updates": true,
    "ddns-override-no-update": false,
    "ddns-override-client-update": false,
    "ddns-replace-client-name": "never",
    "ddns-generated-prefix": "myhost",
    "ddns-qualifying-suffix": "",
    "ddns-update-on-renew": false,
    "ddns-conflict-resolution-mode": "check-with-dhcid",
    "hostname-char-set": "",
    "hostname-char-replacement": "",
    ...
}

```

There are two parameters which determine whether *kea-dhcp4* can generate DDNS requests to D2: the existing `dhcp-ddns:enable-updates` parameter, which now only controls whether *kea-dhcp4* connects to D2; and the new behavioral parameter, `ddns-send-updates`, which determines whether DDNS updates are enabled at a given level (i.e. global, shared-network, or subnet). The following table shows how the two parameters function together:

Table 6: Enabling and disabling DDNS updates

dhcp-ddns: enable-updates	Global ddns-send-updates	Outcome
false (default)	false	no updates at any scope
false	true (default)	no updates at any scope
true	false	updates only at scopes with a local value of true for <code>ddns-enable-updates</code>
true	true	updates at all scopes except those with a local value of false for <code>ddns-enable-updates</code>

Kea 1.9.1 added two new parameters; the first is `ddns-update-on-renew`. Normally, when leases are renewed, the server only updates DNS if the DNS information for the lease (e.g. FQDN, DNS update direction flags) has changed. Setting `ddns-update-on-renew` to `true` instructs the server to always update the DNS information when a lease is renewed, even if its DNS information has not changed. This allows Kea to "self-heal" if it was previously unable to add DNS entries or they were somehow lost by the DNS server.

Note

Setting `ddns-update-on-renew` to `true` may impact performance, especially for servers with numerous clients that renew often.

The second parameter added in Kea 1.9.1 is `ddns-use-conflict-resolution`. This boolean parameter was passed through to D2 and enabled or disabled conflict resolution as described in [RFC 4703](#). Beginning with Kea 2.5.0, it is deprecated and replaced by `ddns-conflict-resolution-mode`, which offers four modes of conflict resolution-related behavior:

- **check-with-dhcid** - This mode, the default, instructs D2 to carry out RFC 4703-compliant conflict resolution. Existing DNS entries may only be overwritten if they have a DHCID record and it matches the client's DHCID.

This is equivalent to `ddns-use-conflict-resolution: true`;

- `no-check-with-dhcid` - Existing DNS entries may be overwritten by any client, whether those entries include a DHCID record or not. The new entries will include a DHCID record for the client to whom they belong. This is equivalent to `ddns-use-conflict-resolution: false`;
- `check-exists-with-dhcid` - Existing DNS entries may only be overwritten if they have a DHCID record. The DHCID record need not match the client's DHCID. This mode provides a way to protect static DNS entries (those that do not have a DHCID record) while allowing dynamic entries (those that do have a DHCID record) to be overwritten by any client. This behavior was not supported prior to Kea 2.4.0.
- `no-check-without-dhcid` - Existing DNS entries may be overwritten by any client; new entries will not include DHCID records. This behavior was not supported prior to Kea 2.4.0.

Note

For backward compatibility, `ddns-use-conflict-resolution` is still accepted in JSON configuration. The server replaces the value internally with `ddns-conflict-resolution-mode` and an appropriate value: `check-with-dhcid` for `true` and `no-check-with-dhcid` for `false`.

Note

Setting `ddns-conflict-resolution-mode` to any value other than `check-with-dhcid` disables the overwrite safeguards that the rules of conflict resolution (from [RFC 4703](#)) are intended to prevent. This means that existing entries for an FQDN or an IP address made for Client-A can be deleted or replaced by entries for Client-B. Furthermore, there are two scenarios by which entries for multiple clients for the same key (e.g. FQDN or IP) can be created.

1. Client-B uses the same FQDN as Client-A but a different IP address. In this case, the forward DNS entries (A and DHCID RRs) for Client-A will be deleted as they match the FQDN, and new entries for Client-B will be added. The reverse DNS entries (PTR and DHCID RRs) for Client-A, however, will not be deleted as they belong to a different IP address, while new entries for Client-B will still be added.
2. Client-B uses the same IP address as Client-A but a different FQDN. In this case, the reverse DNS entries (PTR and DHCID RRs) for Client-A will be deleted as they match the IP address, and new entries for Client-B will be added. The forward DNS entries (A and DHCID RRs) for Client-A, however, will not be deleted, as they belong to a different FQDN, while new entries for Client-B will still be added.

Disabling conflict resolution should be done only after careful review of specific use cases. The best way to avoid unwanted DNS entries is to always ensure that lease changes are processed through Kea, whether they are released, expire, or are deleted via the `lease4-del` command, prior to reassigning either FQDNs or IP addresses. Doing so causes `kea-dhcp4` to generate DNS removal requests to D2.

The DNS entries Kea creates contain a value for TTL (time to live). By default, the `kea-dhcp4` server calculates that value based on [RFC 4702, Section 5](#), which suggests that the TTL value be 1/3 of the lease's lifetime, with a minimum value of 10 minutes. There are four optional parameters which may be used to influence the TTL calculation:

- `ddns-ttl` - If specified and greater than 0 use it as the value of TTL unconditionally. It is specified in seconds. By default it is unspecified.
- `ddns-ttl-percent` - If specified and greater than zero use it as the percentage of the lease life time otherwise use 1/3 (per RFC 4702) of the lease life time. It is specified as a decimal percentage (e.g. 0.75, 0.50) and is unspecified by default.
- `ddns-ttl-min` - If specified and greater than 0 use it the minimum otherwise use a minimum of 600 seconds (per RFC 4702). If the calculated TTL is less than the minimum return the minimum. It is specified in seconds.

By default it is unspecified.

- `ddns-ttl-max` - If specified and greater than zero limit the calculated TTL to its value. It is specified in seconds. By default it is unspecified.

Note

A value of zero for any of the TTL parameters can be used to suppress (i.e. "unspecify") the value of that parameter inherited from a higher scope.

9.2.21.1 DHCP-DDNS Server Connectivity

For NCRs to reach the D2 server, *kea-dhcp4* must be able to communicate with it. *kea-dhcp4* uses the following configuration parameters to control this communication:

- `enable-updates` - Enables connectivity to *kea-dhcp-ddns* such that DDNS updates can be constructed and sent. It must be `true` for NCRs to be generated and sent to D2. It defaults to `false`.
- `server-ip` - This is the IP address on which D2 listens for requests. The default is the local loopback interface at address 127.0.0.1. Either an IPv4 or IPv6 address may be specified.
- `server-port` - This is the port on which D2 listens for requests. The default value is 53001.
- `sender-ip` - This is the IP address which *kea-dhcp4* uses to send requests to D2. The default value is blank, which instructs *kea-dhcp4* to select a suitable address.
- `sender-port` - This is the port which *kea-dhcp4* uses to send requests to D2. The default value of 0 instructs *kea-dhcp4* to select a suitable port.
- `max-queue-size` - This is the maximum number of requests allowed to queue while waiting to be sent to D2. This value guards against requests accumulating uncontrollably if they are being generated faster than they can be delivered. If the number of requests queued for transmission reaches this value, DDNS updating is turned off until the queue backlog has been sufficiently reduced. The intent is to allow the *kea-dhcp4* server to continue lease operations without running the risk that its memory usage grows without limit. The default value is 1024.
- `ncr-protocol` - This specifies the socket protocol to use when sending requests to D2. Currently only UDP is supported.
- `ncr-format` - This specifies the packet format to use when sending requests to D2. Currently only JSON format is supported.

By default, *kea-dhcp-ddns* is assumed to be running on the same machine as *kea-dhcp4*, and all of the default values mentioned above should be sufficient. If, however, D2 has been configured to listen on a different address or port, these values must be altered accordingly. For example, if D2 has been configured to listen on 192.168.1.10 port 900, the following configuration is required:

```
"Dhcp4": {
  "dhcp-ddns": {
    "server-ip": "192.168.1.10",
    "server-port": 900,
    ...
  },
  ...
}
```

9.2.21.2 When Does the kea-dhcp4 Server Generate a DDNS Request?

The *kea-dhcp4* server follows the behavior prescribed for DHCP servers in [RFC 4702](#). It is important to keep in mind that *kea-dhcp4* makes the initial decision of when and what to update and forwards that information to D2 in the form of NCRs. Carrying out the actual DNS updates and dealing with such things as conflict resolution are within the purview of D2 itself (see [The DHCP-DDNS Server](#)). This section describes when *kea-dhcp4* generates NCRs and the configuration parameters that can be used to influence this decision. It assumes that both the connectivity parameter `enable-updates` and the behavioral parameter `ddns-send-updates`, are `true`.

In general, *kea-dhcp4* generates DDNS update requests when:

1. A new lease is granted in response to a DHCPREQUEST;
2. An existing lease is renewed but the FQDN associated with it has changed; or
3. An existing lease is released in response to a DHCPRELEASE.

In the second case, lease renewal, two DDNS requests are issued: one request to remove entries for the previous FQDN, and a second request to add entries for the new FQDN. In the third case, a lease release - a single DDNS request - to remove its entries is made.

As for the first case, the decisions involved when granting a new lease are more complex. When a new lease is granted, *kea-dhcp4* generates a DDNS update request if the DHCPREQUEST contains either the FQDN option (code 81) or the Host Name option (code 12). If both are present, the server uses the FQDN option. By default, *kea-dhcp4* respects the FQDN N and S flags specified by the client as shown in the following table:

Table 7: Default FQDN flag behavior

Client Flags:N-S	Client Intent	Server Response	Server Flags:N-S-O
0-0	Client wants to do forward updates, server should do reverse updates	Server generates reverse-only request	1-0-0
0-1	Server should do both forward and reverse updates	Server generates request to update both directions	0-1-0
1-0	Client wants no updates done	Server does not generate a request	1-0-0

The first row in the table above represents "client delegation." Here the DHCP client states that it intends to do the forward DNS updates and the server should do the reverse updates. By default, *kea-dhcp4* honors the client's wishes and generates a DDNS request to the D2 server to update only reverse DNS data. The parameter `ddns-override-client-update` can be used to instruct the server to override client delegation requests. When this parameter is `true`, *kea-dhcp4* disregards requests for client delegation and generates a DDNS request to update both forward and reverse DNS data. In this case, the N-S-O flags in the server's response to the client will be 0-1-1 respectively.

(Note that the flag combination N=1, S=1 is prohibited according to [RFC 4702](#). If such a combination is received from the client, the packet is dropped by *kea-dhcp4*.)

To override client delegation, set the following values in the configuration file:

```
"Dhcp4": {
  "ddns-override-client-update": true,
  ...
}
```

The third row in the table above describes the case in which the client requests that no DNS updates be done. The parameter `ddns-override-no-update` can be used to instruct the server to disregard the client's wishes. When this parameter is `true`, *kea-dhcp4* generates DDNS update requests to *kea-dhcp-ddns* even if the client requests that no updates be done. The N-S-O flags in the server's response to the client will be 0-1-1.

To override client delegation, issue the following commands:

```
"Dhcp4": {  
    "ddns-override-no-update": true,  
    ...  
}
```

The *kea-dhcp4* server always generates DDNS update requests if the client request only contains the Host Name option. In addition, it includes an FQDN option in the response to the client with the FQDN N-S-O flags set to 0-1-0, respectively. The domain name portion of the FQDN option is the name submitted to D2 in the DDNS update request.

9.2.21.3 *kea-dhcp4* Name Generation for DDNS Update Requests

Each NameChangeRequest must of course include the fully qualified domain name whose DNS entries are to be affected. *kea-dhcp4* can be configured to supply a portion or all of that name, based on what it receives from the client in the DHCPREQUEST.

The default rules for constructing the FQDN that will be used for DNS entries are:

1. If the DHCPREQUEST contains the client FQDN option, take the candidate name from there; otherwise, take it from the Host Name option.
2. If the candidate name is a partial (i.e. unqualified) name, then add a configurable suffix to the name and use the result as the FQDN.
3. If the candidate name provided is empty, generate an FQDN using a configurable prefix and suffix.
4. If the client provides neither option, then take no DNS action.

These rules can be amended by setting the `ddns-replace-client-name` parameter, which provides the following modes of behavior:

- **never** - use the name the client sent. If the client sent no name, do not generate one. This is the default mode.
- **always** - replace the name the client sent. If the client sent no name, generate one for the client.
- **when-present** - replace the name the client sent. If the client sent no name, do not generate one.
- **when-not-present** - use the name the client sent. If the client sent no name, generate one for the client.

Note

In early versions of Kea, this parameter was a boolean and permitted only values of `true` and `false`. Boolean values have been deprecated and are no longer accepted. Administrators currently using booleans must replace them with the desired mode name. A value of `true` maps to `when-present`, while `false` maps to `never`.

For example, to instruct *kea-dhcp4* to always generate the FQDN for a client, set the parameter `ddns-replace-client-name` to `always` as follows:

```
"Dhcp4": {  
    "ddns-replace-client-name": "always",  
    ...  
}
```

The prefix used in the generation of an FQDN is specified by the `ddns-generated-prefix` parameter. The default value is `"myhost"`. To alter its value, simply set it to the desired string:

```
"Dhcp4": {
  "ddns-generated-prefix": "another.host",
  ...
}
```

The suffix used when generating an FQDN, or when qualifying a partial name, is specified by the `ddns-qualifying-suffix` parameter. It is strongly recommended that the user supply a value for the qualifying suffix when DDNS updates are enabled. For obvious reasons, we cannot supply a meaningful default.

```
"Dhcp4": {
  "ddns-qualifying-suffix": "foo.example.org",
  ...
}
```

When qualifying a partial name, *kea-dhcp4* constructs the name in the format:

```
[candidate-name].[ddns-qualifying-suffix].
```

where `candidate-name` is the partial name supplied in the DHCPREQUEST. For example, if the FQDN domain name value is "some-computer" and the `ddns-qualifying-suffix` is "example.com", the generated FQDN is:

```
some-computer.example.com.
```

When generating the entire name, *kea-dhcp4* constructs the name in the format:

```
[ddns-generated-prefix]-[address-text].[ddns-qualifying-suffix].
```

where `address-text` is simply the lease IP address converted to a hyphenated string. For example, if the lease address is 172.16.1.10, the qualifying suffix is "example.com", and the default value is used for `ddns-generated-prefix`, the generated FQDN is:

```
myhost-172-16-1-10.example.com.
```

Note

When the client sends the host name option, *kea-dhcp4* never adds a dot to the host name in the returned option. It will only end with a dot if the client sent it already qualified and ending with a dot. When the client sends the FQDN option, the FQDN returned in the response will always end with a dot.

9.2.21.4 Sanitizing Client Host Name and FQDN Names

Some DHCP clients may provide values in the Host Name option (option code 12) or FQDN option (option code 81) that contain undesirable characters. It is possible to configure *kea-dhcp4* to sanitize these values. The most typical use case is ensuring that only characters that are permitted by RFC 1035 be included: A-Z, a-z, 0-9, and "-". This may be accomplished with the following two parameters:

- `hostname-char-set` - a regular expression describing the invalid character set. This can be any valid, regular expression using POSIX extended expression syntax. Embedded nulls (0x00) are always considered an invalid character to be replaced (or omitted). The default is "[^A-Za-z0-9.-]". This matches any character that is not a letter, digit, dot, hyphen, or null.
- `hostname-char-replacement` - a string of zero or more characters with which to replace each invalid character in the host name. An empty string causes invalid characters to be OMITTED rather than replaced. The default is "".

The following configuration replaces anything other than a letter, digit, dot, or hyphen with the letter "x":

```
"Dhcp4": {
    "hostname-char-set": "[^A-Za-z0-9.-]",
    "hostname-char-replacement": "x",
    ...
}
```

Thus, a client-supplied value of "myhost-\$(123.org" would become "myhost-xx123.org". Sanitizing is performed only on the portion of the name supplied by the client, and it is performed before applying a qualifying suffix (if one is defined and needed).

Note

Name sanitizing is meant to catch the more common cases of invalid characters through a relatively simple character-replacement scheme. It is difficult to devise a scheme that works well in all cases, for both Host Name and FQDN options. Administrators who find they have clients with odd corner cases of character combinations that cannot be readily handled with this mechanism should consider writing a hook that can carry out sufficiently complex logic to address their needs.

If clients include domain names in the Host Name option and the administrator wants these preserved, they need to make sure that the dot, ".", is considered a valid character by the `hostname-char-set` expression, such as this: "[^A-Za-z0-9.-]". This does not affect dots in FQDN Option values. When scrubbing FQDNs, dots are treated as delimiters and used to separate the option value into individual domain labels that are scrubbed and then re-assembled.

If clients are sending values that differ only by characters considered as invalid by the `hostname-char-set`, be aware that scrubbing them will yield identical values. In such cases, DDNS conflict rules will permit only one of them to register the name.

Finally, given the latitude clients have in the values they send, it is virtually impossible to guarantee that a combination of these two parameters will always yield a name that is valid for use in DNS. For example, using an empty value for `hostname-char-replacement` could yield an empty domain label within a name, if that label consists only of invalid characters.

Note

It is possible to specify `hostname-char-set` and/or `hostname-char-replacement` at the global scope.

The Kea hook library *libdhcp_ddns_tuning.so* provides the ability for both *kea-dhcp4* and *kea-dhcp6* to generate host names procedurally based on an expression, to skip DDNS updates on a per-client basis, or to fine-tune various DNS update aspects. Please refer to the *libdhcp_ddns_tuning.so: DDNS Tuning* documentation for the configuration options.

Note

Beginning with Kea 3.0.2, when sanitizing results in an empty string, the option (hostname or FQDN) will be ignored as if the client had not sent it. This can happen if the value sent by the client contains only characters that are deemed invalid by the expression in 'hostname-char-set' and 'hostname-char-replacement' is empty, "".

9.2.22 Next Server (siaddr)

In some cases, clients want to obtain configuration from a TFTP server. Although there is a dedicated option for it, some devices may use the `siaddr` field in the DHCPv4 packet for that purpose. That specific field can be configured using the `next-server` directive. It is possible to define it in the global scope or for a given subnet only. If both are defined, the subnet value takes precedence. The value in the subnet can be set to "0.0.0.0", which means that `next-server` should not be sent. It can also be set to an empty string, which is equivalent to it not being defined at all; that is, it uses the global value.

The `server-hostname` (which conveys a server hostname, can be up to 64 bytes long, and is in the `sname` field) and `boot-file-name` (which conveys the configuration file, can be up to 128 bytes long, and is sent using the `file` field) directives are handled the same way as `next-server`.

```
"Dhcp4": {
  "next-server": "192.0.2.123",
  "boot-file-name": "/dev/null",
  "subnet4": [
    {
      "next-server": "192.0.2.234",
      "server-hostname": "some-name.example.org",
      "boot-file-name": "bootfile.efi",
      ...
    }
  ],
  ...
}
```

9.2.23 Echoing Client-ID (RFC 6842)

The original DHCPv4 specification ([RFC 2131](#)) states that the DHCPv4 server must not send back client-id options when responding to clients. However, in some cases that results in confused clients that do not have a MAC address or client-id; see [RFC 6842](#) for details. That behavior changed with the publication of [RFC 6842](#), which updated [RFC 2131](#). That update states that the server must send the client-id if the client sent it, and that is Kea's default behavior. However, in some cases older devices that do not support [RFC 6842](#) may refuse to accept responses that include the client-id option. To enable backward compatibility, an optional configuration parameter has been introduced. To configure it, use the following configuration statement:

```
"Dhcp4": {
  "echo-client-id": false,
  ...
}
```

9.2.24 Using Client Identifier and Hardware Address

The DHCP server must be able to identify the client from which it receives the message and distinguish it from other clients. There are many reasons why this identification is required; the most important ones are:

- When the client contacts the server to allocate a new lease, the server must store the client identification information in the lease database as a search key.
- When the client tries to renew or release the existing lease, the server must be able to find the existing lease entry in the database for this client, using the client identification information as a search key.
- Some configurations use static reservations for the IP addresses and other configuration information. The server's administrator uses client identification information to create these static assignments.

- In dual-stack networks there is often a need to correlate the lease information stored in DHCPv4 and DHCPv6 servers for a particular host. Using common identification information by the DHCPv4 and DHCPv6 clients allows the network administrator to achieve this correlation and better administer the network. Beginning with release 2.1.2, Kea supports DHCPv6 DUIDs embedded within DHCPv4 Client Identifier options as described in [RFC 4361](#).

DHCPv4 uses two distinct identifiers which are placed by the client in the queries sent to the server and copied by the server to its responses to the client: `chaddr` and `client-identifier`. The former was introduced as a part of the BOOTP specification and it is also used by DHCP to carry the hardware address of the interface used to send the query to the server (MAC address for the Ethernet). The latter is carried in the `client-identifier` option, introduced in [RFC 2132](#).

[RFC 2131](#) indicates that the server may use both of these identifiers to identify the client but the client identifier, if present, takes precedence over `chaddr`. One of the reasons for this is that the client identifier is independent from the hardware used by the client to communicate with the server. For example, if the client obtained the lease using one network card and then the network card is moved to another host, the server will wrongly identify this host as the one which obtained the lease. Moreover, [RFC 4361](#) gives the recommendation to use a DUID (see [RFC 8415](#), the DHCPv6 specification) carried as a client identifier when dual-stack networks are in use to provide consistent identification information for the client, regardless of the type of protocol it is using. Kea adheres to these specifications, and the client identifier by default takes precedence over the value carried in the `chaddr` field when the server searches, creates, updates, or removes the client's lease.

When the server receives a DHCPDISCOVER or DHCPREQUEST message from the client, it tries to find out if the client already has a lease in the database; if it does, the server hands out that lease rather than allocates a new one. Each lease in the lease database is associated with the client identifier and/or `chaddr`. The server first uses the client identifier (if present) to search for the lease; if one is found, the server treats this lease as belonging to the client, even if the current `chaddr` and the `chaddr` associated with the lease do not match. This facilitates the scenario when the network card on the client system has been replaced and thus the new MAC address appears in the messages sent by the DHCP client. If the server fails to find the lease using the client identifier, it performs another lookup using the `chaddr`. If this lookup returns no result, the client is considered to not have a lease and a new lease is created.

A common problem reported by network operators is that poor client implementations do not use stable client identifiers, instead generating a new client identifier each time the client connects to the network. Another well-known case is when the client changes its client identifier during the multi-stage boot process (PXE). In such cases, the MAC address of the client's interface remains stable, and using the `chaddr` field to identify the client guarantees that the particular system is considered to be the same client, even though its client identifier changes.

To address this problem, Kea includes a configuration option which enables client identification using `chaddr` only. This instructs the server to ignore the client identifier during lease lookups and allocations for a particular subnet. Consider the following simplified server configuration:

```
{
  "Dhcp4": {
    "match-client-id": true,
    "subnet4": [
      {
        "id": 1,
        "subnet": "192.0.10.0/24",
        "pools": [ { "pool": "192.0.2.23-192.0.2.87" } ],
        "match-client-id": false
      },
      {
        "id": 1,
        "subnet": "10.0.0.0/8",
        "pools": [ { "pool": "10.0.0.23-10.0.2.99" } ]
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    ]
  }
}

```

The `match-client-id` parameter is a boolean value which controls this behavior. The default value of `true` indicates that the server will use the client identifier for lease lookups and `chaddr` if the first lookup returns no results. `false` means that the server will only use the `chaddr` to search for the client's lease. Whether the DHCID for DNS updates is generated from the client identifier or `chaddr` is controlled through the same parameter.

The `match-client-id` parameter may appear both in the global configuration scope and/or under any subnet declaration. In the example shown above, the effective value of the `match-client-id` will be `false` for the subnet 192.0.10.0/24, because the subnet-specific setting of the parameter overrides the global value of the parameter. The effective value of the `match-client-id` for the subnet 10.0.0.0/8 will be set to `true`, because the subnet declaration lacks this parameter and the global setting is by default used for this subnet. In fact, the global entry for this parameter could be omitted in this case, because `true` is the default value.

It is important to understand what happens when the client obtains its lease for one setting of the `match-client-id` and then renews it when the setting has been changed. First, consider the case when the client obtains the lease and the `match-client-id` is set to `true`. The server stores the lease information, including the client identifier (if supplied) and `chaddr`, in the lease database. When the setting is changed and the client renews the lease, the server will determine that it should use the `chaddr` to search for the existing lease. If the client has not changed its MAC address, the server should successfully find the existing lease. The client identifier associated with the returned lease will be ignored and the client will be allowed to use this lease. When the lease is renewed only the `chaddr` will be recorded for this lease, according to the new server setting.

In the second case, the client has the lease with only a `chaddr` value recorded. When the `match-client-id` setting is changed to `true`, the server will first try to use the client identifier to find the existing client's lease. This will return no results because the client identifier was not recorded for this lease. The server will then use the `chaddr` and the lease will be found. If the lease appears to have no client identifier recorded, the server will assume that this lease belongs to the client and that it was created with the previous setting of the `match-client-id`. However, if the lease contains a client identifier which is different from the client identifier used by the client, the lease will be assumed to belong to another client and a new lease will be allocated.

For a more visual representation of how Kea recognizes the same client, please refer to [How Kea Recognizes the Same Client In Different DHCP Messages](#).

9.2.25 Authoritative DHCPv4 Server Behavior

The original DHCPv4 specification ([RFC 2131](#)) states that if a client requests an address in the INIT-REBOOT state of which the server has no knowledge, the server must remain silent, except if the server knows that the client has requested an IP address from the wrong network. By default, Kea follows the behavior of the ISC `dhcpcd` daemon instead of the specification and also remains silent if the client requests an IP address from the wrong network, because configuration information about a given network segment is not known to be correct. Kea only rejects a client's DHCPREQUEST with a DHCPNAK message if it already has a lease for the client with a different IP address. Administrators can override this behavior through the boolean `authoritative` (`false` by default) setting.

In authoritative mode, `authoritative` set to `true`, Kea always rejects INIT-REBOOT requests from unknown clients with DHCPNAK messages. The `authoritative` setting can be specified in global, shared-network, and subnet configuration scope and is automatically inherited from the parent scope, if not specified. All subnets in a shared-network must have the same `authoritative` setting.

9.2.26 DHCPv4-over-DHCPv6: DHCPv4 Side

The support of DHCPv4-over-DHCPv6 transport is described in [RFC 7341](#) and is implemented using cooperating DHCPv4 and DHCPv6 servers. This section is about the configuration of the DHCPv4 side (the DHCPv6 side is described in *DHCPv4-over-DHCPv6: DHCPv6 Side*).

Note

The DHCPv4-over-DHCPv6 details of the inter-process communication may change; for instance, the support of port relay (RFC 8357) introduced an incompatible change. Both the DHCPv4 and DHCPv6 sides should be running the same version of Kea.

The `dhcp4o6-port` global parameter specifies the first of the two consecutive ports of the UDP sockets used for the communication between the DHCPv6 and DHCPv4 servers. The DHCPv4 server is bound to `::1` on `port + 1` and connected to `::1` on `port`.

With DHCPv4-over-DHCPv6, the DHCPv4 server does not have access to several of the identifiers it would normally use to select a subnet. To address this issue, three new configuration entries are available; the presence of any of these allows the subnet to be used with DHCPv4-over-DHCPv6. These entries are:

- `4o6-subnet`: takes a prefix (i.e., an IPv6 address followed by a slash and a prefix length) which is matched against the source address.
- `4o6-interface-id`: takes a relay interface ID option value.
- `4o6-interface`: takes an interface name which is matched against the incoming interface name.

ISC tested the following configuration:

```
{
# DHCPv4 conf
"Dhcp4": {
    "interfaces-config": {
        "interfaces": [ "eno33554984" ]
    },
    "lease-database": {
        "type": "memfile",
        "name": "leases4"
    },
    "valid-lifetime": 4000,
    "subnet4": [
    {
        "id": 1,
        "subnet": "10.10.10.0/24",
        "4o6-interface": "eno33554984",
        "4o6-subnet": "2001:db8:1:1::/64",
        "pools": [ { "pool": "10.10.10.100 - 10.10.10.199" } ]
    }
    ],
    "dhcp4o6-port": 786,
```

(continues on next page)

(continued from previous page)

```
"loggers": [  
  {  
    "name": "kea-dhcp4",  
    "output-options": [  
      {  
        "output": "kea-dhcp4.log"  
      }  
    ],  
    "severity": "DEBUG",  
    "debuglevel": 0  
  }  
]  
}
```

9.2.27 Sanity Checks in DHCPv4

An important aspect of a well-running DHCP system is an assurance that the data remains consistent; however, in some cases it may be convenient to tolerate certain inconsistent data. For example, a network administrator who temporarily removes a subnet from a configuration would not want all the leases associated with it to disappear from the lease database. Kea has a mechanism to implement sanity checks for situations like this.

Kea supports a configuration scope called `sanity-checks`. A parameter, called `lease-checks`, governs the verification carried out when a new lease is loaded from a lease file. This mechanism permits Kea to attempt to correct inconsistent data.

Every subnet has a `subnet-id` value; this is how Kea internally identifies subnets. Each lease has a `subnet-id` parameter as well, which identifies the subnet it belongs to. However, if the configuration has changed, it is possible that a lease could exist with a `subnet-id` but without any subnet that matches it. Also, it is possible that the subnet's configuration has changed and the `subnet-id` now belongs to a subnet that does not match the lease.

Kea's corrective algorithm first checks to see if there is a subnet with the `subnet-id` specified by the lease. If there is, it verifies whether the lease belongs to that subnet. If not, depending on the `lease-checks` setting, the lease is discarded, a warning is displayed, or a new subnet is selected for the lease that matches it topologically.

There are five levels which are supported:

- **none** - do no special checks; accept the lease as is.
- **warn** - if problems are detected display a warning, but accept the lease data anyway. This is the default value.
- **fix** - if a data inconsistency is discovered, try to correct it. If the correction is not successful, insert the incorrect data anyway.
- **fix-del** - if a data inconsistency is discovered, try to correct it. If the correction is not successful, reject the lease. This setting ensures the data's correctness, but some incorrect data may be lost. Use with care.
- **del** - if any inconsistency is detected, reject the lease. This is the strictest mode; use with care.

This feature is currently implemented for the memfile backend. The sanity check applies to the lease database in memory, not to the lease file, i.e. inconsistent leases stay in the lease file.

An example configuration that sets this parameter looks as follows:

```
"Dhcp4": {
  "sanity-checks": {
    "lease-checks": "fix-del"
  },
  ...
}
```

9.2.28 Storing Extended Lease Information

To support such features as DHCP Leasequery (RFC 4388) and stash agent options (*Stash Agent Options*), additional information must be stored with each lease. Because the amount of information for each lease has ramifications in terms of performance and system resource consumption, storage of this additional information is configurable through the `store-extended-info` parameter. It defaults to `false` and may be set at the global, shared-network, and subnet levels.

```
"Dhcp4": {
  "store-extended-info": true,
  ...
}
```

When set to `true`, information relevant to the DHCPREQUEST asking for the lease is added into the lease's user context as a map element labeled "ISC". When the DHCPREQUEST received contains the option (DHCP Option 82), the map contains the `relay-agent-info` map with the content option (DHCP Option 82) in the `sub-options` entry and, when present, the `remote-id` and `relay-id` options. Since DHCPREQUESTs sent as renewals are not likely to contain this information, the values taken from the last DHCPREQUEST that did contain it are retained on the lease. The lease's user context looks something like this:

```
{ "ISC": { "relay-agent-info": { "sub-options": "0x0104AABBCCDD" } } }
```

Or with remote and relay sub-options:

```
{
  "ISC": {
    "relay-agent-info": {
      "sub-options": "0x02030102030C03AABBCC",
      "remote-id": "03010203",
      "relay-id": "AABBCC"
    }
  }
}
```

Note

It is possible that other hook libraries are already using `user-context`. Enabling `store-extended-info` should not interfere with any other `user-context` content, as long as it does not also use an element labeled "ISC". In other words, `user-context` is intended to be a flexible container serving multiple purposes. As long as no other purpose also writes an "ISC" element to `user-context` there should not be a conflict.

Extended lease information is also subject to configurable sanity checking. The parameter in the `sanity-checks` scope is named `extended-info-checks` and supports these levels:

- **none** - do no check nor upgrade. This level should be used only when extended info is not used at all or when no badly formatted extended info, including using the old format, is expected.

- **fix** - fix some common inconsistencies and upgrade extended info using the old format to the new one. It is the default level and is convenient when the Leasequery hook library is not loaded.
- **strict** - fix all inconsistencies which have an impact on the (Bulk) Leasequery hook library.
- **pedantic** - enforce full conformance to the format produced by the Kea code; for instance, no extra entries are allowed with the exception of **comment**.

Note

This feature is currently implemented only for the memfile backend. The sanity check applies to the lease database in memory, not to the lease file, i.e. inconsistent leases stay in the lease file.

9.2.29 Stash Agent Options

This global parameter was added in version 2.5.8 to mirror a feature that was previously available in ISC DHCP. When the `stash-agent-options` parameter is `true`, the server records the relay agent information options sent during the client's initial DHCPREQUEST message (when the client was in the `SELECTING` state) and behaves as if those options are included in all subsequent DHCPREQUEST messages sent in the `RENEWING` state. This works around a problem with relay agent information options, which do not usually appear in DHCPREQUEST messages sent by the client in the `RENEWING` state; such messages are unicast directly to the server and are not sent through a relay agent.

The default is `false`.

9.2.30 Multi-Threading Settings

The Kea server can be configured to process packets in parallel using multiple threads. These settings can be found under the `multi-threading` structure and are represented by:

- **enable-multi-threading** - use multiple threads to process packets in parallel. The default is `true`.
- **thread-pool-size** - specify the number of threads to process packets in parallel. It may be set to `0` (auto-detect), or any positive number that explicitly sets the thread count. The default is `0`.
- **packet-queue-size** - specify the size of the queue used by the thread pool to process packets. It may be set to `0` (unlimited), or any positive number that explicitly sets the queue size. The default is `64`.

An example configuration that sets these parameters looks as follows:

```
"Dhcp4": {
  "multi-threading": {
    "enable-multi-threading": true,
    "thread-pool-size": 4,
    "packet-queue-size": 16
  },
  ...
}
```

9.2.31 Multi-Threading Settings With Different Database Backends

The Kea DHCPv4 server is benchmarked by ISC to determine which settings give the best performance. Although this section describes our results, they are merely recommendations and are very dependent on the particular hardware used for benchmarking. We strongly advise that administrators run their own performance benchmarks.

A full report of performance results for the latest stable Kea version can be found [here](#). This includes hardware and benchmark scenario descriptions, as well as current results.

After enabling multi-threading, the number of threads is set by the `thread-pool-size` parameter. Results from our experiments show that the best settings for *kea-dhcp4* are:

- `thread-pool-size`: 4 when using `memfile` for storing leases.
- `thread-pool-size`: 12 or more when using `mysql` for storing leases.
- `thread-pool-size`: 8 when using `postgresql`.

Another very important parameter is `packet-queue-size`; in our benchmarks we used it as a multiplier of `thread-pool-size`. The actual setting strongly depends on `thread-pool-size`.

We saw the best results in our benchmarks with the following settings:

- `packet-queue-size`: $7 * \text{thread-pool-size}$ when using `memfile` for storing leases; in our case it was $7 * 4 = 28$. This means that at any given time, up to 28 packets could be queued.
- `packet-queue-size`: $66 * \text{thread-pool-size}$ when using `mysql` for storing leases; in our case it was $66 * 12 = 792$. This means that up to 792 packets could be queued.
- `packet-queue-size`: $11 * \text{thread-pool-size}$ when using `postgresql` for storing leases; in our case it was $11 * 8 = 88$.

9.2.32 IPv6-Only Preferred Networks

[RFC 8925](#), recently published by the IETF, specifies a DHCPv4 option to indicate that a host supports an IPv6-only mode and is willing to forgo obtaining an IPv4 address if the network provides IPv6 connectivity. The general idea is that a network administrator can enable this option to signal to compatible dual-stack devices that IPv6 connectivity is available and they can shut down their IPv4 stack. The new option `v6-only-preferred` content is a 32-bit unsigned integer and specifies for how long the device should disable its stack. The value is expressed in seconds.

The RFC mentions the `V6ONLY_WAIT` timer. This is implemented in Kea by setting the value of the `v6-only-preferred` option. This follows the usual practice of setting options; the option value can be specified on the pool, subnet, shared network, or global levels, or even via host reservations.

There is no special processing involved; it follows the standard Kea option processing regime. The option is not sent back unless the client explicitly requests it. For example, to enable the option for the whole subnet, the following configuration can be used:

```
{
  "subnet4": [
    {
      "id": 1,
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
      "subnet": "192.0.2.0/24",
      "option-data": [
        {
          // This will make the v6-only capable devices to disable their
          // v4 stack for half an hour and then try again
          "name": "v6-only-preferred",
          "data": "1800"
        }
      ]
    }
  ],
  ...
}
```

9.2.33 Lease Caching

Clients that attempt multiple renewals in a short period can cause the server to update and write to the database frequently, resulting in a performance impact on the server. The cache parameters instruct the DHCP server to avoid updating leases too frequently, thus avoiding this behavior. Instead, the server assigns the same lease (i.e. reuses it) with no modifications except for CLTT (Client Last Transmission Time), which does not require disk operations.

The two parameters are the `cache-threshold` double and the `cache-max-age` integer. These parameters can be configured at the global, shared-network, and subnet levels. The subnet level has precedence over the shared-network level, while the global level is used as a last resort. For example:

```
{
  "subnet4": [
    {
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
      "subnet": "192.0.2.0/24",
      "cache-threshold": .25,
      "cache-max-age": 600,
      "valid-lifetime": 2000,
      ...
    }
  ],
  ...
}
```

Note

Beginning with Kea 2.7.8, the default value for `cache-threshold` is 0.25 while `cache-max-age` is left unspecified. This means that lease caching is enabled by default. It may be disabled at the desired scope either by setting `cache-threshold` to 0.0 or setting `cache-max-age` to 0.

When an already-assigned lease can fulfill a client query:

- any important change, e.g. for DDNS parameter, hostname, or valid lifetime reduction, makes the lease not reusable.
- lease age, i.e. the difference between the creation or last modification time and the current time, is computed (elapsed duration).
- if `cache-max-age` is explicitly configured, it is compared with the lease age; leases that are too old are not reusable. This means that the value 0 for `cache-max-age` disables the lease cache feature.
- if `cache-threshold` is explicitly configured and is between 0.0 and 1.0, it expresses the percentage of the lease valid lifetime which is allowed for the lease age. Values below and including 0.0 and values greater than 1.0 disable the lease cache feature.

In our example, a lease with a valid lifetime of 2000 seconds can be reused if it was committed less than 500 seconds ago. With a lifetime of 3000 seconds, a maximum age of 600 seconds applies.

In outbound client responses (e.g. DHCPACK messages), the `dhcp-lease-time` option is set to the reusable valid lifetime, i.e. the expiration date does not change. Other options based on the valid lifetime e.g. `dhcp-renewal-time` and `dhcp-rebinding-time`, also depend on the reusable lifetime.

9.2.34 Temporary Allocation on DHCPDISCOVER

By default, *kea-dhcp4* does not allocate or store a lease when offering an address to a client in response to a DHCPDISCOVER. In general, *kea-dhcp4* can fulfill client demands faster by deferring lease allocation and storage until it receives DHCPREQUESTs for them. The `offer-lifetime` parameter in *kea-dhcp4* (when not zero) instructs the server to allocate and persist a lease when generating a DHCPOFFER. In addition:

- The persisted lease's lifetime is equal to `offer-lifetime` (in seconds).
- The lifetime sent to the client in the DHCPOFFER via option 51 is still based on `valid-lifetime`. This avoids issues with clients that may reject offers whose lifetimes they perceive as too short.
- DDNS updates are not performed. As with the default behavior, those updates occur on DHCPREQUEST.
- Updates are not sent to HA peers.
- Assigned lease statistics are incremented.
- Expiration processing and reclamation behave just as they do for leases allocated during DHCPREQUEST processing.
- Lease caching, if enabled, is honored.
- In sites running multiple instances of *kea-dhcp4* against a single, shared lease store, races for given address values are lost during DHCPDISCOVER processing rather than during DHCPREQUEST processing. Servers that lose the race for the address simply do not respond to the client, rather than NAK them. The client in turn simply retries its DHCPDISCOVER. This should reduce the amount of traffic such conflicts incur.
- Clients repeating DHCPDISCOVERs are offered the same address each time.

An example subnet configuration is shown below:

```
{
  "subnet4": [
    {
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
      "subnet": "192.0.2.0/24",
      "offer-lifetime": 60,
      "valid-lifetime": 2000,
      ...
    }
  ],
  ...
}
```

Here `offer-lifetime` has been configured to be 60 seconds, with a `valid-lifetime` of 2000 seconds. This instructs *kea-dhcp4* to persist leases for 60 seconds when sending them back in DHCPOFFERs, and then extend them to 2000 seconds when clients DHCPREQUEST them.

The value, which defaults to 0, is supported at the global, shared-network, subnet, and class levels. Choosing an appropriate value for `offer-lifetime` is extremely site-dependent, but a value between 60 and 120 seconds is a reasonable starting point.

9.2.35 DNR (Discovery of Network-designated Resolvers) Options for DHCPv4

The Discovery of Network-designated Resolvers, or DNR option, was introduced in [RFC 9463](#) as a way to communicate location of DNS resolvers available over means other than the classic DNS over UDP over port 53. As of spring 2024, the supported technologies are DoT (DNS-over-TLS), DoH (DNS-over-HTTPS), and DoQ (DNS-over-QUIC), but the option was designed to be extensible to accommodate other protocols in the future.

The DHCPv4 option and its corresponding DHCPv6 options are almost exactly the same, with the exception of cardinality: only one DHCPv4 option is allowed, while multiple options are allowed for DHCPv6. To be able to convey multiple entries, the DHCPv4 DNR is an array that allows multiple DNS instances. Each instance is logically equal to one DHCPv6 option; the only difference is that it uses IPv4 rather than IPv6 addresses. DNR DHCPv4 options allow more than one DNR instance to be configured, and the DNR instances are separated with the "pipe" (0x7C) character.

For a detailed example of how to configure the DNR option, see [DNR \(Discovery of Network-designated Resolvers\) Options for DHCPv6](#).

For each DNR instance, comma-delimited fields must be provided in the following order:

- Service Priority (mandatory),
- ADN FQDN (mandatory),
- IP address(es) (optional; if more than one, they must be separated by spaces)
- SvcParams as a set of key=value pairs (optional; if more than one - they must be separated by spaces) To provide more than one alpn-id, separate them with double backslash-escaped commas as in the example below).

Example usage:

```
{
  "name": "v4-dnr",
  // 2 DNR Instances:
  // - Service priority 2, ADN, resolver IPv4 address and Service Parameters
  // - Service priority 3, ADN - this is ADN-only mode as per RFC 9463 3.1.6
  "data": "2, resolver.example., 10.0.5.6, alpn=dot\\,doq port=8530 | 3, fooexp.resolver.
  →example."
}
```

Note

If "comma" or "pipe" characters are used as text rather than as field delimiters, they must be escaped with double backslashes (\\, or \\|). Escaped commas must be used when configuring more than one ALPN protocol, to separate them. The "pipe" (0x7C) character can be used in the dohpath service parameter, as it is allowed in a URI.

Examples for DNR DHCPv4 options are provided in the Kea sources, in *all-options.json* in the *doc/examples/kea4* directory.

9.3 Host Reservations in DHCPv4

There are many cases where it is useful to provide a configuration on a per-host basis. The most obvious one is to reserve a specific, static address for exclusive use by a given client (host); the returning client receives the same address from the server every time, and other clients generally do not receive that address. Host reservations are also convenient when a host has specific requirements, e.g. a printer that needs additional DHCP options. Yet another possible use case is to define unique names for hosts.

There may be cases when a new reservation has been made for a client for an address currently in use by another client. We call this situation a "conflict." These conflicts get resolved automatically over time, as described in subsequent sections. Once a conflict is resolved, the correct client will receive the reserved configuration when it renews.

Host reservations are defined as parameters for each subnet. Each host must have its own unique identifier, such as the hardware/MAC address. There is an optional `reservations` array in the `subnet4` structure; each element in that array is a structure that holds information about reservations for a single host. In particular, the structure has an identifier that uniquely identifies a host. In the DHCPv4 context, the identifier is usually a hardware or MAC address. In most

cases an IP address will be specified. It is also possible to specify a hostname, host-specific options, or fields carried within the DHCPv4 message such as `siaddr`, `sname`, or `file`.

Note

The reserved address must be within the subnet.

The following example shows how to reserve addresses for specific hosts in a subnet:

```
{
  "subnet4": [
    {
      "id": 1,
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.200" } ],
      "subnet": "192.0.2.0/24",
      "interface": "eth0",
      "reservations": [
        {
          "hw-address": "1a:1b:1c:1d:1e:1f",
          "ip-address": "192.0.2.202"
        },
        {
          "duid": "0a:0b:0c:0d:0e:0f",
          "ip-address": "192.0.2.100",
          "hostname": "alice-laptop"
        },
        {
          "circuit-id": "'charter950'",
          "ip-address": "192.0.2.203"
        },
        {
          "client-id": "01:11:22:33:44:55:66",
          "ip-address": "192.0.2.204"
        }
      ]
    }
  ],
  ...
}
```

The first entry reserves the 192.0.2.202 address for the client that uses a MAC address of 1a:1b:1c:1d:1e:1f. The second entry reserves the address 192.0.2.100 and the hostname of "alice-laptop" for the client using a DUID 0a:0b:0c:0d:0e:0f. (If DNS updates are planned, it is strongly recommended that the hostnames be unique.) The third example reserves address 192.0.3.203 for a client whose request would be relayed by a relay agent that inserts a `circuit-id` option with the value "charter950". The fourth entry reserves address 192.0.2.204 for a client that uses a client identifier with value 01:11:22:33:44:55:66.

The above example is used for illustrational purposes only; in actual deployments it is recommended to use as few types as possible (preferably just one). See *Fine-Tuning DHCPv4 Host Reservation* for a detailed discussion of this point.

Making a reservation for a mobile host that may visit multiple subnets requires a separate host definition in each subnet that host is expected to visit. It is not possible to define multiple host definitions with the same hardware address in a single subnet. Multiple host definitions with the same hardware address are valid if each is in a different subnet.

Adding host reservations incurs a performance penalty. In principle, when a server that does not support host reservation

responds to a query, it needs to check whether there is a lease for a given address being considered for allocation or renewal. The server that does support host reservation has to perform additional checks: not only whether the address is currently used (i.e., if there is a lease for it), but also whether the address could be used by someone else (i.e., if there is a reservation for it). That additional check incurs extra overhead.

9.3.1 Address Reservation Types

In a typical Kea scenario there is an IPv4 subnet defined, e.g. 192.0.2.0/24, with a certain part of it dedicated for dynamic allocation by the DHCPv4 server. That dynamic part is referred to as a dynamic pool or simply a pool. In principle, a host reservation can reserve any address that belongs to the subnet. The reservations that specify addresses that belong to configured pools are called "in-pool reservations." In contrast, those that do not belong to dynamic pools are called "out-of-pool reservations." There is no formal difference in the reservation syntax and both reservation types are handled uniformly.

Kea supports global host reservations. These are reservations that are specified at the global level within the configuration and that do not belong to any specific subnet. Kea still matches inbound client packets to a subnet as before, but when the subnet's reservation mode is set to "global", Kea looks for host reservations only among the global reservations defined. Typically, such reservations would be used to reserve hostnames for clients which may move from one subnet to another.

Note

Global reservations, while useful in certain circumstances, have aspects that must be given due consideration when using them. Please see [Conflicts in DHCPv4 Reservations](#) for more details.

Note

Since Kea 1.9.1, reservation mode has been replaced by three boolean flags, `reservations-global`, `reservations-in-subnet`, and `reservations-out-of-pool`, which allow the configuration of host reservations both globally and in a subnet. In such cases a subnet host reservation has preference over a global reservation when both exist for the same client.

9.3.2 Conflicts in DHCPv4 Reservations

As reservations and lease information are stored separately, conflicts may arise. Consider the following series of events: the server has configured the dynamic pool of addresses from the range of 192.0.2.10 to 192.0.2.20. Host A requests an address and gets 192.0.2.10. Now the system administrator decides to reserve address 192.0.2.10 for Host B. In general, reserving an address that is currently assigned to someone else is not recommended, but there are valid use cases where such an operation is warranted.

The server now has a conflict to resolve. If Host B boots up and requests an address, the server cannot immediately assign the reserved address 192.0.2.10. A naive approach would be to immediately remove the existing lease for Host A and create a new one for Host B. That would not solve the problem, though, because as soon as Host B gets the address, it will detect that the address is already in use (by Host A) and will send a DHCPDECLINE message. Therefore, in this situation, the server has to temporarily assign a different address from the dynamic pool (not matching what has been reserved) to Host B.

When Host A renews its address, the server will discover that the address being renewed is now reserved for another host - Host B. The server will inform Host A that it is no longer allowed to use it by sending a DHCPNAK message. The server will not remove the lease, though, as there's a small chance that the DHCPNAK will not be delivered if the network is lossy. If that happens, the client will not receive any responses, so it will retransmit its DHCPREQUEST packet. Once the DHCPNAK is received by Host A, it will revert to server discovery and will eventually get a different

address. Besides allocating a new lease, the server will also remove the old one. As a result, address 192.0.2.10 will become free.

When Host B tries to renew its temporarily assigned address, the server will detect that it has a valid lease, but will note that there is a reservation for a different address. The server will send DHCPNAK to inform Host B that its address is no longer usable, but will keep its lease (again, the DHCPNAK may be lost, so the server will keep it until the client returns for a new address). Host B will revert to the server discovery phase and will eventually send a DHCPREQUEST message. This time the server will find that there is a reservation for that host and that the reserved address 192.0.2.10 is not used, so it will be granted. It will also remove the lease for the temporarily assigned address that Host B previously obtained.

This recovery will succeed, even if other hosts attempt to get the reserved address. If Host C requests the address 192.0.2.10 after the reservation is made, the server will either offer a different address (when responding to DHCPDISCOVER) or send DHCPNAK (when responding to DHCPREQUEST).

This mechanism allows the server to fully recover from a case where reservations conflict with existing leases; however, this procedure takes roughly as long as the value set for `renew-timer`. The best way to avoid such a recovery is not to define new reservations that conflict with existing leases. Another recommendation is to use out-of-pool reservations; if the reserved address does not belong to a pool, there is no way that other clients can get it.

Note

The conflict-resolution mechanism does not work for global reservations. Although the global address reservations feature may be useful in certain settings, it is generally recommended not to use global reservations for addresses. Administrators who do choose to use global reservations must manually ensure that the reserved addresses are not in dynamic pools.

9.3.3 Reserving a Hostname

When the reservation for a client includes the `hostname`, the server returns this hostname to the client in the Client FQDN or Hostname option. The server responds with the Client FQDN option only if the client has included the Client FQDN option in its message to the server. The server responds with the Hostname option if the client included the Hostname option in its message to the server, or if the client requested the Hostname option using the Parameter Request List option. The server returns the Hostname option even if it is not configured to perform DNS updates. The reserved hostname always takes precedence over the hostname supplied by the client or the autogenerated (from the IPv4 address) hostname.

The server qualifies the reserved hostname with the value of the `ddns-qualifying-suffix` parameter. For example, the following subnet configuration:

```
{
  "subnet4": [
    {
      "id": 1,
      "subnet": "10.0.0.0/24",
      "pools": [ { "pool": "10.0.0.10-10.0.0.100" } ],
      "ddns-qualifying-suffix": "example.isc.org.",
      "reservations": [
        {
          "hw-address": "aa:bb:cc:dd:ee:ff",
          "hostname": "alice-laptop"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "dhcp-ddns": {
        "enable-updates": true
    }
}

```

will result in the "alice-laptop.example.isc.org." hostname being assigned to the client using the MAC address "aa:bb:cc:dd:ee:ff". If the `ddns-qualifying-suffix` is not specified, the default (empty) value will be used, and in this case the value specified as a `hostname` will be treated as a fully qualified name. Thus, by leaving the `ddns-qualifying-suffix` empty it is possible to qualify hostnames for different clients with different domain names:

```

{
  "subnet4": [
    {
      "id": 1,
      "subnet": "10.0.0.0/24",
      "pools": [ { "pool": "10.0.0.10-10.0.0.100" } ],
      "reservations": [
        {
          "hw-address": "aa:bb:cc:dd:ee:ff",
          "hostname": "alice-laptop.isc.org."
        },
        {
          "hw-address": "12:34:56:78:99:AA",
          "hostname": "mark-desktop.example.org."
        }
      ]
    }
  ],
  "dhcp-ddns": {
    "enable-updates": true
  }
}

```

The above example results in the assignment of the "alice-laptop.isc.org." hostname to the client using the MAC address "aa:bb:cc:dd:ee:ff", and the hostname "mark-desktop.example.org." to the client using the MAC address "12:34:56:78:99:AA".

9.3.4 Including Specific DHCPv4 Options in Reservations

Kea offers the ability to specify options on a per-host basis. These options follow the same rules as any other options. These can be standard options (see *Standard DHCPv4 Options*), custom options (see *Custom DHCPv4 Options*), or vendor-specific options (see *DHCPv4 Vendor-Specific Options*). The following example demonstrates how standard options can be defined:

```

{
  "subnet4": [
    {
      "reservations": [
        {
          "hw-address": "aa:bb:cc:dd:ee:ff",
          "ip-address": "192.0.2.1",
          "option-data": [

```

(continues on next page)

(continued from previous page)

```

        {
            "name": "cookie-servers",
            "data": "10.1.1.202,10.1.1.203"
        },
        {
            "name": "log-servers",
            "data": "10.1.1.200,10.1.1.201"
        }
    ]
}
],
...
}
],
...
}

```

Vendor-specific options can be reserved in a similar manner:

```

{
    "subnet4": [
        {
            "reservations": [
                {
                    "hw-address": "aa:bb:cc:dd:ee:ff",
                    "ip-address": "10.0.0.7",
                    "option-data": [
                        {
                            "name": "vivso-suboptions",
                            "data": "4491"
                        },
                        {
                            "name": "tftp-servers",
                            "space": "vendor-4491",
                            "data": "10.1.1.202,10.1.1.203"
                        }
                    ]
                }
            ]
        }
    ],
    ...
}
],
...
}

```

Options defined at the host level have the highest priority. In other words, if there are options defined with the same type on the global, subnet, class, and host levels, the host-specific values are used.

9.3.5 Reserving Next Server, Server Hostname, and Boot File Name

BOOTP/DHCPv4 messages include "siaddr", "sname", and "file" fields. Even though DHCPv4 includes corresponding options, such as option 66 and option 67, some clients may not support these options. For this reason, server administrators often use the "siaddr", "sname", and "file" fields instead.

With Kea, it is possible to make static reservations for these DHCPv4 message fields:

```
{
  "subnet4": [
    {
      "reservations": [
        {
          "hw-address": "aa:bb:cc:dd:ee:ff",
          "next-server": "10.1.1.2",
          "server-hostname": "server-hostname.example.org",
          "boot-file-name": "/usr/local/share/kea/bootfile.efi"
        }
      ],
      ...
    }
  ],
  ...
}
```

Note that those parameters can be specified in combination with other parameters for a reservation, such as a reserved IPv4 address. These parameters are optional; a subset of them can be specified, or all of them can be omitted.

9.3.6 Reserving Client Classes in DHCPv4

Using Expressions in Classification explains how to configure the server to assign classes to a client, based on the content of the options that this client sends to the server. Host reservation mechanisms also allow for the static assignment of classes to clients. The definitions of these classes are placed in the Kea configuration file or a database. The following configuration snippet shows how to specify that a client belongs to the classes `reserved-class1` and `reserved-class2`. Those classes are associated with specific options sent to the clients which belong to them.

```
{
  "client-classes": [
    {
      "name": "reserved-class1",
      "option-data": [
        {
          "name": "routers",
          "data": "10.0.0.200"
        }
      ]
    },
    {
      "name": "reserved-class2",
      "option-data": [
        {
          "name": "domain-name-servers",
          "data": "10.0.0.201"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "subnet4": [
  {
    "id": 1,
    "subnet": "10.0.0.0/24",
    "pools": [ { "pool": "10.0.0.10-10.0.0.100" } ],
    "reservations": [
    {
      "hw-address": "aa:bb:cc:dd:ee:ff",

      "client-classes": [ "reserved-class1", "reserved-class2" ]

    }
    ]
  }
  ]
}

```

In some cases the host reservations can be used in conjunction with client classes specified within the Kea configuration. In particular, when a host reservation exists for a client within a given subnet, the "KNOWN" built-in class is assigned to the client. Conversely, when there is no static assignment for the client, the "UNKNOWN" class is assigned to the client. Class expressions within the Kea configuration file can refer to "KNOWN" or "UNKNOWN" classes using the "member" operator. For example:

```

{
  "client-classes": [
    {
      "name": "dependent-class",
      "test": "member('KNOWN')",
      "only-in-additional-list": true
    }
  ]
}

```

The `only-in-additional-list` parameter is needed here to force evaluation of the class after the lease has been allocated, and thus the reserved class has been also assigned.

Note

The classes specified in non-global host reservations are assigned to the processed packet after all classes with the `only-in-additional-list` parameter set to `false` have been evaluated. This means that these classes must not depend on the statically assigned classes from the host reservations. If such a dependency is needed, the `only-in-additional-list` parameter must be set to `true` for the dependent classes. Such classes are evaluated after the static classes have been assigned to the packet. This, however, imposes additional configuration overhead, because all classes marked as `only-in-additional-list` must be listed in the `evaluate-additional-classes` list for every subnet where they are used.

Note

Client classes specified within the Kea configuration file may depend on the classes specified within the global host

reservations. In such a case the `only-in-additional-list` parameter is not needed. Refer to [Pool Selection with Client Class Reservations](#) and [Subnet Selection with Client Class Reservations](#) for specific use cases.

9.3.7 Storing Host Reservations in MySQL or PostgreSQL

Kea can store host reservations in MySQL or PostgreSQL. See [Hosts Storage](#) for information on how to configure Kea to use reservations stored in MySQL or PostgreSQL. Kea provides a dedicated hook for managing reservations in a database; section [libdhcp_host_cmds.so: Host Commands](#) provides detailed information. The [Kea wiki](#) provides some examples of how to conduct common host reservation operations.

Note

In Kea, the maximum length of an option specified per-host-reservation is arbitrarily set to 4096 bytes.

9.3.8 Fine-Tuning DHCPv4 Host Reservation

The host reservation capability introduces additional restrictions for the allocation engine (the component of Kea that selects an address for a client) during lease selection and renewal. In particular, three major checks are necessary. First, when selecting a new lease, it is not sufficient for a candidate lease to simply not be in use by another DHCP client; it also must not be reserved for another client. Similarly, when renewing a lease, an additional check must be performed to see whether the address being renewed is reserved for another client. Finally, when a host renews an address, the server must check whether there is a reservation for this host, which would mean the existing (dynamically allocated) address should be revoked and the reserved one be used instead.

Some of those checks may be unnecessary in certain deployments, and not performing them may improve performance. The Kea server provides the `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool` configuration parameters to select the types of reservations allowed for a particular subnet. Each reservation type has different constraints for the checks to be performed by the server when allocating or renewing a lease for the client.

Configuration flags are:

- `reservations-in-subnet` - when set to `true`, it enables in-pool host reservation types. This setting is the default value, and is the safest and most flexible. However, as all checks are conducted, it is also the slowest. It does not check against global reservations. This flag defaults to `true`.
- `reservations-out-of-pool` - when set to `true`, it allows only out-of-pool host reservations. In this case the server assumes that all host reservations are for addresses that do not belong to the dynamic pool. Therefore, it can skip the reservation checks when dealing with in-pool addresses, thus improving performance. Do not use this mode if any reservations use in-pool addresses. Caution is advised when using this setting; Kea does not sanity-check the reservations against `reservations-out-of-pool` and misconfiguration may cause problems. This flag defaults to `false`.
- `reservations-global` - allows global host reservations. With this setting in place, the server searches for reservations for a client among the defined global reservations. If an address is specified, the server skips the reservation checks carried out in other modes, thus improving performance. Caution is advised when using this setting; Kea does not sanity-check the reservations when `reservations-global` is set to `true`, and misconfiguration may cause problems. This flag defaults to `false`.

Note: setting all flags to false disables host reservation support.

As there are no reservations, the server skips all checks. Any reservations defined are completely ignored. As checks are skipped, the server may operate faster in this mode.

Since Kea 1.9.1 the `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool` flags are supported.

(continued from previous page)

		in the respective	
		subnet?	
	+-----+		
	yes	no	
			V
+-----+		+--> "in-subnet"	

An example configuration that disables reservations looks as follows:

```
{
  "Dhcp4": {
    "subnet4": [
      {
        "id": 1,
        "pools": [
          {
            "pool": "192.0.2.10-192.0.2.100"
          }
        ],
        "reservations-global": false,
        "reservations-in-subnet": false,
        "subnet": "192.0.2.0/24"
      }
    ]
  }
}
```

An example configuration using global reservations is shown below:

```
{
  "Dhcp4": {
    "reservations-global": true,
    "reservations": [
      {
        "hostname": "host-one",
        "hw-address": "01:bb:cc:dd:ee:ff"
      },
      {
        "hostname": "host-two",
        "hw-address": "02:bb:cc:dd:ee:ff"
      }
    ],
    "subnet4": [
      {
        "id": 1,
        "pools": [
          {
            "pool": "192.0.2.10-192.0.2.100"
          }
        ],
        "subnet": "192.0.2.0/24"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

The meaning of the reservation flags are:

- **reservations-global**: fetch global reservations.
- **reservations-in-subnet**: fetch subnet reservations. For a shared network this includes all subnet members of the shared network.
- **reservations-out-of-pool**: this makes sense only when the **reservations-in-subnet** flag is **true**. When **reservations-out-of-pool** is **true**, the server assumes that all host reservations are for addresses that do not belong to the dynamic pool. Therefore, it can skip the reservation checks when dealing with in-pool addresses, thus improving performance. The server will not assign reserved addresses that are inside the dynamic pools to the respective clients. This also means that the addresses matching the respective reservations from inside the dynamic pools (if any) can be dynamically assigned to any client.

The disabled configuration corresponds to:

```

{
  "Dhcp4": {
    "reservations-global": false,
    "reservations-in-subnet": false
  }
}

```

The global configuration using `reservations-global` corresponds to:

```

{
  "Dhcp4": {
    "reservations-global": true,
    "reservations-in-subnet": false
  }
}

```

The out-of-pool configuration using `reservations-out-of-pool` corresponds to:

```

{
  "Dhcp4": {
    "reservations-global": false,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": true
  }
}

```

And the in-subnet configuration using `reservations-in-subnet` corresponds to:

```

{
  "Dhcp4": {
    "reservations-global": false,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": false
  }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

To activate both global and in-subnet, the following combination can be used:

```
{
  "Dhcp4": {
    "reservations-global": true,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": false
  }
}
```

To activate both global and out-of-pool, the following combination can be used:

```
{
  "Dhcp4": {
    "reservations-global": true,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": true
  }
}
```

Enabling out-of-pool and disabling in-subnet at the same time is not recommended because out-of-pool applies to host reservations in a subnet, which are fetched only when the in-subnet flag is true.

The parameter can be specified at the global, subnet, and shared-network levels.

An example configuration that disables reservations looks as follows:

```
{
  "Dhcp4": {
    "subnet4": [
      {
        "reservations-global": false,
        "reservations-in-subnet": false,
        "subnet": "192.0.2.0/24",
        "id": 1
      }
    ]
  }
}
```

An example configuration using global reservations is shown below:

```
{
  "Dhcp4": {
    "reservations": [
      {
        "hostname": "host-one",
        "hw-address": "01:bb:cc:dd:ee:ff"
      },
      {
        "hostname": "host-two",
```

(continues on next page)

(continued from previous page)

```

        "hw-address": "02:bb:cc:dd:ee:ff"
    }
],
"reservations-global": true,
"reservations-in-subnet": false,
"subnet4": [
    {
        "pools": [
            {
                "pool": "192.0.2.10-192.0.2.100"
            }
        ],
        "subnet": "192.0.2.0/24",
        "id": 1
    }
]
}
}

```

For more details regarding global reservations, see *Global Reservations in DHCPv4*.

Another aspect of host reservations is the different types of identifiers. Kea currently supports four types of identifiers: `hw-address`, `duid`, `client-id`, and `circuit-id`. This is beneficial from a usability perspective; however, there is one drawback. For each incoming packet, Kea has to extract each identifier type and then query the database to see if there is a reservation by this particular identifier. If nothing is found, the next identifier is extracted and the next query is issued. This process continues until either a reservation is found or all identifier types have been checked. Over time, with an increasing number of supported identifier types, Kea would become slower and slower.

To address this problem, a parameter called `host-reservation-identifiers` is available. It takes a list of identifier types as a parameter. Kea checks only those identifier types enumerated in `host-reservation-identifiers`. From a performance perspective, the number of identifier types should be kept to a minimum, ideally one. If the deployment uses several reservation types, please enumerate them from most- to least-frequently used, as this increases the chances of Kea finding the reservation using the fewest queries. An example of a `host-reservation-identifiers` configuration looks as follows:

```

{
"host-reservation-identifiers": [ "circuit-id", "hw-address", "duid", "client-id" ],
"subnet4": [
    {
        "subnet": "192.0.2.0/24",
        ...
    }
],
...
}

```

If not specified, the default value is:

```
"host-reservation-identifiers": [ "hw-address", "duid", "circuit-id", "client-id" ]
```

Note

As soon as a host reservation is found, the search is stopped; when a client has two host reservations using different

enabled identifier types, the first is always returned and the second ignored. In other words, this is usually a configuration error. In those rare cases when having two reservations for the same host makes sense, the one to be used can be specified by ordering the list of identifier types in `host-reservation-identifiers`.

9.3.9 Global Reservations in DHCPv4

In some deployments, such as mobile networks, clients can roam within the network and certain parameters must be specified regardless of the client's current location. To meet such a need, Kea offers a global reservation mechanism. The idea behind it is that regular host reservations are tied to specific subnets, by using a specific subnet ID. Kea can specify a global reservation that can be used in every subnet that has global reservations enabled.

This feature can be used to assign certain parameters, such as hostname or other dedicated, host-specific options. It can also be used to assign addresses.

An address assigned via global host reservation must be feasible for the subnet the server selects for the client. In other words, the address must lie within the subnet; otherwise, it is ignored and the server will attempt to dynamically allocate an address. If the selected subnet belongs to a shared network, the server checks for feasibility against the subnet's siblings, selecting the first in-range subnet. If no such subnet exists, the server falls back to dynamically allocating the address.

Note

Prior to release 2.3.5, the server did not perform feasibility checks on globally reserved addresses, which allowed the server to be configured to hand out nonsensical leases for arbitrary address values. Later versions of Kea perform these checks.

To use global host reservations, a configuration similar to the following can be used:

```
"Dhcp4": {
  # This specifies global reservations.
  # They will apply to all subnets that
  # have global reservations enabled.

  "reservations": [
    {
      "hw-address": "aa:bb:cc:dd:ee:ff",
      "hostname": "hw-host-dynamic"
    },
    {
      "hw-address": "01:02:03:04:05:06",
      "hostname": "hw-host-fixed",

      # Use of IP addresses in global reservations is risky.
      # If used outside of a matching subnet, such as 192.0.1.0/24,
      # it will result in a broken configuration being handed
      # to the client.
      "ip-address": "192.0.1.77"
    },
    {
      "duid": "01:02:03:04:05",
      "hostname": "duid-host"
    },
  ],
}
```

(continues on next page)

(continued from previous page)

```

{
  "circuit-id": "'charter950'",
  "hostname": "circuit-id-host"
},
{
  "client-id": "01:11:22:33:44:55:66",
  "hostname": "client-id-host"
}
],
"valid-lifetime": 600,
"subnet4": [
{
  "id": 1,
  "subnet": "10.0.0.0/24",
  # Specify if the server should look up global reservations.
  "reservations-global": true,
  # Specify if the server should look up in-subnet reservations.
  "reservations-in-subnet": false,
  # Specify if the server can assume that all reserved addresses
  # are out-of-pool. It can be ignored because "reservations-in-subnet"
  # is false.
  # "reservations-out-of-pool": false,
  "pools": [ { "pool": "10.0.0.10-10.0.0.100" } ]
}
]
}

```

When using database backends, the global host reservations are distinguished from regular reservations by using a `subnet-id` value of 0.

9.3.10 Pool Selection with Client Class Reservations

Client classes can be specified in the Kea configuration file and/or via host reservations. The classes specified in the Kea configuration file are evaluated immediately after receiving the DHCP packet and therefore can be used to influence subnet selection using the `client-classes` parameter specified in the subnet scope. The classes specified within the host reservations are fetched and assigned to the packet after the server has already selected a subnet for the client. This means that the client class specified within a host reservation cannot be used to influence subnet assignment for this client, unless the subnet belongs to a shared network. If the subnet belongs to a shared network, the server may dynamically change the subnet assignment while trying to allocate a lease. If the subnet does not belong to a shared network, the subnet is not changed once selected.

If the subnet does not belong to a shared network, it is possible to use host reservation-based client classification to select a pool within the subnet as follows:

```

{
  "Dhcp4": {
    "client-classes": [
      {
        "name": "reserved_class"
      },
      {
        "name": "unreserved_class",
        "test": "not member('reserved_class')"
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "reservations": [
        {
          "hw-address": "aa:bb:cc:dd:ee:fe",
          "client-classes": [ "unreserved_class" ]
        }
      ],
      "pools": [
        {
          "pool": "192.0.2.10 - 192.0.2.20",
          "client-classes": [ "unreserved_class" ]
        },
        {
          "pool": "192.0.2.30 - 192.0.2.40",
          "client-classes": [ "reserved_class" ]
        },
        {
          "pool": "192.0.2.50 - 192.0.2.60"
        }
      ]
    }
  ]
}

```

`reserved_class` is declared without the `test` parameter because it may only be assigned to a client via the host reservation mechanism. The second class, `unreserved_class`, is assigned to clients which do not belong to `reserved_class`.

The first pool in the subnet is used for clients not having such a reservation. The second pool is only used for clients having a reservation for `reserved_class`. The third pool is an unrestricted pool for any clients, comprising of both `reserved_class` clients and `unreserved_class`.

The configuration snippet includes one host reservation which causes the client with the MAC address `aa:bb:cc:dd:ee:fe` to be assigned to `reserved_class`. Thus, this client will be given an IP address from the second address pool.

Reservations defined on a subnet that belongs to a shared network are not visible to an otherwise matching client, so they cannot be used to select pools, nor subnets for that matter.

9.3.11 Subnet Selection with Client Class Reservations

There is one specific use case when subnet selection may be influenced by client classes specified within host reservations: when the client belongs to a shared network. In such a case it is possible to use classification to select a subnet within this shared network. Consider the following example:

```

{
  "Dhcp4": {
    "client-classes": [

```

(continues on next page)

(continued from previous page)

```

    {
      "name": "reserved_class"
    },
    {
      "name": "unreserved_class",
      "test": "not member('reserved_class')"
    }
  ],
  "reservations": [
    {
      "hw-address": "aa:bb:cc:dd:ee:fe",
      "client-classes": [ "reserved_class" ]
    }
  ],
  "reservations-global": true,
  "reservations-in-subnet": false,
  "shared-networks": [
    {
      "name": "net",
      "subnet4": [
        {
          "id": 1,
          "subnet": "192.0.2.0/24",
          "pools": [
            {
              "pool": "192.0.2.10-192.0.2.20",
              "client-classes": [ "unreserved_class" ]
            },
            {
              "pool": "192.0.2.30-192.0.2.40",
              "client-classes": [ "unreserved_class" ]
            }
          ]
        },
        {
          "id": 2,
          "subnet": "192.0.3.0/24",
          "pools": [
            {
              "pool": "192.0.3.10-192.0.3.20",
              "client-classes": [ "reserved_class" ]
            },
            {
              "pool": "192.0.3.30-192.0.3.40",
              "client-classes": [ "reserved_class" ]
            }
          ]
        },
        {
          "id": 3,
          "subnet": "192.0.4.0/24",
          "pools": [

```

(continues on next page)

(continued from previous page)

```

    {
      "pool": "192.0.4.10-192.0.4.20"
    },
    {
      "pool": "192.0.4.30-192.0.4.40"
    }
  ]
}
]
}
]
}
}

```

This is similar to the example described in *Pool Selection with Client Class Reservations*. This time, however, there are three subnets, of which the first two have a pool associated with a different class each.

The clients that do not have a reservation for `reserved_class` are assigned an address from the first subnet and when that is filled from the third subnet. Clients with a reservation for `reserved_class` are assigned an address from the second subnet and when that is filled from the third subnet.

The subnets must belong to the same shared network.

For a subnet to be restricted to a certain class, or skipped, all of the pools inside that subnet must be guarded by `reserved_class` or `unreserved_class` respectively.

In addition, the reservation for the client class must be specified at the global scope (global reservation) and `reservations-global` must be set to `true`.

In the example above, the `client-classes` configuration parameter could also be specified at the subnet level rather than the pool level, and would yield the same effect.

If the subnets were defined outside shared networks, and `client-classes` were specified at the subnet level, then `early-global-reservations-lookup` would also need to be enabled in order for subnet selection to work.

9.3.12 Multiple Reservations for the Same IP

Host reservations were designed to preclude the creation of multiple reservations for the same IP address within a particular subnet, to avoid having two different clients compete for the same address. When using the default settings, the server returns a configuration error when it finds two or more reservations for the same IP address within a subnet in the Kea configuration file. `libdhcp_host_cmds.so` returns an error in response to the `reservation-add` command when it detects that the reservation exists in the database for the IP address for which the new reservation is being added.

In some deployments a single host can select one of several network interfaces to communicate with the DHCP server, and the server must assign the same IP address to the host regardless of the interface used. Since each interface is assigned a different MAC address, it implies that several host reservations must be created to associate all of the MAC addresses present on this host with IP addresses. Using different IP addresses for each interface is impractical and is considered a waste of the IPv4 address space, especially since the host typically uses only one interface for communication with the server, hence only one IP address is in use.

This causes a need to create multiple host reservations for a single IP address within a subnet; this is supported since the Kea 1.9.1 release as an optional mode of operation, enabled with the `ip-reservations-unique` global parameter.

`ip-reservations-unique` is a boolean parameter that defaults to `true`, which forbids the specification of more than one reservation for the same IP address within a given subnet. Setting this parameter to `false` allows such reservations to be created both in the Kea configuration file and in the host database backend, via `libdhcp_host_cmds.so`.

Setting `ip-reservations-unique` to `false` when using `memfile`, `MySQL`, or `PostgreSQL` is supported. This setting is not supported when using `Host Cache` (see [libdhcp_host_cache.so: Host Cache Reservations for Improved Performance](#)) or the `RADIUS` backend (see [libdhcp_radius.so: RADIUS Server Support](#)). These reservation backends do not support multiple reservations for the same IP; if either of these hooks is loaded and `ip-reservations-unique` is set to `false`, then a configuration error is emitted and the server fails to start.

Note

When `ip-reservations-unique` is set to `true` (the default value), the server ensures that IP reservations are unique for a subnet within a single host backend and/or Kea configuration file. It does not guarantee that the reservations are unique across multiple backends. On server startup, only IP reservations defined in the Kea configuration file are checked for uniqueness.

The following is an example configuration with two reservations for the same IP address but different MAC addresses:

```
"Dhcp4": {
  "ip-reservations-unique": false,
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "reservations": [
        {
          "hw-address": "1a:1b:1c:1d:1e:1f",
          "ip-address": "192.0.2.11"
        },
        {
          "hw-address": "2a:2b:2c:2d:2e:2f",
          "ip-address": "192.0.2.11"
        }
      ]
    }
  ]
}
```

It is possible to control the `ip-reservations-unique` parameter via the [Configuration Backend in DHCPv4](#). If the new setting of this parameter conflicts with the currently used backends (i.e. backends do not support the new setting), the new setting is ignored and a warning log message is generated. The backends continue to use the default setting, expecting that IP reservations are unique within each subnet. To allow the creation of non-unique IP reservations, the administrator must remove the backends which lack support for them from the configuration file.

Administrators must be careful when they have been using multiple reservations for the same IP address and later decide to return to the default mode in which this is no longer allowed. They must make sure that at most one reservation for a given IP address exists within a subnet, prior to switching back to the default mode. If such duplicates are left in the configuration file, the server reports a configuration error. Leaving such reservations in the host databases does not cause configuration errors but may lead to lease allocation errors during the server's operation, when it unexpectedly finds multiple reservations for the same IP address.

Note

Currently, the Kea server does not verify whether multiple reservations for the same IP address exist in `MySQL` and/or `PostgreSQL` host databases when `ip-reservations-unique` is updated from `false` to `true`. This may cause issues with lease allocations. The administrator must ensure that there is at most one reservation for each IP

address within each subnet, prior to the configuration update.

The `reservations-lookup-first` is a boolean parameter which controls whether host reservations lookup should be performed before lease lookup. This parameter has effect only when multi-threading is disabled. When multi-threading is enabled, host reservations lookup is always performed first to avoid lease-lookup resource locking. The `reservations-lookup-first` parameter defaults to `false` when multi-threading is disabled.

9.3.13 Host Reservations as Basic Access Control

It is possible to define a host reservation that contains just an identifier, without any address, options, or values. In some deployments this is useful, as the hosts that have a reservation belong to the `KNOWN` class while others do not. This can be used as a basic access control mechanism.

The following example demonstrates this concept. It indicates a single IPv4 subnet and all clients will get an address from it. However, only known clients (those that have reservations) will get their default router configured. Empty reservations, i.e. reservations that only have the identification criterion, can be useful as a way of making the clients known.

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "KNOWN",
      "option-data": [
        {
          "name": "routers",
          "data": "192.0.2.250"
        }
      ]
    }
  ],
  "reservations": [
    // Clients on this list will be added to the KNOWN class.
    { "hw-address": "aa:bb:cc:dd:ee:fe" },
    { "hw-address": "11:22:33:44:55:66" }
  ],
  "reservations-in-subnet": true,

  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          "pool": "192.0.2.1-192.0.2.200"
        }
      ]
    }
  ]
}
```

This concept can be extended further. A good real-life scenario might be a situation where some customers of an ISP have not paid their bills. A new class can be defined to use an alternative default router that, instead of relaying traffic, redirects those customers to a captive portal urging them to bring their accounts up to date.

```

"Dhcp4": {
  "client-classes": [
    {
      "name": "blocked",
      "option-data": [
        {
          "name": "routers",
          "data": "192.0.2.251"
        }
      ]
    }
  ],
  "reservations": [
    // Clients on this list will be added to the KNOWN class. Some
    // will also be added to the blocked class.
    { "hw-address": "aa:bb:cc:dd:ee:fe",
      "client-classes": [ "blocked" ] },
    { "hw-address": "11:22:33:44:55:66" }
  ],
  "reservations-in-subnet": true,

  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          "pool": "192.0.2.1-192.0.2.200"
        }
      ],
      "option-data": [
        {
          "name": "routers",
          "data": "192.0.2.250"
        }
      ]
    }
  ]
}

```

9.4 Shared Networks in DHCPv4

DHCP servers use subnet information in two ways. It is used to both determine the point of attachment, i.e. where the client is connected to the network, and to group information pertaining to a specific location in the network. Sometimes it is useful to have more than one logical IP subnet deployed on the same physical link. Understanding that two or more subnets are used on the same link requires additional logic in the DHCP server. This capability is called "shared networks" in Kea, and sometimes also "shared subnets"; in Microsoft's nomenclature it is called "multinet."

There are many cases where the shared networks feature is useful; here we explain just a handful of the most common ones. The first and by far most common use case is an existing IPv4 network that has grown and is running out of available address space. Rather than migrating all devices to a new, larger subnet, it is easier to simply configure additional subnets on top of the existing one. Sometimes, due to address space fragmentation (e.g. only many disjointed /24s are available), this is the only choice. Also, configuring additional subnets has the advantage of not disrupting the

operation of existing devices.

Another very frequent use case comes from cable networks. There are two types of devices in cable networks: cable modems and the end-user devices behind them. It is a common practice to use different subnets for cable modems to prevent users from tinkering with them. In this case, the distinction is based on the type of device, rather than on address-space exhaustion.

A client connected to a shared network may be assigned an address from any of the pools defined within the subnets belonging to the shared network. Internally, the server selects one of the subnets belonging to a shared network and tries to allocate an address from this subnet. If the server is unable to allocate an address from the selected subnet (e.g., due to address-pool exhaustion), it uses another subnet from the same shared network and tries to allocate an address from this subnet. The server typically allocates all addresses available in a given subnet before it starts allocating addresses from other subnets belonging to the same shared network. However, in certain situations the client can be allocated an address from another subnet before the address pools in the first subnet get exhausted; this sometimes occurs when the client provides a hint that belongs to another subnet, or the client has reservations in a subnet other than the default.

Note

Deployments should not assume that Kea waits until it has allocated all the addresses from the first subnet in a shared network before allocating addresses from other subnets.

To define a shared network, an additional configuration scope is introduced:

```
{
  "Dhcp4": {
    "shared-networks": [
      {
        # Name of the shared network. It may be an arbitrary string
        # and it must be unique among all shared networks.
        "name": "my-secret-lair-level-1",

        # The subnet selector can be specified at the shared-network level.
        # Subnets from this shared network will be selected for directly
        # connected clients sending requests to the server's "eth0" interface.
        "interface": "eth0",

        # This starts a list of subnets in this shared network.
        # There are two subnets in this example.
        "subnet4": [
          {
            "id": 1,
            "subnet": "10.0.0.0/8",
            "pools": [ { "pool": "10.0.0.1 - 10.0.0.99" } ]
          },
          {
            "id": 2,
            "subnet": "192.0.2.0/24",
            "pools": [ { "pool": "192.0.2.100 - 192.0.2.199" } ]
          }
        ]
      }
    ],
    # end of shared-networks
  }
}
```

(continues on next page)

(continued from previous page)

```

# It is likely that in the network there will be a mix of regular,
# "plain" subnets and shared networks. It is perfectly valid to mix
# them in the same configuration file.
#
# This is a regular subnet. It is not part of any shared network.
"subnet4": [
    {
        "id": 3,
        "subnet": "192.0.3.0/24",
        "pools": [ { "pool": "192.0.3.1 - 192.0.3.200" } ],
        "interface": "eth1"
    }
]
}
}

```

As demonstrated in the example, it is possible to mix shared and regular ("plain") subnets. Each shared network must have a unique name. This is similar to the ID for subnets, but gives administrators more flexibility. It is used for logging, but also internally for identifying shared networks.

In principle it makes sense to define only shared networks that consist of two or more subnets. However, for testing purposes, an empty subnet or a network with just a single subnet is allowed. This is not a recommended practice in production networks, as the shared network logic requires additional processing and thus lowers the server's performance. To avoid unnecessary performance degradation, shared subnets should only be defined when required by the deployment.

Shared networks provide the ability to specify many parameters in the shared network scope that apply to all subnets within it. If necessary, it is possible to specify a parameter in the shared-network scope and then override its value in the subnet scope. For example:

```

{
  "shared-networks": [
    {
      "name": "lab-network3",

      "interface": "eth0",

      # This applies to all subnets in this shared network, unless
      # values are overridden on subnet scope.
      "valid-lifetime": 600,

      # This option is made available to all subnets in this shared
      # network.
      "option-data": [
        {
          "name": "log-servers",
          "data": "1.2.3.4"
        }
      ],

      "subnet4": [
        {

```

(continues on next page)

(continued from previous page)

```

        "id": 1,
        "subnet": "10.0.0.0/8",
        "pools": [ { "pool": "10.0.0.1 - 10.0.0.99" } ],

        # This particular subnet uses different values.
        "valid-lifetime": 1200,
        "option-data": [
            {
                "name": "log-servers",
                "data": "10.0.0.254"
            },
            {
                "name": "routers",
                "data": "10.0.0.254"
            }
        ]
    },
    {
        "id": 2,
        "subnet": "192.0.2.0/24",
        "pools": [ { "pool": "192.0.2.100 - 192.0.2.199" } ],

        # This subnet does not specify its own valid-lifetime value,
        # so it is inherited from shared network scope.
        "option-data": [
            {
                "name": "routers",
                "data": "192.0.2.1"
            }
        ]
    }
]
}
...
}

```

In this example, there is a `log-servers` option defined that is available to clients in both subnets in this shared network. Also, the valid lifetime is set to 10 minutes (600s). However, the first subnet overrides some of the values (the valid lifetime is 20 minutes, there is a different IP address for `log-servers`), but also adds its own option (the router address). Assuming a client asking for router and `log-servers` options is assigned a lease from this subnet, it will get a lease for 20 minutes and a `log-servers` and routers value of 10.0.0.254. If the same client is assigned to the second subnet, it will get a 10-minute lease, a `log-servers` value of 1.2.3.4, and routers set to 192.0.2.1.

9.4.1 Local and Relayed Traffic in Shared Networks

It is possible to specify an interface name at the shared-network level, to tell the server that this specific shared network is reachable directly (not via relays) using the local network interface. As all subnets in a shared network are expected to be used on the same physical link, it is a configuration error to attempt to define a shared network using subnets that are reachable over different interfaces. In other words, all subnets within the shared network must have the same value for the `interface` parameter. The following configuration is an example of what **NOT** to do:

```

{
  "shared-networks": [

```

(continues on next page)

(continued from previous page)

```

{
  "name": "office-floor-2",
  "subnet4": [
    {
      "id": 1,
      "subnet": "10.0.0.0/8",
      "pools": [ { "pool": "10.0.0.1 - 10.0.0.99" } ],
      "interface": "eth0"
    },
    {
      "id": 2,
      "subnet": "192.0.2.0/24",
      "pools": [ { "pool": "192.0.2.100 - 192.0.2.199" } ],

      # Specifying the different interface name is a configuration
      # error. This value should rather be "eth0" or the interface
      # name in the other subnet should be "eth1".
      "interface": "eth1"
    }
  ]
}
],
...
}

```

To minimize the chance of configuration errors, it is often more convenient to simply specify the interface name once, at the shared-network level, as shown in the example below.

```

{
  "shared-networks": [
    {
      "name": "office-floor-2",

      # This tells Kea that the whole shared network is reachable over a
      # local interface. This applies to all subnets in this network.
      "interface": "eth0",

      "subnet4": [
        {
          "id": 1,
          "subnet": "10.0.0.0/8",
          "pools": [ { "pool": "10.0.0.1 - 10.0.0.99" } ]
        },
        {
          "id": 2,
          "subnet": "192.0.2.0/24",
          "pools": [ { "pool": "192.0.2.100 - 192.0.2.199" } ]
        }
      ]
    }
  ]
},
...

```

(continues on next page)

(continued from previous page)

}

With relayed traffic, subnets are typically selected using the relay agents' addresses. If the subnets are used independently (not grouped within a shared network), a different relay address can be specified for each of these subnets. When multiple subnets belong to a shared network they must be selected via the same relay address and, similarly to the case of the local traffic described above, it is a configuration error to specify different relay addresses for the respective subnets in the shared network. The following configuration is another example of what **NOT** to do:

```
{
"shared-networks": [
  {
    "name": "kakapo",
    "subnet4": [
      {
        "id": 1,
        "subnet": "192.0.2.0/26",
        "relay": {
          "ip-addresses": [ "192.1.1.1" ]
        },
        "pools": [ { "pool": "192.0.2.63 - 192.0.2.63" } ]
      },
      {
        "id": 2,
        "subnet": "10.0.0.0/24",
        "relay": {
          # Specifying a different relay address for this
          # subnet is a configuration error. In this case
          # it should be 192.1.1.1 or the relay address
          # in the previous subnet should be 192.2.2.2.
          "ip-addresses": [ "192.2.2.2" ]
        },
        "pools": [ { "pool": "10.0.0.16 - 10.0.0.16" } ]
      }
    ]
  }
],
...
}
```

Again, it is better to specify the relay address at the shared-network level; this value will be inherited by all subnets belonging to the shared network.

```
{
"shared-networks": [
  {
    "name": "kakapo",
    "relay": {
      # This relay address is inherited by both subnets.
      "ip-addresses": [ "192.1.1.1" ]
    },
    "subnet4": [
      {
```

(continues on next page)

(continued from previous page)

```

        "id": 1,
        "subnet": "192.0.2.0/26",
        "pools": [ { "pool": "192.0.2.63 - 192.0.2.63" } ]
    },
    {
        "id": 2,
        "subnet": "10.0.0.0/24",
        "pools": [ { "pool": "10.0.0.16 - 10.0.0.16" } ]
    }
]
}
],
...
}

```

Even though it is technically possible to configure two (or more) subnets within the shared network to use different relay addresses, this will almost always lead to a different behavior than what the user expects. In this case, the Kea server will initially select one of the subnets by matching the relay address in the client's packet with the subnet's configuration. However, it MAY end up using the other subnet (even though it does not match the relay address) if the client already has a lease in this subnet or has a host reservation in this subnet, or simply if the initially selected subnet has no more addresses available. Therefore, it is strongly recommended to always specify subnet selectors (interface or relay address) at the shared-network level if the subnets belong to a shared network, as it is rarely useful to specify them at the subnet level and may lead to the configuration errors described above.

9.4.2 Client Classification in Shared Networks

Sometimes it is desirable to segregate clients into specific subnets based on certain properties. This mechanism is called client classification and is described in *Client Classification*.

Client classification can be applied to subnets belonging to shared networks in the same way as it is used for subnets specified outside of shared networks. It is important to understand how the server selects subnets for clients when client classification is in use, to ensure that the appropriate subnet is selected for a given client type.

If a subnet is associated with one or more classes, only the clients belonging to at least one of these classes may use this subnet. If there are no classes specified for a subnet, any client connected to a given shared network can use this subnet. A common mistake is to assume that a subnet that includes a client class is preferred over subnets without client classes.

The `client-classes` parameter may be specified at the shared network, subnet, and/or pool scopes. If specified for a shared network, clients must belong to at least one of the classes specified for that network to be considered for subnets within that network. If specified for a subnet, clients must belong to at least one of the classes specified for that subnet to be considered for any of that subnet's pools or host reservations. If specified for a pool, clients must belong to at least one of the classes specified for that pool to be given a lease from that pool.

Consider the following example:

```

{
    "client-classes": [
        {
            "name": "b-devices",
            "test": "option[93].hex == 0x0002"
        }
    ],
    "shared-networks": [

```

(continues on next page)

(continued from previous page)

```

{
  "name": "galah",
  "interface": "eth0",
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/26",
      "pools": [ { "pool": "192.0.2.1 - 192.0.2.63" } ]
    },
    {
      "id": 2,
      "subnet": "10.0.0.0/24",
      "pools": [ { "pool": "10.0.0.2 - 10.0.0.250" } ],
      "client-classes": [ "b-devices" ]
    }
  ]
}

```

If the client belongs to the "b-devices" class (because it includes option 93 with a value of 0x0002), that does not guarantee that the subnet 10.0.0.0/24 will be used (or preferred) for this client. The server can use either of the two subnets, because the subnet 192.0.2.0/26 is also allowed for this client. The client classification used in this case should be perceived as a way to restrict access to certain subnets, rather than as a way to express subnet preference. For example, if the client does not belong to the "b-devices" class, it may only use the subnet 192.0.2.0/26 and will never use the subnet 10.0.0.0/24.

A typical use case for client classification is in a cable network, where cable modems should use one subnet and other devices should use another subnet within the same shared network. In this case it is necessary to apply classification on all subnets. The following example defines two classes of devices, and the subnet selection is made based on option 93 values.

```

{
  "client-classes": [
    {
      "name": "a-devices",
      "test": "option[93].hex == 0x0001"
    },
    {
      "name": "b-devices",
      "test": "option[93].hex == 0x0002"
    }
  ],
  "shared-networks": [
    {
      "name": "galah",
      "interface": "eth0",
      "subnet4": [
        {
          "id": 1,
          "subnet": "192.0.2.0/26",
          "pools": [ { "pool": "192.0.2.1 - 192.0.2.63" } ],

```

(continues on next page)

(continued from previous page)

```

        "client-classes": [ "a-devices" ]
    },
    {
        "id": 2,
        "subnet": "10.0.0.0/24",
        "pools": [ { "pool": "10.0.0.2 - 10.0.0.250" } ],
        "client-classes": [ "b-devices" ]
    }
]
}
]
}

```

In this example each class has its own restriction. Only clients that belong to class "a-devices" are able to use subnet 192.0.2.0/26 and only clients belonging to "b-devices" are able to use subnet 10.0.0.0/24. Care should be taken not to define too-restrictive classification rules, as clients that are unable to use any subnets will be refused service. However, this may be a desired outcome if one wishes to provide service only to clients with known properties (e.g. only VoIP phones allowed on a given link).

It is possible to achieve an effect similar to the one presented in this section without the use of shared networks. If the subnets are placed in the global subnets scope, rather than in the shared network, the server will still use classification rules to pick the right subnet for a given class of devices. The major benefit of placing subnets within the shared network is that common parameters for the logically grouped subnets can be specified once in the shared-network scope, e.g. the `interface` or `relay` parameter. All subnets belonging to this shared network will inherit those parameters.

9.4.3 Host Reservations in Shared Networks

Subnets that are part of a shared network allow host reservations, similar to regular subnets:

```

{
  "shared-networks": [
    {
      "name": "frog",
      "interface": "eth0",
      "subnet4": [
        {
          "subnet": "192.0.2.0/26",
          "id": 100,
          "pools": [ { "pool": "192.0.2.1 - 192.0.2.63" } ],
          "reservations": [
            {
              "hw-address": "aa:bb:cc:dd:ee:ff",
              "ip-address": "192.0.2.28"
            }
          ]
        }
      ],
    },
    {
      "subnet": "10.0.0.0/24",
      "id": 101,
      "pools": [ { "pool": "10.0.0.1 - 10.0.0.254" } ],
      "reservations": [
        {
          "hw-address": "11:22:33:44:55:66",

```

(continues on next page)

(continued from previous page)

```

        "ip-address": "10.0.0.29"
      }
    ]
  }
]
}

```

It is worth noting that Kea conducts additional checks when processing a packet if shared networks are defined. First, instead of simply checking whether there is a reservation for a given client in its initially selected subnet, Kea looks through all subnets in a shared network for a reservation. This is one of the reasons why defining a shared network may impact performance. If there is a reservation for a client in any subnet, that particular subnet is selected for the client. Although it is technically not an error, it is considered bad practice to define reservations for the same host in multiple subnets belonging to the same shared network.

While not strictly mandatory, it is strongly recommended to use explicit "id" values for subnets if database storage will be used for host reservations. If an ID is not specified, the values for it are auto-generated, i.e. Kea assigns increasing integer values starting from 1. Thus, the auto-generated IDs are not stable across configuration changes.

9.5 Server Identifier in DHCPv4

The DHCPv4 protocol uses a "server identifier" to allow clients to discriminate between several servers present on the same link; this value is an IPv4 address of the server. The server chooses the IPv4 address of the interface on which the message from the client (or relay) has been received. A single server instance uses multiple server identifiers if it is receiving queries on multiple interfaces.

It is possible to override the default server identifier values by specifying the `dhcp-server-identifier` option. This option configuration is only supported at the subnet, shared network, client class, and global levels. It must not be specified at the host-reservation level. When configuring the `dhcp-server-identifier` option at client-class level, the class must not set the `only-in-additional-list` flag, because this class would not be evaluated before the server determines if the received DHCP message should be accepted for processing. Such classes are evaluated after subnet selection. See [Additional Classification](#) for details.

The following example demonstrates how to override the server identifier for a subnet:

```

{
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "option-data": [
        {
          "name": "dhcp-server-identifier",
          "data": "10.2.5.76"
        }
      ],
      ...
    }
  ],
  ...
}

```


9.6 How the DHCPv4 Server Selects a Subnet for the Client

The DHCPv4 server differentiates among directly connected clients, clients trying to renew leases, and clients sending their messages through relays. For directly connected clients, the server checks the configuration for the interface on which the message has been received and, if the server configuration does not match any configured subnet, the message is discarded.

An optional interface parameter is available within a subnet definition to designate that a given subnet is local, i.e. reachable directly over the specified interface. For example, a server that is intended to serve a local subnet over eth0 may be configured as follows:

```
"Dhcp4": {
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          "pool": "192.0.2.100 - 192.0.2.199"
        }
      ],
      "interface": "eth0"
    }
  ],
  ...
}
```

Assuming that the server's interface is configured with the IPv4 address 192.0.2.3, the server only processes messages received through this interface from a directly connected client if there is a subnet configured to which this IPv4 address belongs, such as 192.0.2.0/24. The server uses this subnet to assign an IPv4 address for the client.

The rule above does not apply when the client unicasts its message, i.e. is trying to renew its lease; such a message is accepted through any interface. The renewing client sets `ciaddr` to the currently used IPv4 address, and the server uses this address to select the subnet for the client (in particular, to extend the lease using this address).

If the message is relayed it is accepted through any interface. The `giaddr` set by the relay agent is used to select the subnet for the client.

It is also possible to specify a relay IPv4 address for a given subnet. It can be used to match incoming packets into a subnet in uncommon configurations, e.g. shared networks. See *Using a Specific Relay Agent for a Subnet* for details.

Note

The subnet selection mechanism described in this section is based on the assumption that client classification is not used. The classification mechanism alters the way in which a subnet is selected for the client, depending on the classes to which the client belongs.

Note

When the selected subnet is a member of a shared network, the whole shared network is selected.

9.6.1 Using a Specific Relay Agent for a Subnet

A relay must have an interface connected to the link on which the clients are being configured. Typically the relay has an IPv4 address configured on that interface, which belongs to the subnet from which the server assigns addresses. Normally, the server is able to use the IPv4 address inserted by the relay (in the `giaddr` field of the DHCPv4 packet) to select the appropriate subnet.

However, that is not always the case. In certain uncommon — but valid — deployments, the relay address may not match the subnet. This usually means that there is more than one subnet allocated for a given link. The two most common examples of this are long-lasting network renumbering (where both old and new address spaces are still being used) and a cable network. In a cable network, both cable modems and the devices behind them are physically connected to the same link, yet they use distinct addressing. In such a case, the DHCPv4 server needs additional information (the IPv4 address of the relay) to properly select an appropriate subnet.

The following example assumes that there is a subnet 192.0.2.0/24 that is accessible via a relay that uses 10.0.0.1 as its IPv4 address. The server is able to select this subnet for any incoming packets that come from a relay that has an address in the 192.0.2.0/24 subnet. It also selects that subnet for a relay with address 10.0.0.1.

```
{
  "Dhcp4": {
    "subnet4": [
      {
        "id": 1,
        "subnet": "192.0.2.0/24",
        "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
        "relay": {
          "ip-addresses": [ "10.0.0.1" ]
        }
      }
    ]
  }
}
```

If `relay` is specified, the `ip-addresses` parameter within it is mandatory. The `ip-addresses` parameter supports specifying a list of addresses.

9.6.2 Segregating IPv4 Clients in a Cable Network

In certain cases, it is useful to mix relay address information (introduced in *Using a Specific Relay Agent for a Subnet*) with client classification (explained in *Client Classification*). One specific example is in a cable network, where modems typically get addresses from a different subnet than all the devices connected behind them.

Let us assume that there is one Cable Modem Termination System (CMTS) with one CM MAC (a physical link that modems are connected to). We want the modems to get addresses from the 10.1.1.0/24 subnet, while everything connected behind the modems should get addresses from the 192.0.2.0/24 subnet. The CMTS that acts as a relay uses address 10.1.1.1. The following configuration can serve that situation:

```
"Dhcp4": {
  "subnet4": [
    {
      "id": 1,
      "subnet": "10.1.1.0/24",
      "pools": [ { "pool": "10.1.1.2 - 10.1.1.20" } ],
      "client-classes": [ "docsis3.0" ],
      "relay": {
        "ip-addresses": [ "10.1.1.1" ]
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    }
  },
  {
    "id": 2,
    "subnet": "192.0.2.0/24",
    "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
    "relay": {
      "ip-addresses": [ "10.1.1.1" ]
    }
  }
],
...
}

```

9.7 Duplicate Addresses (DHCPDECLINE Support)

The DHCPv4 server is configured with a certain pool of addresses that it is expected to hand out to DHCPv4 clients. It is assumed that the server is authoritative and has complete jurisdiction over those addresses. However, for various reasons such as misconfiguration or a faulty client implementation that retains its address beyond the valid lifetime, there may be devices connected that use those addresses without the server's approval or knowledge.

Such an unwelcome event can be detected by legitimate clients (using ARP or ICMP Echo Request mechanisms) and reported to the DHCPv4 server using a DHCPDECLINE message. The server does a sanity check (to see whether the client declining an address really was supposed to use it) and then conducts a clean-up operation. Any DNS entries related to that address are removed, the event is logged, and hooks are triggered. After that is complete, the address is marked as declined (which indicates that it is used by an unknown entity and thus not available for assignment) and a probation time is set on it. Unless otherwise configured, the probation period lasts 24 hours; after that time, the server will recover the lease (i.e. put it back into the available state) and the address will be available for assignment again. It should be noted that if the underlying issue of a misconfigured device is not resolved, the duplicate-address scenario will repeat. If reconfigured correctly, this mechanism provides an opportunity to recover from such an event automatically, without any system administrator intervention.

To configure the decline probation period to a value other than the default, the following syntax can be used:

```

"Dhcp4": {
  "decline-probation-period": 3600,
  "subnet4": [
    {
      ...
    },
    ...
  ],
  ...
}

```

The parameter is expressed in seconds, so the example above instructs the server to recycle declined leases after one hour.

There are several statistics and hook points associated with the decline handling procedure. The `lease4_decline` hook point is triggered after the incoming DHCPDECLINE message has been sanitized and the server is about to decline the lease. The `declined-addresses` statistic is increased after the hook returns (both the global and subnet-specific variants). (See *Statistics in the DHCPv4 Server* and *Hook Libraries* for more details on DHCPv4 statistics and Kea hook points.)

Once the probation time elapses, the declined lease is recovered using the standard expired-lease reclamation procedure, with several additional steps. In particular, both `declined-addresses` statistics (global and subnet-specific) are decreased. At the same time, `reclaimed-declined-addresses` statistics (again in two variants, global and subnet-specific) are increased.

A note about statistics: The Kea server does not decrease the `assigned-addresses` statistics when a DHCPDECLINE is received and processed successfully. While technically a declined address is no longer assigned, the primary usage of the `assigned-addresses` statistic is to monitor pool utilization. Most people would forget to include `declined-addresses` in the calculation, and would simply use `assigned-addresses/total-addresses`. This would cause a bias towards under-representing pool utilization. As this has a potential to cause serious confusion, ISC decided not to decrease `assigned-addresses` immediately after receiving DHCPDECLINE, but to do it later when Kea recovers the address back to the available pool.

9.8 Statistics in the DHCPv4 Server

The DHCPv4 server supports the following statistics:

Table 8: DHCPv4 statistics

Statistic	Data Type	Description
<code>pkt4-received</code>	integer	Number of DHCPv4 packets received. This includes all packets: valid, bogus, corrupted, rejected, etc. This statistic is expected to grow rapidly.
<code>pkt4-discover-received</code>	integer	Number of DHCPDISCOVER packets received. This statistic is expected to grow; its increase means that clients that just booted started their configuration process and their initial packets reached the Kea server.
<code>pkt4-offer-received</code>	integer	Number of DHCPOFFER packets received. This statistic is expected to remain zero at all times, as DHCPOFFER packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards DHCPOFFER messages towards the server, rather than back to the clients.
<code>pkt4-request-received</code>	integer	Number of DHCPREQUEST packets received. This statistic is expected to grow. Its increase means that clients that just booted received the server's response (DHCPOFFER) and accepted it, and are now requesting an address (DHCPREQUEST).
<code>pkt4-ack-received</code>	integer	Number of DHCPACK packets received. This statistic is expected to remain zero at all times, as DHCPACK packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards DHCPACK messages towards the server, rather than back to the clients.
<code>pkt4-nak-received</code>	integer	Number of DHCPNAK packets received. This statistic is expected to remain zero at all times, as DHCPNAK packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards DHCPNAK messages towards the server, rather than back to the clients.
<code>pkt4-release-received</code>	integer	Number of DHCPRELEASE packets received. This statistic is expected to grow. Its increase means that clients that had an address are shutting down or ceasing to use their addresses.
<code>pkt4-decline-received</code>	integer	Number of DHCPDECLINE packets received. This statistic is expected to remain close to zero. Its increase means that a client leased an address, but discovered that the address is currently used by an unknown device elsewhere in the network.

continues on next page

Table 8 – continued from previous page

Statistic	Data Type	Description
pkt4-inform-received	integer	Number of DHCPINFORM packets received. This statistic is expected to grow. Its increase means that there are clients that either do not need an address or already have an address and are interested only in getting additional configuration parameters.
pkt4-unknown-received	integer	Number of packets received of an unknown type. A non-zero value of this statistic indicates that the server received a packet that it was not able to recognize, either with an unsupported type or possibly malformed (without a message-type option).
pkt4-sent	integer	Number of DHCPv4 packets sent. This statistic is expected to grow every time the server transmits a packet. In general, it should roughly match <code>pkt4-received</code> , as most incoming packets cause the server to respond. There are exceptions (e.g. DHCPRELEASE), so do not worry if it is less than <code>pkt4-received</code> .
pkt4-offer-sent	integer	Number of DHCPOFFER packets sent. This statistic is expected to grow in most cases after a DHCPDISCOVER is processed. There are certain uncommon but valid cases where incoming DHCPDISCOVER packets are dropped, but in general this statistic is expected to be close to <code>pkt4-discover-received</code> .
pkt4-ack-sent	integer	Number of DHCPACK packets sent. This statistic is expected to grow in most cases after a DHCPREQUEST is processed; there are certain cases where DHCPNAK is sent instead. In general, the sum of <code>pkt4-ack-sent</code> and <code>pkt4-nak-sent</code> should be close to <code>pkt4-request-received</code> .
pkt4-nak-sent	integer	Number of DHCPNAK packets sent. This statistic is expected to grow when the server chooses not to honor the address requested by a client. In general, the sum of <code>pkt4-ack-sent</code> and <code>pkt4-nak-sent</code> should be close to <code>pkt4-request-received</code> .
pkt4-parse-failed	integer	Number of incoming packets that could not be parsed. A non-zero value of this statistic indicates that the server received a malformed or truncated packet. This may indicate problems in the network, faulty clients, or a bug in the server.
pkt4-receive-drop	integer	Number of incoming packets that were dropped. The exact reason for dropping packets is logged, but the most common reasons may be that an unacceptable packet type was received, direct responses are forbidden, or the server ID sent by the client does not match the server's server ID.
subnet[id].total-addresses	integer	Total number of addresses available for DHCPv4 management for a given subnet; in other words, this is the count of all addresses in all configured pools. This statistic changes only during configuration updates. It does not take into account any addresses that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pool[pid].total-addresses	integer	Total number of addresses available for DHCPv4 management for a given subnet pool; in other words, this is the count of all addresses in configured subnet pool. This statistic changes only during configuration updates. It does not take into account any addresses that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of a given pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
cumulative-assigned-addresses	integer	Cumulative number of addresses that have been assigned since server startup. It is incremented each time an address is assigned and is not reset when the server is reconfigured.
subnet[id].cumulative-assigned-addresses	integer	Cumulative number of assigned addresses in a given subnet. It increases every time a new lease is allocated (as a result of receiving a DHCPREQUEST message) and never decreases. The <i>id</i> is the subnet ID of the subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.

continues on next page

Table 8 – continued from previous page

Statistic	Data Type	Description
subnet[id].pool[pid].cumulative-assigned-addresses	integer	Cumulative number of assigned addresses in a given subnet pool. It increases every time a new lease is allocated (as a result of receiving a DHCPREQUEST message) and never decreases. The <i>id</i> is the subnet ID of the subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
subnet[id].assigned-addresses	integer	Number of assigned addresses in a given subnet. It increases every time a new lease is allocated (as a result of receiving a DHCPREQUEST message) and decreases every time a lease is released (a DHCPRELEASE message is received) or expires. The <i>id</i> is the subnet ID of the subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pool[pid].assigned-addresses	integer	Number of assigned addresses in a given subnet pool. It increases every time a new lease is allocated (as a result of receiving a DHCPREQUEST message) and decreases every time a lease is released (a DHCPRELEASE message is received) or expires. The <i>id</i> is the subnet ID of the subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
reclaimed-leases	integer	Number of expired leases that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed and never decreases. It can be used as a long-term indicator of how many actual leases have been reclaimed. This is a global statistic that covers all subnets.
subnet[id].reclaimed-leases	integer	Number of expired leases associated with a given subnet that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
subnet[id].pool[pid].reclaimed-leases	integer	Number of expired leases associated with a given subnet pool that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately.
declined-addresses	integer	Number of IPv4 addresses that are currently declined; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. This is a global statistic that covers all subnets.
subnet[id].declined-addresses	integer	Number of IPv4 addresses that are currently declined in a given subnet; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
subnet[id].pool[pid].declined-addresses	integer	Number of IPv4 addresses that are currently declined in a given subnet pool; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately.
reclaimed-declined-addresses	integer	Number of IPv4 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. This is a global statistic that covers all subnets.

continues on next page

Table 8 – continued from previous page

Statistic	Data Type	Description
subnet[id].reclaimed-declined-addresses	integer	Number of IPv4 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
subnet[id].pool[pid].reclaimed-declined-addresses	integer	Number of IPv4 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately.
pkt4-lease-query-received	integer	Number of IPv4 DHCPLEASEQUERY packets received. (Only exists if the Leasequery hook library is loaded.)
pkt4-lease-query-response-unknown-sent	integer	Number of IPv4 DHCPLEASEUNKNOWN responses sent. (Only exists if the Leasequery hook library is loaded.)
pkt4-lease-query-response-unassigned-sent	integer	Number of IPv4 DHCPLEASEUNASSIGNED responses sent. (Only exists if the Leasequery hook library is loaded.)
pkt4-lease-query-response-active-sent	integer	Number of IPv4 DHCPLEASEACTIVE responses sent. (Only exists if the Leasequery hook library is loaded.)
v4-allocation-fail	integer	Number of total address allocation failures for a particular client. This consists of the number of lease allocation attempts that the server made before giving up, if it was unable to use any of the address pools. This is a global statistic that covers all subnets.
subnet[id].v4-allocation-fail	integer	Number of total address allocation failures for a particular client. This consists of the number of lease allocation attempts that the server made before giving up, if it was unable to use any of the address pools. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v4-allocation-fail-shared-network	integer	Number of address allocation failures for a particular client connected to a shared network. This is a global statistic that covers all subnets.
subnet[id].v4-allocation-fail-shared-network	integer	Number of address allocation failures for a particular client connected to a shared network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v4-allocation-fail-subnet	integer	Number of address allocation failures for a particular client connected to a subnet that does not belong to a shared network. This is a global statistic that covers all subnets.
subnet[id].v4-allocation-fail-subnet	integer	Number of address allocation failures for a particular client connected to a subnet that does not belong to a shared network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v4-allocation-fail-no-pools	integer	Number of address allocation failures because the server could not use any configured pools for a particular client. It is also possible that all of the subnets from which the server attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects that some clients should be assigned dynamic addresses. This is a global statistic that covers all subnets.
subnet[id].v4-allocation-fail-no-pools	integer	Number of address allocation failures because the server could not use any configured pools for a particular client. It is also possible that all of the subnets from which the server attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects that some clients should be assigned dynamic addresses. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.

continues on next page

Table 8 – continued from previous page

Statistic	Data Type	Description
v4-allocation-fail-classes	integer	Number of address allocation failures when the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned: for example, if all pools require packets to belong to certain classes and an incoming packet does not belong to any. Another case where this information may be useful is to indicate that the pool reserved for a given class has run out of addresses. This is a global statistic that covers all subnets.
subnet[id].v4-allocation-fail-classes	integer	Number of address allocation failures when the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned: for example, if all pools require packets to belong to certain classes and an incoming packet does not belong to any. Another case where this information may be useful is to indicate that the pool reserved for a given class has run out of addresses. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v4-lease-reuses	integer	Number of times an IPv4 lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is global.
subnet[id].v4-lease-reuses	integer	Number of times an IPv4 lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is on a per-subnet basis. The <i>id</i> is the subnet ID of a given subnet.
v4-reservation-conflicts	integer	Number of host reservation allocation conflicts which have occurred across every subnet. When a client sends a DHCP Discover and is matched to a host reservation which is already leased to another client, this counter is increased by 1.
subnet[id].v4-reservation-conflicts	integer	Number of host reservation allocation conflicts which have occurred in a specific subnet. When a client sends a DHCP Discover and is matched to a host reservation which is already leased to another client, this counter is increased by 1.

Note

The pool ID can be configured on each pool by explicitly setting the `pool-id` parameter in the pool parameter map. If not configured, `pool-id` defaults to 0. The statistics related to pool ID 0 refer to all the statistics of all the pools that have an unconfigured `pool-id`. The pool ID does not need to be unique within the subnet or across subnets. The statistics regarding a specific pool ID within a subnet are combined with the other statistics of all other pools with the same pool ID in the respective subnet.

Note

This section describes DHCPv4-specific statistics. For a general overview and usage of statistics, see [Statistics](#).

The DHCPv4 server provides two global parameters to control the default sample limits of statistics:

- `statistic-default-sample-count` - determines the default maximum number of samples to be kept. The special value of 0 indicates that a default maximum age should be used.
- `statistic-default-sample-age` - determines the default maximum age, in seconds, of samples to be kept.

For instance, to reduce the statistic-keeping overhead, set the default maximum sample count to 1 so only one sample is kept:


```
"Dhcp4": {
  "statistic-default-sample-count": 1,
  "subnet4": [
    {
      ...
    },
    ...
  ],
  ...
}
```

Statistics can be retrieved periodically to gain more insight into Kea operations. One tool that leverages that capability is ISC Stork. See *Monitoring Kea With Stork* for details.

9.9 Management API for the DHCPv4 Server

The management API allows the issuing of specific management commands, such as statistics retrieval, reconfiguration, or shutdown. For more details, see *Management API*. By default there are no sockets open; to instruct Kea to open a socket, the following entry in the configuration file can be used:

```
"Dhcp4": {
  "control-sockets": [
    {
      "socket-type": "unix",
      "socket-name": "/path/to/the/unix/socket"
    }
  ],

  "subnet4": [
    {
      ...
    },
    ...
  ],
  ...
}
```

9.9.1 UNIX Control Socket

Until Kea server 2.7.2 the only supported communication channel type was the UNIX stream socket with `socket-type` set to `unix` and `socket-name` to the file path of the UNIX/LOCAL socket.

The length of the path specified by the `socket-name` parameter is restricted by the maximum length for the UNIX socket name on the administrator's operating system, i.e. the size of the `sun_path` field in the `sockaddr_un` structure, decreased by 1. This value varies on different operating systems, between 91 and 107 characters. Typical values are 107 on Linux and 103 on FreeBSD.

Kea supports only one `unix` control socket in the "control-sockets" list.

Note

As of Kea 2.7.9, control sockets may only reside in the directory determined during compilation as "[`kea-install-dir`]/var/run/kea", which must also have `0750` access rights. This path may be overrid-

den at startup by setting the environment variable `KEA_CONTROL_SOCKET_DIR` to the desired path. If a path other than this value is used in `socket-name`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in simply omit the path component from `socket-name`.

Communication over the control channel is conducted using JSON structures. See the [Control Channel](#) section in the [Kea Developer's Guide](#) for more details.

The DHCPv4 server supports the following operational commands:

- `build-report`
- `config-get`
- `config-hash-get`
- `config-reload`
- `config-set`
- `config-test`
- `config-write`
- `dhcp-disable`
- `dhcp-enable`
- `leases-reclaim`
- `list-commands`
- `shutdown`
- `status-get`
- `version-get`

as described in *Commands Supported by Both the DHCPv4 and DHCPv6 Servers*. In addition, it supports the following statistics-related commands:

- `statistic-get`
- `statistic-reset`
- `statistic-remove`
- `statistic-get-all`
- `statistic-reset-all`
- `statistic-remove-all`
- `statistic-sample-age-set`
- `statistic-sample-age-set-all`
- `statistic-sample-count-set`
- `statistic-sample-count-set-all`

as described in *Commands for Manipulating Statistics*.

9.9.2 HTTP/HTTPS Control Socket

The `socket-type` must be `http` or `https` (when the type is `https` TLS is required). The `socket-address` (default 127.0.0.1) and `socket-port` (default 8000) specify an IP address and port to which the HTTP service will be bound.

Since Kea 2.7.5 the `http-headers` parameter specifies a list of extra HTTP headers to add to HTTP responses.

The `trust-anchor`, `cert-file`, `key-file`, and `cert-required` parameters specify the TLS setup for HTTP, i.e. HTTPS. If these parameters are not specified, HTTP is used. The TLS/HTTPS support in Kea is described in [TLS/HTTPS Support](#).

Basic HTTP authentication protects against unauthorized uses of the control agent by local users. For protection against remote attackers, HTTPS and reverse proxy of [Secure Connections](#) provide stronger security.

The authentication is described in the `authentication` block with the mandatory `type` parameter, which selects the authentication. Currently only the basic HTTP authentication (type `basic`) is supported.

The `realm` authentication parameter (default `kea-dhcpv4-server` is used for error messages when the basic HTTP authentication is required but the client is not authorized.

When the `clients` authentication list is configured and not empty, basic HTTP authentication is required. Each element of the list specifies a user ID and a password. The user ID is mandatory, must not be empty, and must not contain the colon (:) character. The password is optional; when it is not specified an empty password is used.

Note

The basic HTTP authentication user ID and password are encoded in UTF-8, but the current Kea JSON syntax only supports the Latin-1 (i.e. 0x00..0xff) Unicode subset.

To avoid exposing the user ID and/or the associated password, these values can be read from files. The syntax is extended by:

- The `directory` authentication parameter, which handles the common part of file paths. The default value is the empty string.
- The `password-file` client parameter, which, alongside the `directory` parameter, specifies the path of a file that can contain the password, or when no user ID is given, the whole basic HTTP authentication secret.
- The `user-file` client parameter, which, with the `directory` parameter, specifies the path of a file where the user ID can be read.

Since Kea-2.7.6 Kea supports multiple HTTP/HTTPS connections. Both IPv4 and IPv6 addresses can be used. The server will issue an error when changing the socket type from HTTP to HTTPS or from HTTPS to HTTP using the same address and port. This action is not allowed as it might introduce a security issue accidentally caused by a user mistake. A different address or port must be specified when using the "config-set" command to switch from HTTP to HTTPS or from HTTPS to HTTP. The same applies when modifying the configuration file and then running "config-reload" command.

When files are used, they are read when the configuration is loaded, to detect configuration errors as soon as possible.

```
"Dhcp4": {
  "control-sockets": [
    {
      "socket-type": "https",
      "socket-address": "10.20.30.40",
      "socket-port": 8004,
      "http-headers": [
        {
```

(continues on next page)

(continued from previous page)

```

        "name": "Strict-Transport-Security",
        "value": "max-age=31536000"
    }
],
"trust-anchor": "/path/to/the/ca-cert.pem",
"cert-file": "/path/to/the/agent-cert.pem",
"key-file": "/path/to/the/agent-key.pem",
"cert-required": true,
"authentication": {
    "type": "basic",
    "realm": "kea-dhcpv4-server",
    "clients": [
        {
            "user": "admin",
            "password": "1234"
        }
    ]
}
},
{
    "socket-type": "http",
    "socket-address": "2010:30:40::50",
    "socket-port": 8004
}
],
"subnet4": [
    {
        ...
    },
    ...
],
...
}

```

9.10 User Contexts in IPv4

Kea allows the loading of hook libraries that can sometimes benefit from additional parameters. If such a parameter is specific to the whole library, it is typically defined as a parameter for the hook library. However, sometimes there is a need to specify parameters that are different for each pool.

See *Comments and User Context* for additional background regarding the user-context idea. See *User Contexts in Hooks* for a discussion from the hooks perspective.

User contexts can be specified at global scope; at the shared-network, subnet, pool, client-class, option-data, or definition level; and via host reservation. One other useful feature is the ability to store comments or descriptions.

Let's consider an imaginary case of devices that have colored LED lights. Depending on their location, they should glow red, blue, or green. It would be easy to write a hook library that would send specific values, maybe as a vendor option. However, the server has to have some way to specify that value for each pool. This need is addressed by user contexts. In essence, any user data can be specified in the user context as long as it is a valid JSON map. For example, the aforementioned case of LED devices could be configured in the following way:

```

"Dhcp4": {
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          "pool": "192.0.2.10 - 192.0.2.20",
          # This is pool specific user context
          "user-context": { "color": "red" }
        }
      ],
      # This is a subnet-specific user context. Any type
      # of information can be entered here as long as it is valid JSON.
      "user-context": {
        "comment": "network on the second floor",
        "last-modified": "2017-09-04 13:32",
        "description": "you can put anything you like here",
        "phones": [ "x1234", "x2345" ],
        "devices-registered": 42,
        "billing": false
      }
    }
  ]
}

```

Kea does not interpret or use the user-context information; it simply stores it and makes it available to the hook libraries. It is up to each hook library to extract that information and use it. The parser translates a `comment` entry into a user context with the entry, which allows a comment to be attached inside the configuration itself.

9.11 Supported DHCP Standards

The following standards are currently supported in Kea:

- *BOOTP Vendor Information Extensions*, [RFC 1497](#): This requires the open source BOOTP hook to be loaded. See *libdhcp_bootp.so: Support for BOOTP Clients* for details.
- *Dynamic Host Configuration Protocol*, [RFC 1531](#): This RFC is obsolete and was replaced by RFC 1541, which in turn was replaced by RFC 2131. Kea supports all three RFCs.
- *Clarifications and Extensions for the Bootstrap Protocol*, [RFC 1532](#): This RFC has an editorial error and was quickly superseded by RFC 1542. Kea supports them both.
- *DHCP Options and BOOTP Vendor Extensions*, [RFC 1533](#): This RFC is obsolete and was replaced by RFC 2132. Nevertheless, Kea supports the options defined in it.
- *Dynamic Host Configuration Protocol*, [RFC 1541](#): This RFC is obsolete and was replaced by RFC 2131. Kea supports both.
- *Clarifications and Extensions for the Bootstrap Protocol*, [RFC 1542](#): This RFC is supported.
- *Dynamic Host Configuration Protocol*, [RFC 2131](#): Supported messages are DHCPDISCOVER (1), DHCPOFFER (2), DHCPREQUEST (3), DHCPRELEASE (7), DHCPINFORM (8), DHCPACK (5), and DHCPNAK(6).
- *DHCP Options and BOOTP Vendor Extensions*, [RFC 2132](#): Supported options are PAD (0), END(255), Message Type(53), DHCP Server Identifier (54), Domain Name (15), DNS Servers (6), IP Address Lease Time (51),

Subnet Mask (1), and Routers (3).

- *DHCP Options for Novell Directory Services*, [RFC 2241](#): All three options are supported.
- *Management of IP numbers by peg-dhcp*, [RFC 2322](#): This RFC is supported, although additional hardware is required for full deployment.
- *DHCP Option for The Open Group's User Authentication Protocol*, [RFC 2485](#): The option is supported.
- *DHCP Option to Disable Stateless Auto-Configuration in IPv4 Clients*, [RFC 2563](#): The option is supported.
- *DHCP Options for Service Location Protocol*, [RFC 2610](#): Both options are supported.
- *The Name Service Search Option for DHCP*, [RFC 2937](#): The option is supported.
- *The User Class Option for DHCP*, [RFC 3004](#): The user class is supported and can be used in any expression. The option's structure is not parsed and has to be referenced using hex.
- *The IPv4 Subnet Selection Option for DHCP*, [RFC 3011](#): The subnet-selection option is supported; if received in a packet, it is used in the subnet-selection process.
- *DHCP Relay Agent Information Option*, [RFC 3046](#): Relay Agent Information, Circuit ID, and Remote ID options are supported.
- *The DOCSIS (Data-Over-Cable Service Interface Specifications) Device Class DHCP (Dynamic Host Configuration Protocol) Relay Agent Information Sub-option*, [RFC 3256](#): The DOCSIS sub-option is supported and can be used in any expression.
- *Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)*, [RFC 3396](#): The Kea server can both receive and send long options. The long options can be configured and Kea will send them as separate instances if the payload length is longer than 255 octets.
- *Dynamic Host Configuration Protocol (DHCP) Domain Search Option*, [RFC 3397](#): The option is supported.
- *The Classless Static Route Option for Dynamic Host Configuration Protocol (DHCP) version 4*, [RFC 3442](#): The option is supported.
- *Dynamic Host Configuration Protocol (DHCP) Option for CableLabs Client Configuration*, [RFC 3495](#): The option and its suboptions 1, 2, 4, 5, 6, 7 and 8 are supported. See [CableLabs Client Conf Suboptions](#) for details.
- *Link Selection sub-option for the Relay Agent Option*, [RFC 3527](#): The link selection sub-option is supported.
- *PacketCable Security Ticket Control Sub-Option for the DHCP CableLabs Client Configuration (CCC) Option*, [RFC 3594](#): The Security Ticket Control sub-option is supported.
- *Key Distribution Center (KDC) Server Address Sub-option for the Dynamic Host Configuration Protocol (DHCP) CableLabs Client Configuration (CCC) Option*, [RFC 3634](#): The Key Distribution Center IP Address sub-option is supported.
- *Unused Dynamic Host Configuration Protocol (DHCP) Option Codes*, [RFC 3679](#): Kea does not support any of the old options that were obsoleted by this RFC.
- *Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4*, [RFC 3925](#): The Vendor-Identifying Vendor Class and Vendor-Identifying Vendor-Specific Information options are supported.
- *Reclassifying Dynamic Host Configuration Protocol version 4 (DHCPv4) Options*, [RFC 3942](#): Kea supports options with codes greater than 127.
- *Subscriber-ID Suboption for the DHCP Relay Agent Option*, [RFC 3993](#): The Subscriber-ID option is supported.
- *Dynamic Host Configuration Protocol (DHCP) Options for Broadcast and Multicast Control Servers*, [RFC 4280](#): The DHCPv4 options are supported.
- *Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)*, [RFC 4361](#): The DUID in DHCPv4 is supported.

- *Dynamic Host Configuration Protocol (DHCP) Leasequery*, RFC 4388: The server functionality is supported. This requires leasequery hook. See [libdhcp_lease_query.so: Leasequery Support](#) for details.
- *Dynamic Host Configuration Protocol (DHCP) Options for the Intel Preboot eXecution Environment (PXE)*, RFC 4578: All three options defined are supported.
- *A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)*, RFC 4701: The DHCPv4 server supports DHCID records. The DHCP-DDNS server must be running to add, update, and/or delete DHCID records.
- *The Dynamic Host Configuration Protocol (DHCP) Client Fully Qualified Domain Name (FQDN) Option*, RFC 4702: The Kea server is able to handle the Client FQDN option. Also, it is able to use the [kea-dhcp-ddns](#) component to initiate appropriate DNS Update operations.
- *Resolution of Fully Qualified Domain Name (FQDN) Conflicts among Dynamic Host Configuration Protocol (DHCP) Clients*, RFC 4703: The DHCPv6 server uses a DHCP-DDNS server to resolve conflicts.
- *Timezone Options for DHCP*: RFC 4833: Both DHCPv4 options are supported.
- *The Dynamic Host Configuration Protocol Version 4 (DHCPv4) Relay Agent Flags Suboption*: RFC 5010: The Relay Agent Flags sub-option is understood by Kea and can be used in an expression.
- *Server Identifier Override sub-option for the Relay Agent Option*, RFC 5107: The server identifier override sub-option is supported. The implementation is not complete according to the RFC, because the server does not store the RAI, but the functionality handles expected use cases.
- *DHCP Options for Protocol for Carrying Authentication for Network Access (PANA) Authentication Agents*: RFC 5192: The PANA option is supported.
- *Discovering Location-to-Service Translation (LoST) Servers Using the Dynamic Host Configuration Protocol (DHCP)*: RFC 5223: The LOST option is supported.
- *Control And Provisioning of Wireless Access Points (CAPWAP) Access Controller DHCP Option*: RFC 5417: The CAPWAP for IPv4 option is supported.
- *DHCPv4 Lease Query by Relay Agent Remote ID*, RFC 6148: The leasequery by remote-id is supported. This requires leasequery hook. See [libdhcp_lease_query.so: Leasequery Support](#) for details.
- *Client Identifier Option in DHCP Server Replies*, RFC 6842: The server by default sends back the client-id option. That capability can be disabled. See [Echoing Client-ID \(RFC 6842\)](#) for details.
- *The DHCPv4 Relay Agent Identifier Sub-Option*, RFC 6925: The relay-id option is supported and can be used in all features that are using expressions (client classification, flex-id reservations, etc.).
- *DHCPv4 Bulk Leasequery*, RFC 6926: The server functionality (TCP connections, new query types, multiple responses, etc.) is supported. This requires leasequery hook. See [libdhcp_lease_query.so: Leasequery Support](#) for details.
- *Generalized UDP Source Port for the DHCP Relay Agent Option*, RFC 8357: The Kea server handles the Relay Agent Information Source Port sub-option in a received message, remembers the UDP port, and sends back a reply to the same relay agent using this UDP port.
- *Captive-Portal Identification in DHCP and Router Advertisements (RAs)*, RFC 8910: The Kea server can configure both v4 and v6 versions of the captive portal options.
- *IPv6-Only Preferred Option for DHCPv4*, RFC 8925: The Kea server is able to designate its pools and subnets as IPv6-Only Preferred and send back the v6-only-preferred option to clients that requested it.
- *DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)*, RFC 9463: The Kea server supports the DNR option.

9.11.1 Known RFC Violations

In principle, Kea aspires to be a reference implementation and aims to implement 100% of the RFC standards. However, in some cases there are practical aspects that prevent Kea from completely adhering to the text of all RFC documents.

- [RFC 2131](#), page 30, says that if the incoming DHCPREQUEST packet has no "requested IP address" option and `ciaddr` is not set, the server is supposed to respond with NAK. However, broken clients exist that will always send a DHCPREQUEST without those options indicated. In that event, Kea accepts the DHCPREQUEST, assigns an address, and responds with an ACK.
- [RFC 2131](#), table 5, says that messages of type DHCPDECLINE or DHCPRELEASE must have the server identifier set and should be dropped if that option is missing. However, ISC DHCP does not enforce this, presumably as a compatibility effort for broken clients, and the Kea team decided to follow suit.

9.12 DHCPv4 Server Limitations

These are the current known limitations of the Kea DHCPv4 server software. Most of them are reflections of the current stage of development and should be treated as “not yet implemented,” rather than as actual limitations. However, some of them are implications of the design choices made. Those are clearly marked as such.

- On the Linux and BSD system families, DHCP messages are sent and received over raw sockets (using LPF and BPF) and all packet headers (including data link layer, IP, and UDP headers) are created and parsed by Kea, rather than by the system kernel. Currently, Kea can only parse the data-link layer headers with a format adhering to the IEEE 802.3 standard, and assumes this data-link-layer header format for all interfaces. Thus, Kea does not work on interfaces which use different data-link-layer header formats (e.g. Infiniband).

9.13 Kea DHCPv4 Server Examples

A collection of simple-to-use examples for the DHCPv4 component of Kea is available with the source files, located in the `doc/examples/kea4` directory.

9.14 Configuration Backend in DHCPv4

In the [Kea Configuration Backend](#) section we have described the Configuration Backend (CB) feature, its applicability, and its limitations. This section focuses on the usage of the CB with the Kea DHCPv4 server. It lists the supported parameters, describes limitations, and gives examples of DHCPv4 server configurations to take advantage of the CB. Please also refer to the corresponding section [Configuration Backend in DHCPv6](#) for DHCPv6-specific usage of the CB.

9.14.1 Supported Parameters

The ultimate goal for the CB is to serve as a central configuration repository for one or multiple Kea servers connected to a database. In currently supported Kea versions, only a subset of the DHCPv4 server parameters can be configured in the database. All other parameters must be specified in the JSON configuration file, if required.

All supported parameters can be configured via `libdhcp_cb_cmds.so`. The general rule is that scalar global parameters are set using `remote-global-parameter4-set`; shared-network-specific parameters are set using `remote-network4-set`; and subnet-level and pool-level parameters are set using `remote-subnet4-set`. Whenever there is an exception to this general rule, it is highlighted in the table. Non-scalar global parameters have dedicated commands; for example, the global DHCPv4 options (`option-data`) are modified using `remote-option4-global-set`. Client classes, together with class-specific option definitions and DHCPv4 options, are configured using the `remote-class4-set` command.

The [Configuration Sharing and Server Tags](#) section explains the concept of shareable and non-shareable configuration elements and the limitations for sharing them between multiple servers. In the DHCP configuration (both DHCPv4

and DHCPv6), the shareable configuration elements are subnets and shared networks. Thus, they can be explicitly associated with multiple server tags. The global parameters, option definitions, and global options are non-shareable and can be associated with only one server tag. This rule does not apply to the configuration elements associated with all servers. Any configuration element associated with all servers (using the all keyword as a server tag) is used by all servers connecting to the configuration database.

The following table lists DHCPv4-specific parameters supported by the configuration backend, with an indication of the level of the hierarchy at which it is currently supported.

Table 9: List of DHCPv4 parameters supported by the configuration backend

Parameter	Global	Client Class	Shared Network	Subnet	Pool
4o6-interface	n/a	n/a	n/a	yes	n/a
4o6-interface-id	n/a	n/a	n/a	yes	n/a
4o6-subnet	n/a	n/a	n/a	yes	n/a
allocator	yes	n/a	yes	yes	n/a
boot-file-name	yes	yes	yes	yes	n/a
cache-max-age	yes	n/a	no	no	n/a
cache-threshold	yes	n/a	no	no	n/a
calculate-tee-times	yes	n/a	yes	yes	n/a
client-class	n/a	n/a	yes	yes	yes
ddns-send-update	yes	n/a	yes	yes	n/a
ddns-override-no-update	yes	n/a	yes	yes	n/a
ddns-override-client-update	yes	n/a	yes	yes	n/a
ddns-replace-client-name	yes	n/a	yes	yes	n/a
ddns-generated-prefix	yes	n/a	yes	yes	n/a
ddns-qualifying-suffix	yes	n/a	yes	yes	n/a
decline-probation-period	yes	n/a	n/a	n/a	n/a
dhcp4o6-port	yes	n/a	n/a	n/a	n/a
echo-client-id	yes	n/a	n/a	n/a	n/a
hostname-char-set	yes	n/a	yes	yes	n/a
hostname-char-replacement	yes	n/a	yes	yes	n/a
interface	n/a	n/a	yes	yes	n/a
match-client-id	yes	n/a	yes	yes	n/a
min-valid-lifetime	yes	yes	yes	yes	n/a
max-valid-lifetime	yes	yes	yes	yes	n/a
next-server	yes	yes	yes	yes	n/a
option-data	yes (via remote-option4-global-set)	yes	yes	yes	yes
option-def	yes (via remote-option-def4-set)	yes	n/a	n/a	n/a
rebind-timer	yes	n/a	yes	yes	n/a
renew-timer	yes	n/a	yes	yes	n/a
server-hostname	yes	yes	yes	yes	n/a
valid-lifetime	yes	yes	yes	yes	n/a
relay	n/a	n/a	yes	yes	n/a
require-client-classes	no	n/a	yes	yes	yes
evaluate-additional-classes	no	n/a	yes	yes	yes
reservations-global	yes	n/a	yes	yes	n/a
reservations-in-subnet	yes	n/a	yes	yes	n/a
reservations-out-of-pool	yes	n/a	yes	yes	n/a
t1-percent	yes	n/a	yes	yes	n/a
t2-percent	yes	n/a	yes	yes	n/a

- yes - indicates that the parameter is supported at the given level of the hierarchy and can be configured via the

configuration backend.

- **no** - indicates that a parameter is supported at the given level of the hierarchy but cannot be configured via the configuration backend.
- **n/a** - indicates that a given parameter is not applicable at the particular level of the hierarchy or that the server does not support the parameter at that level.

Some scalar parameters contained by top-level global maps are supported by the configuration backend.

Table 10: List of DHCPv4 map parameters supported by the configuration backend

Parameter name (flat naming format)	Global map	Parameter name
compatibility.ignore-dhcp-server-identifier	compatibility	ignore-dhcp-server-identifier
compatibility.ignore-rai-link-selection	compatibility	ignore-rai-link-selection
compatibility.lenient-option-parsing	compatibility	lenient-option-parsing
compatibility.exclude-first-last-24	compatibility	exclude-first-last-24
dhcp-ddns.enable-updates	dhcp-ddns	enable-updates
dhcp-ddns.max-queue-size	dhcp-ddns	max-queue-size
dhcp-ddns.ncr-format	dhcp-ddns	ncr-format
dhcp-ddns.ncr-protocol	dhcp-ddns	ncr-protocol
dhcp-ddns.sender-ip	dhcp-ddns	sender-ip
dhcp-ddns.sender-port	dhcp-ddns	sender-port
dhcp-ddns.server-ip	dhcp-ddns	server-ip
dhcp-ddns.server-port	dhcp-ddns	server-port
expired-leases-processing.flush-reclaimed-timer-wait-time	expired-leases-processing	flush-reclaimed-timer-wait-time
expired-leases-processing.hold-reclaimed-time	expired-leases-processing	hold-reclaimed-time
expired-leases-processing.max-reclaim-leases	expired-leases-processing	max-reclaim-leases
expired-leases-processing.max-reclaim-time	expired-leases-processing	max-reclaim-time
expired-leases-processing.reclaim-timer-wait-time	expired-leases-processing	reclaim-timer-wait-time
expired-leases-processing.unwarned-reclaim-cycles	expired-leases-processing	unwarned-reclaim-cycles
multi-threading.enable-multi-threading	multi-threading	enable-multi-threading
multi-threading.thread-pool-size	multi-threading	thread-pool-size
multi-threading.packet-queue-size	multi-threading	packet-queue-size
sanity-checks.lease-checks	sanity-checks	lease-checks
sanity-checks.extended-info-checks	sanity-checks	extended-info-checks
dhcp-queue-control.enable-queue	dhcp-queue-control	enable-queue
dhcp-queue-control.queue-type	dhcp-queue-control	queue-type
dhcp-queue-control.capacity	dhcp-queue-control	capacity

9.14.2 Enabling the Configuration Backend

Consider the following configuration snippet, which uses a MySQL configuration database:

```
{
  "Dhcp4": {
    "server-tag": "my DHCPv4 server",
```

(continues on next page)

(continued from previous page)

```

    "config-control": {
      "config-databases": [
        {
          "type": "mysql",
          "name": "kea",
          "user": "kea",
          "password": "1234",
          "host": "192.0.2.1",
          "port": 3302
        }
      ],
      "config-fetch-wait-time": 20
    },
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_mysql.so"
      },
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_cb_cmds.so"
      }
    ]
  }
}

```

The following snippet illustrates the use of a PostgreSQL database:

```

{
  "Dhcp4": {
    "server-tag": "my DHCPv4 server",
    "config-control": {
      "config-databases": [
        {
          "type": "postgresql",
          "name": "kea",
          "user": "kea",
          "password": "1234",
          "host": "192.0.2.1",
          "port": 3302
        }
      ],
      "config-fetch-wait-time": 20
    },
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_pgsql.so"
      },
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_cb_cmds.so"
      }
    ]
  }
}

```

The `config-control` map contains two parameters. `config-databases` is a list that contains one element, which includes the database type, its location, and the credentials to be used to connect to this database. (Note that the parameters specified here correspond to the database specification for the lease database backend and hosts database backend.) Currently only one database connection can be specified on the `config-databases` list. The server connects to this database during startup or reconfiguration, and fetches the configuration available for this server from the database. This configuration is merged into the configuration read from the configuration file.

Note

Whenever there is a conflict between the parameters specified in the configuration file and the database, the parameters from the database take precedence. We strongly recommend avoiding the duplication of parameters in the file and the database, but this recommendation is not enforced by the Kea servers. In particular, if the subnets' configuration is sourced from the database, we recommend that all subnets be specified in the database and that no subnets be specified in the configuration file. It is possible to specify the subnets in both places, but the subnets in the configuration file with overlapping IDs and/or prefixes with the subnets from the database will be superseded by those from the database.

Once the Kea server is configured, it starts periodically polling the database for configuration changes. The polling frequency is controlled by the `config-fetch-wait-time` parameter, expressed in seconds; it is the period between the time when the server completed its last poll (and possibly the local configuration update) and the time when it will begin polling again. In the example above, this period is set to 20 seconds. This means that after adding a new configuration into the database (e.g. adding a new subnet), it will take up to 20 seconds (plus the time needed to fetch and apply the new configuration) before the server starts using this subnet. The lower the `config-fetch-wait-time` value, the shorter the time for the server to react to incremental configuration updates in the database. On the other hand, polling the database too frequently may impact the DHCP server's performance, because the server needs to make at least one query to the database to discover any pending configuration updates. The default value of `config-fetch-wait-time` is 30 seconds.

The `config-backend-pull` command can be used to force the server to immediately poll any configuration changes from the database and avoid waiting for the next fetch cycle.

In the configuration examples above, two hook libraries are loaded. The first is a library which implements the configuration backend for a specific database type: `libdhcp_mysql.so` provides support for MySQL and `libdhcp_pgsql.so` provides support for PostgreSQL. The library loaded must match the database type specified within the `config-control` parameter; otherwise an error is logged when the server attempts to load its configuration, and the load fails.

The second hook library, `libdhcp_cb_cmds.so`, is optional. It should be loaded when the Kea server instance is to be used to manage the configuration in the database. See the `libdhcp_cb_cmds.so: Configuration Backend Commands` section for details.

9.15 Kea DHCPv4 Compatibility Configuration Parameters

ISC's intention is for Kea to follow the RFC documents to promote better standards compliance. However, many buggy DHCP implementations already exist that cannot be easily fixed or upgraded. Therefore, Kea provides an easy-to-use compatibility mode for broken or non-compliant clients. For that purpose, the compatibility option must be enabled to permit uncommon practices:

```
{
  "Dhcp4": {
    "compatibility": {
    }
  }
}
```

9.15.1 Lenient Option Parsing

By default, tuple fields defined in custom options are parsed as a set of length-value pairs.

With `"lenient-option-parsing": true`, if a length ever exceeded the rest of the option's buffer, previous versions of Kea returned a log message unable to parse the opaque data tuple, the buffer length is `x`, but the tuple length is `y` with `x < y`; this no longer occurs. Instead, the value is considered to be the rest of the buffer, or in terms of the log message above, the tuple length `y` becomes `x`.

```
{
  "Dhcp4": {
    "compatibility": {
      "lenient-option-parsing": true
    }
  }
}
```

Starting with Kea version 2.5.8, this parsing is extended to silently ignore FQDN (81) options with some invalid domain names.

9.15.2 Ignore DHCP Server Identifier

With `"ignore-dhcp-server-identifier": true`, the server does not check the address in the DHCP Server Identifier option, i.e. whether a query is sent to this server or another one (and in the second case dropping the query).

```
{
  "Dhcp4": {
    "compatibility": {
      "ignore-dhcp-server-identifier": true
    }
  }
}
```

9.15.3 Ignore RAI Link Selection

With `"ignore-rai-link-selection": true`, Relay Agent Information Link Selection sub-option data is not used for subnet selection. In this case, normal logic drives the subnet selection, instead of attempting to use the subnet specified by the sub-option. This option is not RFC-compliant and is set to `false` by default. Setting this option to `true` can help with subnet selection in certain scenarios; for example, when DHCP relays do not allow the administrator to specify which sub-options are included in the Relay Agent Information option, and include incorrect Link Selection information.

```
{
  "Dhcp4": {
    "compatibility": {
      "ignore-rai-link-selection": true
    }
  }
}
```

9.15.4 Exclude First Last Addresses in /24 Subnets or Larger

The `exclude-first-last-24` compatibility flag is described in *Configuration of IPv4 Address Pools* (when true .0 and .255 addresses are excluded from subnets with prefix length less than or equal to 24).

9.16 Address Allocation Strategies in DHCPv4

A DHCP server follows a complicated algorithm to select an IPv4 address for a client. It prefers assigning specific addresses requested by the client and the addresses for which the client has reservations.

If the client requests no particular address and has no reservations, or other clients are already using any requested addresses, the server must find another available address within the configured pools. A server function called an "allocator" is responsible in Kea for finding an available address in such a case.

The Kea DHCPv4 server provides configuration parameters to select different allocators at the global, shared-network, and subnet levels. Consider the following example:

```
{
  "Dhcp4": {
    "allocator": "random",
    "subnet4": [
      {
        "id": 1,
        "subnet": "10.0.0.0/8",
        "allocator": "iterative"
      },
      {
        "id": 2,
        "subnet": "192.0.2.0/24"
      }
    ]
  }
}
```

This allocator overrides the default iterative allocation strategy at the global level and selects the random allocation instead. The random allocation will be used for the subnet with ID 2, while the iterative allocation will be used for the subnet with ID 1.

The following sections describe the supported allocators and their recommended uses.

9.16.1 Allocators Comparison

In the table below, we briefly compare the supported allocators, all of which are described in detail in later sections.

Table 11: Comparison of the lease allocators supported by Kea DHCPv4

Allocator	Low Utilization Performance	High Utilization Performance	Lease Ran- domization	Startup/Configurat	Memory Usage
Iterative	very high	low	no	very fast	low
Random	high	low	yes	very fast	high (vary- ing)
Free Lease Queue	high	high	yes	slow (depends on pool sizes)	high (vary- ing)

9.16.2 Iterative Allocator

This is the default allocator used by the Kea DHCPv4 server. It remembers the last offered address and offers this address, increased by one, to the next client. For example, it may offer addresses in this order: 192.0.2.10, 192.0.2.11, 192.0.2.12, and so on. The time to find and offer the next address is very short; thus, this is the most performant allocator when pool utilization is low and there is a high probability that the next address is available.

The iterative allocation underperforms when multiple DHCP servers share a lease database or are connected to a cluster. The servers tend to offer and allocate the same blocks of addresses to different clients independently, which causes many allocation conflicts between the servers and retransmissions by clients. A random allocation addresses this issue by dispersing the allocation order.

9.16.3 Random Allocator

The random allocator uses a uniform randomization function to select offered addresses from subnet pools. It is suitable in deployments where multiple servers are connected to a shared database or a database cluster. By dispersing the offered addresses, the servers minimize the risk of allocating the same address to two different clients at the same or nearly the same time. In addition, it improves the server's resilience against attacks based on allocation predictability.

The random allocator is, however, slightly slower than the iterative allocator. Moreover, it increases the server's memory consumption because it must remember randomized addresses to avoid offering them repeatedly. Memory consumption grows with the number of offered addresses; in other words, larger pools and more clients increase memory consumption by random allocation.

The following configuration snippet shows how to select the random allocator for a subnet:

```
{
  "Dhcp4": {
    "allocator": "random",
    "subnet4": [
      {
        "id": 1,
        "subnet": "10.0.0.0/8",
        "allocator": "random"
      }
    ]
  }
}
```

9.16.4 Free Lease Queue Allocator

This is a sophisticated allocator whose use should be considered in subnets with highly utilized address pools. In such cases, it can take a considerable amount of time for the iterative or random allocator to find an available address, because they must repeatedly check whether there is a valid lease for an address they will offer. The number of checks can be as high as the number of addresses in the subnet when the subnet pools are exhausted, which can have a direct negative impact on the DHCP response time for each request.

The Free Lease Queue (FLQ) allocator tracks lease allocations and de-allocations and maintains a running list of available addresses for each address pool. It allows an available lease to be selected within a constant time, regardless of the subnet pools' utilization. The allocator continuously updates the list of free leases by removing any allocated leases and adding released or reclaimed ones.

The following configuration snippet shows how to select the FLQ allocator for a subnet:

```
{
  "Dhcp4": {
```

(continues on next page)

(continued from previous page)

```
"subnet4": [
  {
    "id": 1,
    "subnet": "192.0.2.0/24",
    "allocator": "flq"
  }
]
```

There are several considerations that the administrator should take into account before using this allocator. The FLQ allocator can heavily impact the server's startup and reconfiguration time, because the allocator has to populate the list of free leases for each subnet where it is used. These delays can be observed both during the configuration reload and when the subnets are created using *libdhcp_subnet_cmds.so*. This allocator increases memory consumption to hold the list of free leases, proportional to the total size of the address pools for which the FLQ allocator is used. Finally, lease reclamation must be enabled with a low value of the `reclaim-timer-wait-time` parameter, to ensure that the server frequently collects expired leases and makes them available for allocation via the free lease queue; expired leases are not considered free by the allocator until they are reclaimed by the server. See [Lease Reclamation](#) for more details about the lease reclamation process.

We recommend that the FLQ allocator be selected only after careful consideration. For example, using it for a subnet with a /8 pool may delay the server's startup by 15 seconds or more. On the other hand, the startup delay and the memory consumption increase should be acceptable for subnets with a /16 pool or smaller. We also recommend specifying another allocator type in the global configuration settings and overriding this selection at the subnet or shared-network level, to use the FLQ allocator only for selected subnets. That way, when a new subnet is added without an allocator specification, the global setting is used, thus avoiding unnecessary impact on the server's startup time.

Warning

The FLQ allocator is not suitable for use with a shared lease database (i.e., when multiple Kea servers store leases in the same database). The servers are unaware of the expired leases reclaimed by the sibling servers and never return them to their local free lease queues. As a result, the servers will not be able to offer some of the available leases to the clients. Only a server reclaiming a particular lease will be able to offer it.

THE DHCPV6 SERVER

10.1 Starting and Stopping the DHCPv6 Server

It is recommended that the Kea DHCPv6 server be started and stopped using *keactrl* (described in *Managing Kea with keactrl*); however, it is also possible to run the server directly via the *kea-dhcp6* command, which accepts the following command-line switches:

- **-c file** - specifies the configuration file. This is the only mandatory switch.
- **-d** - specifies whether the server logging should be switched to debug/verbose mode. In verbose mode, the logging severity and debuglevel specified in the configuration file are ignored; "debug" severity and the maximum debuglevel (99) are assumed. The flag is convenient for temporarily switching the server into maximum verbosity, e.g. when debugging.
- **-p server-port** - specifies the local UDP port on which the server listens. This is only useful during testing, as a DHCPv6 server listening on ports other than the standard ones is not able to handle regular DHCPv6 queries.
- **-P client-port** - specifies the remote UDP port to which the server sends all responses. This is only useful during testing, as a DHCPv6 server sending responses to ports other than the standard ones is not able to handle regular DHCPv6 queries.
- **-t file** - specifies a configuration file to be tested. *kea-dhcp6* loads it, checks it, and exits. During the test, log messages are printed to standard output and error messages to standard error. The result of the test is reported through the exit code (0 = configuration looks OK, 1 = error encountered). The check is not comprehensive; certain checks are possible only when running the server.
- **-T file** - specifies a configuration file to be tested. *kea-dhcp6* loads it, checks it, and exits. It performs extra checks beyond what **-t** offers, such as establishing database connections (for the lease backend, host reservations backend, configuration backend, and forensic logging backend), loading hook libraries, parsing hook-library configurations, etc. It does not open UNIX or TCP/UDP sockets, nor does it open or rotate files, as any of these actions could interfere with a running process on the same machine.
- **-v** - displays the Kea version and exits.
- **-V** - displays the Kea extended version with additional parameters and exits. The listing includes the versions of the libraries dynamically linked to Kea.
- **-W** - displays the Kea configuration report and exits. The report is a copy of the `config.report` file produced by `meson setup`; it is embedded in the executable binary.

The contents of the `config.report` file may also be accessed by examining certain libraries in the installation tree or in the source tree.

```
# from installation using libkea-process.so
$ strings ${prefix}/lib/libkea-process.so | sed -n 's/;;; //p'
```

(continues on next page)

(continued from previous page)

```
# from sources using libkea-process.so
$ strings src/lib/process/.libs/libkea-process.so | sed -n 's/;;;; //p'

# from sources using libkea-process.a
$ strings src/lib/process/.libs/libkea-process.a | sed -n 's/;;;; //p'

# from sources using libcfgrpt.a
$ strings src/lib/process/cfgrpt/.libs/libcfgrpt.a | sed -n 's/;;;; //p'
```

- **-X** - As of Kea 3.0, disables path and permissions restrictions. The server will emit a warning at startup that security restrictions have been disabled. Do not use this mode of operation without careful consideration and taking any necessary precautions. Failure to do so can expose deployments to security vulnerabilities. For more information please read section *Securing a Kea Deployment*.

On startup, the server detects available network interfaces and attempts to open UDP sockets on all interfaces listed in the configuration file. Since the DHCPv6 server opens privileged ports, it requires root access; this daemon must be run as root.

During startup, the server attempts to create a PID file of the form: `[pidfile_dir]/[conf name].kea-dhcp6.pid`, where:

- **pidfile_dir** - is `[prefix]/[localstatedir]/run/kea` where `prefix` and `localstatedir` are the values passed into meson setup using `--prefix` and `--localstatedir` which default to `/usr/local` and `var` respectively. So the whole `pidfile_dir` defaults to `/usr/local/var`. Note that this value may be overridden at runtime by setting the environment variable `KEA_PIDFILE_DIR` intended primarily for testing purposes.
- **conf name**: The configuration file name used to start the server, minus all preceding paths and the file extension. For example, given a pathname of `/usr/local/etc/kea/myconf.txt`, the portion used would be `myconf`.

If the file already exists and contains the PID of a live process, the server issues a `DHCP6_ALREADY_RUNNING` log message and exits. It is possible, though unlikely, that the file is a remnant of a system crash and the process to which the PID belongs is unrelated to Kea. In such a case, it would be necessary to manually delete the PID file.

The server can be stopped using the `kill` command. When running in a console, the server can also be shut down by pressing `Ctrl-c`. Kea detects the key combination and shuts down gracefully.

The reconfiguration of each Kea server is triggered by the `SIGHUP` signal. When a server receives the `SIGHUP` signal it rereads its configuration file and, if the new configuration is valid, uses the new configuration. If the new configuration proves to be invalid, the server retains its current configuration; however, in some cases a fatal error message is logged indicating that the server is no longer providing any service: a working configuration must be loaded as soon as possible.

10.2 DHCPv6 Server Configuration

10.2.1 Introduction

This section explains how to configure the Kea DHCPv6 server using a configuration file.

Before DHCPv6 is started, its configuration file must be created. The basic configuration is as follows:

```
{
# DHCPv6 configuration starts on the next line
"Dhcp6": {

# First we set up global values
  "valid-lifetime": 4000,
  "renew-timer": 1000,
```

(continues on next page)

(continued from previous page)

```

    "rebind-timer": 2000,
    "preferred-lifetime": 3000,

# Next we set up the interfaces to be used by the server.
    "interfaces-config": {
        "interfaces": [ "eth0" ]
    },

# And we specify the type of lease database
    "lease-database": {
        "type": "memfile",
        "persist": true,
        "name": "/var/lib/kea/dhcp6.leases"
    },

# Finally, we list the subnets from which we will be leasing addresses.
    "subnet6": [
        {
            "id": 1,
            "subnet": "2001:db8:1::/64",
            "pools": [
                {
                    "pool": "2001:db8:1::1-2001:db8:1::ffff"
                }
            ]
        }
    ]
# DHCPv6 configuration ends with the next line
}
}

```

The following paragraphs provide a brief overview of the parameters in the above example, along with their format. Subsequent sections of this chapter go into much greater detail for these and other parameters.

The lines starting with a hash (#) are comments and are ignored by the server; they do not impact its operation in any way.

The configuration starts in the first line with the initial opening curly bracket (or brace). Each configuration must contain an object specifying the configuration of the Kea module using it. In the example above, this object is called Dhc6.

The Dhc6 configuration starts with the "Dhc6": { line and ends with the corresponding closing brace (in the above example, the brace after the last comment). Everything defined between those lines is considered to be the Dhc6 configuration.

In general, the order in which those parameters appear does not matter, but there are two caveats. The first one is that the configuration file must be well-formed JSON, meaning that the parameters for any given scope must be separated by a comma, and there must not be a comma after the last parameter. When reordering a configuration file, moving a parameter to or from the last position in a given scope may also require moving the comma. The second caveat is that it is uncommon — although legal JSON — to repeat the same parameter multiple times. If that happens, the last occurrence of a given parameter in a given scope is used, while all previous instances are ignored. This is unlikely to cause any confusion as there are no real-life reasons to keep multiple copies of the same parameter in the configuration file.

The first few DHCPv6 configuration elements define some global parameters. `valid-lifetime` defines how long the addresses (leases) given out by the server are valid; the default is for a client to be allowed to use a given address for 4000 seconds. (Note that integer numbers are specified as is, without any quotes around them.) The address will become deprecated in 3000 seconds, i.e. clients are allowed to keep old connections, but cannot use this address to create new connections. `renew-timer` and `rebind-timer` are values (also in seconds) that define T1 and T2 timers, which govern when the client begins the renewal and rebind procedures.

The `interfaces-config` map specifies the network interfaces on which the server should listen to DHCP messages. The `interfaces` parameter specifies a list of network interfaces on which the server should listen. Lists are opened and closed with square brackets, with elements separated by commas. To listen on two interfaces, the `interfaces-config` element should look like this:

```
{
  "interfaces-config": {
    "interfaces": [ "eth0", "eth1" ]
  },
  ...
}
```

The next lines define the lease database, the place where the server stores its lease information. This particular example tells the server to use `memfile`, which is the simplest and fastest database backend. It uses an in-memory database and stores leases on disk in a CSV (comma-separated values) file. This is a very simple configuration example; usually the lease database configuration is more extensive and contains additional parameters. Note that `lease-database` is an object and opens up a new scope, using an opening brace. Its parameters (just one in this example: `type`) follow. If there were more than one, they would be separated by commas. This scope is closed with a closing brace. As more parameters for the `Dhcp6` definition follow, a trailing comma is present.

Finally, we need to define a list of IPv6 subnets. This is the most important DHCPv6 configuration structure, as the server uses that information to process clients' requests. It defines all subnets from which the server is expected to receive DHCP requests. The subnets are specified with the `subnet6` parameter. It is a list, so it starts and ends with square brackets. Each subnet definition in the list has several attributes associated with it, so it is a structure and is opened and closed with braces. At a minimum, a subnet definition must have at least two parameters: `subnet`, which defines the whole subnet; and `pools`, which is a list of dynamically allocated pools that are governed by the DHCP server.

The example contains a single subnet. If more than one were defined, additional elements in the `subnet6` parameter would be specified and separated by commas. For example, to define two subnets, the following syntax would be used:

```
{
  "subnet6": [
    {
      "id": 1,
      "pools": [ { "pool": "2001:db8:1::/112" } ],
      "subnet": "2001:db8:1::/64"
    },
    {
      "id": 2,
      "pools": [ { "pool": "2001:db8:2::1-2001:db8:2::ffff" } ],
      "subnet": "2001:db8:2::/64"
    }
  ],
  ...
}
```

Note that indentation is optional and is used for aesthetic purposes only. In some cases it may be preferable to use more compact notation.

After all the parameters have been specified, there are two contexts open: `global` and `Dhcp6`; thus, two closing curly brackets must be used to close them.

10.2.2 Lease Storage

All leases issued by the server are stored in the lease database. There are three database backends available: `memfile` (the default), `MySQL`, `PostgreSQL`.

10.2.2.1 Memfile - Basic Storage for Leases

The server is able to store lease data in different repositories. Larger deployments may elect to store leases in a database; [Lease Database Configuration](#) describes this option. In typical smaller deployments, though, the server stores lease information in a CSV file rather than a database. As well as requiring less administration, an advantage of using a file for storage is that it eliminates a dependency on third-party database software.

The configuration of the `memfile` backend is controlled through the `Dhcp6/lease-database` parameters. The `type` parameter is mandatory and specifies which storage for leases the server should use, through the `"memfile"` value. The following list gives additional optional parameters that can be used to configure the `memfile` backend.

- **`persist`**: controls whether the new leases and updates to existing leases are written to the file. It is strongly recommended that the value of this parameter be set to `true` at all times during the server's normal operation. Not writing leases to disk means that if a server is restarted (e.g. after a power failure), it will not know which addresses have been assigned. As a result, it may assign new clients addresses that are already in use. The value of `false` is mostly useful for performance-testing purposes. The default value of the `persist` parameter is `true`, which enables writing lease updates to the lease file.
- **`name`**: specifies the lease file in which new leases and lease updates are recorded. The default value for this parameter is `"[kea-install-dir]/var/lib/kea/kea-leases6.csv"`.

Note

As of Kea 2.7.9, lease files may only be loaded from the data directory determined during compilation: `"[kea-install-dir]/var/lib/kea"`. This path may be overridden at startup by setting the environment variable `KEA_DHCP_DATA_DIR` to the desired path. If a path other than this value is used in `name`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in specifying a custom file name simply omit the path component from `name`.

- **`lfc-interval`**: specifies the interval, in seconds, at which the server will perform a lease file cleanup (LFC). This removes redundant (historical) information from the lease file and effectively reduces the lease file size. The cleanup process is described in more detail later in this section. The default value of the `lfc-interval` is `3600`. A value of `0` disables the LFC.
- **`max-row-errors`**: specifies the number of row errors before the server stops attempting to load a lease file. When the server loads a lease file, it is processed row by row, each row containing a single lease. If a row is flawed and cannot be processed correctly the server logs it, discards the row, and goes on to the next row. This parameter can be used to set a limit on the number of such discards that can occur, after which the server abandons the effort and exits. The default value of `0` disables the limit and allows the server to process the entire file, regardless of how many rows are discarded.

An example configuration of the `memfile` backend is presented below:

```
"Dhcp6": {
  "lease-database": {
    "type": "memfile",
    "persist": true,
    "name": "kea-leases6.csv",
```

(continues on next page)

(continued from previous page)

```
"lfc-interval": 1800,  
"max-row-errors": 100  
}  
}
```

This configuration selects `kea-leases6.csv` as the storage file for lease information and enables persistence (writing lease updates to this file). It also configures the backend to perform a periodic cleanup of the lease file every 1800 seconds (30 minutes) and sets the maximum number of row errors to 100.

10.2.2.2 Why Is Lease File Cleanup Necessary?

It is important to know how the lease file contents are organized to understand why the periodic lease file cleanup is needed. Every time the server updates a lease or creates a new lease for a client, the new lease information must be recorded in the lease file. For performance reasons, the server does not update the existing client's lease in the file, as this would potentially require rewriting the entire file. Instead, it simply appends the new lease information to the end of the file; the previous lease entries for the client are not removed. When the server loads leases from the lease file, e.g. at server startup, it assumes that the latest lease entry for the client is the valid one. Previous entries are discarded, meaning that the server can reconstruct accurate information about the leases even though there may be many lease entries for each client. However, storing many entries for each client results in a bloated lease file and impairs the performance of the server's startup and reconfiguration, as it needs to process a larger number of lease entries.

Lease file cleanup (LFC) removes all previous entries for each client and leaves only the latest ones. The interval at which the cleanup is performed is configurable, and it should be selected according to the frequency of lease renewals initiated by the clients. The more frequent the renewals, the smaller the value of `lfc-interval` should be. Note, however, that the LFC takes time and thus it is possible (although unlikely) that, if the `lfc-interval` is too short, a new cleanup may be started while the previous one is still running. The server would recover from this by skipping the new cleanup when it detected that the previous cleanup was still in progress, but it implies that the actual cleanups will be triggered more rarely than the configured interval. Moreover, triggering a new cleanup adds overhead to the server, which is not able to respond to new requests for a short period of time when the new cleanup process is spawned. Therefore, it is recommended that the `lfc-interval` value be selected in a way that allows the LFC to complete the cleanup before a new cleanup is triggered.

Lease file cleanup is performed by a separate process (in the background) to avoid a performance impact on the server process. To avoid conflicts between two processes using the same lease files, the LFC process starts with Kea opening a new lease file; the actual LFC process operates on the lease file that is no longer used by the server. There are also other files created as a side effect of the lease file cleanup. The detailed description of the LFC process is located later in this Kea Administrator's Reference Manual: *The LFC Process*.

10.2.2.3 Lease Database Configuration

Note

Lease database access information must be configured for the DHCPv6 server, even if it has already been configured for the DHCPv4 server. The servers store their information independently, so each server can use a separate database or both servers can use the same database.

Note

Kea requires the database timezone to match the system timezone. For more details, see *First-Time Creation of the MySQL Database* and *First-Time Creation of the PostgreSQL Database*.

Lease database configuration is controlled through the Dhcp6/lease-database parameters. The database type must be set to memfile, mysql or postgresql, e.g.:

```
"Dhcp6": { "lease-database": { "type": "mysql", ... }, ... }
```

Next, the name of the database to hold the leases must be set; this is the name used when the database was created (see *First-Time Creation of the MySQL Database* or *First-Time Creation of the PostgreSQL Database*).

For MySQL or PostgreSQL:

```
"Dhcp6": { "lease-database": { "name": "database-name" , ... }, ... }
```

If the database is located on a different system from the DHCPv6 server, the database host name must also be specified:

```
"Dhcp6": { "lease-database": { "host": "remote-host-name", ... }, ... }
```

Normally, the database is on the same machine as the DHCPv6 server. In this case, set the value to the empty string:

```
"Dhcp6": { "lease-database": { "host" : "", ... }, ... }
```

Should the database use a port other than the default, it may be specified as well:

```
"Dhcp6": { "lease-database": { "port" : 12345, ... }, ... }
```

Should the database be located on a different system, the administrator may need to specify a longer interval for the connection timeout:

```
"Dhcp6": { "lease-database": { "connect-timeout" : timeout-in-seconds, ... }, ... }
```

The default value of five seconds should be more than adequate for local connections. If a timeout is given, though, it should be an integer greater than zero.

The maximum number of times the server automatically attempts to reconnect to the lease database after connectivity has been lost may be specified:

```
"Dhcp6": { "lease-database": { "max-reconnect-tries" : number-of-tries, ... }, ... }
```

If the server is unable to reconnect to the database after making the maximum number of attempts, the server will exit. A value of 0 (the default) disables automatic recovery and the server will exit immediately upon detecting a loss of connectivity (MySQL and PostgreSQL only).

The number of milliseconds the server waits between attempts to reconnect to the lease database after connectivity has been lost may also be specified:

```
"Dhcp6": { "lease-database": { "reconnect-wait-time" : number-of-milliseconds, ... }, ... }
→ }
```

The default value for MySQL and PostgreSQL is 0, which disables automatic recovery and causes the server to exit immediately upon detecting the loss of connectivity.

```
"Dhcp6": { "lease-database": { "on-fail" : "stop-retry-exit", ... }, ... }
```

The possible values are:

- **stop-retry-exit** - disables the DHCP service while trying to automatically recover lost connections, and shuts down the server on failure after exhausting **max-reconnect-tries**. This is the default value for the lease backend, the host backend, and the configuration backend.

- `serve-retry-exit` - continues the DHCP service while trying to automatically recover lost connections, and shuts down the server on failure after exhausting `max-reconnect-tries`.
- `serve-retry-continue` - continues the DHCP service and does not shut down the server even if the recovery fails. This is the default value for forensic logging.

Note

Automatic reconnection to database backends is configured individually per backend; this allows users to tailor the recovery parameters to each backend they use. We suggest that users enable it either for all backends or none, so behavior is consistent.

Losing connectivity to a backend for which reconnection is disabled results (if configured) in the server shutting itself down. This includes cases when the lease database backend and the hosts database backend are connected to the same database instance.

It is highly recommended not to change the `stop-retry-exit` default setting for the lease manager, as it is critical for the connection to be active while processing DHCP traffic. Change this only if the server is used exclusively as a configuration tool.

```
"Dhcp6": { "lease-database": { "retry-on-startup" : true, ... }, ... }
```

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server starts reconnection attempts even at server startup or on reconfigure events, and honors the action specified in the `on-fail` parameter.

The host parameter is used by the MySQL and PostgreSQL backends.

Finally, the credentials of the account under which the server will access the database should be set:

```
"Dhcp6": {
  "lease-database": {
    "user": "user-name",
    "password": "1234",
    ...
  },
  ...
}
```

If there is no password to the account, set the password to the empty string `""`. (This is the default.)

10.2.2.4 Tuning Database Timeouts

In rare cases, reading or writing to the database may hang. This can be caused by a temporary network issue, or by mis-configuration of the proxy server switching the connection between different database instances. These situations are rare, but users have reported that Kea sometimes hangs while performing database IO operations. Setting appropriate timeout values can mitigate such issues.

MySQL exposes two distinct connection options to configure the read and write timeouts. Kea's corresponding `read-timeout` and `write-timeout` configuration parameters specify the timeouts in seconds. For example:

```
"Dhcp6": { "lease-database": { "read-timeout" : 10, "write-timeout": 20, ... }, ... }
```

Setting these parameters to 0 is equivalent to not specifying them, and causes the Kea server to establish a connection to the database with the MySQL defaults. In this case, Kea waits indefinitely for the completion of the read and write operations.

MySQL versions earlier than 5.6 do not support setting timeouts for read and write operations. Moreover, the `read-timeout` and `write-timeout` parameters can only be specified for the MySQL backend; setting them for any other backend database type causes a configuration error.

To set a timeout in seconds for PostgreSQL, use the `tcp-user-timeout` parameter. For example:

```
"Dhcp6": { "lease-database": { "tcp-user-timeout" : 10, ... }, ... }
```

Specifying this parameter for other backend types causes a configuration error.

Note

The timeouts described here are only effective for TCP connections. Please note that the MySQL client library used by the Kea servers typically connects to the database via a UNIX domain socket when the `host` parameter is `localhost`, but establishes a TCP connection for `127.0.0.1`.

Since Kea.2.7.4, the `libdhcp_mysql.so` hook library must be loaded in order to store leases in the MySQL Lease Database Backend. Specify the lease backend hook library location:

```
"Dhcp6": { "hooks-libraries": [
  {
    // the MySQL lease backend hook library required for lease storage.
    "library": "/opt/lib/kea/hooks/libdhcp_mysql.so"
  }, ... ], ... }
```

Since Kea.2.7.4, the `libdhcp_pgsql.so` hook library must be loaded in order to store leases in the PostgreSQL Lease Database Backend. Specify the lease backend hook library location.

```
"Dhcp6": { "hooks-libraries": [
  {
    // the PostgreSQL lease backend hook library required for lease storage.
    "library": "/opt/lib/kea/hooks/libdhcp_pgsql.so"
  }, ... ], ... }
```

10.2.3 Hosts Storage

Kea is also able to store information about host reservations in the database. The hosts database configuration uses the same syntax as the lease database. In fact, the Kea server opens independent connections for each purpose, be it lease or hosts information, which gives the most flexibility. Kea can keep leases and host reservations separately, but can also point to the same database. Currently the supported hosts database types are MySQL and PostgreSQL.

The following configuration can be used to configure a connection to MySQL:

```
"Dhcp6": {
  "hosts-database": {
    "type": "mysql",
    "name": "kea",
    "user": "kea",
    "password": "1234",
    "host": "localhost",
```

(continues on next page)

(continued from previous page)

```

    "port": 3306
  }
}

```

Depending on the database configuration, many of the parameters may be optional.

Please note that usage of hosts storage is optional. A user can define all host reservations in the configuration file, and that is the recommended way if the number of reservations is small. However, when the number of reservations grows, it is more convenient to use host storage. Please note that both storage methods (the configuration file and one of the supported databases) can be used together. If hosts are defined in both places, the definitions from the configuration file are checked first and external storage is checked later, if necessary.

Host information can be placed in multiple stores. Operations are performed on the stores in the order they are defined in the configuration file, although this leads to a restriction in ordering in the case of a host reservation addition; read-only stores must be configured after a (required) read-write store, or the addition will fail.

Note

Kea requires the database timezone to match the system timezone. For more details, see *First-Time Creation of the MySQL Database* and *First-Time Creation of the PostgreSQL Database*.

10.2.3.1 DHCPv6 Hosts Database Configuration

Hosts database configuration is controlled through the `Dhcp6/hosts-database` parameters. If enabled, the type of database must be set to `mysql` or `postgresql`.

```
"Dhcp6": { "hosts-database": { "type": "mysql", ... }, ... }
```

Next, the name of the database to hold the reservations must be set; this is the name used when the lease database was created (see *Supported Backends* for instructions on how to set up the desired database type):

```
"Dhcp6": { "hosts-database": { "name": "database-name", ... }, ... }
```

If the database is located on a different system than the DHCPv6 server, the database host name must also be specified:

```
"Dhcp6": { "hosts-database": { "host": remote-host-name, ... }, ... }
```

Normally, the database is on the same machine as the DHCPv6 server. In this case, set the value to the empty string:

```
"Dhcp6": { "hosts-database": { "host" : "", ... }, ... }
```

Should the database use a port different than the default, it may be specified as well:

```
"Dhcp6": { "hosts-database": { "port" : 12345, ... }, ... }
```

The maximum number of times the server automatically attempts to reconnect to the host database after connectivity has been lost may be specified:

```
"Dhcp6": { "hosts-database": { "max-reconnect-tries" : number-of-tries, ... }, ... }
```

If the server is unable to reconnect to the database after making the maximum number of attempts, the server will exit. A value of 0 (the default) disables automatic recovery and the server will exit immediately upon detecting a loss of connectivity (MySQL and PostgreSQL only).

The number of milliseconds the server waits between attempts to reconnect to the host database after connectivity has been lost may also be specified:

```
"Dhcp6": { "hosts-database": { "reconnect-wait-time" : number-of-milliseconds, ... }, ...
↪ }
```

The default value for MySQL and PostgreSQL is 0, which disables automatic recovery and causes the server to exit immediately upon detecting the loss of connectivity.

```
"Dhcp6": { "hosts-database": { "on-fail" : "stop-retry-exit", ... }, ... }
```

The possible values are:

- **stop-retry-exit** - disables the DHCP service while trying to automatically recover lost connections. Shuts down the server on failure after exhausting `max-reconnect-tries`. This is the default value for MySQL and PostgreSQL.
- **serve-retry-exit** - continues the DHCP service while trying to automatically recover lost connections. Shuts down the server on failure after exhausting `max-reconnect-tries`.
- **serve-retry-continue** - continues the DHCP service and does not shut down the server even if the recovery fails.

Note

Automatic reconnection to database backends is configured individually per backend. This allows users to tailor the recovery parameters to each backend they use. We suggest that users enable it either for all backends or none, so behavior is consistent.

Losing connectivity to a backend for which reconnection is disabled results (if configured) in the server shutting itself down. This includes cases when the lease database backend and the hosts database backend are connected to the same database instance.

```
"Dhcp6": { "hosts-database": { "retry-on-startup" : true, ... }, ... }
```

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server starts reconnection attempts even at server startup or on reconfigure events, and honors the action specified in the `on-fail` parameter.

Finally, the credentials of the account under which the server will access the database should be set:

```
"Dhcp6": {
  "hosts-database": {
    "user": "user-name",
    "password": "1234",
    ...
  },
  ...
}
```

If there is no password to the account, set the password to the empty string `""`. (This is the default.)

The multiple-storage extension uses a similar syntax; a configuration is placed into a `hosts-databases` list instead of into a `hosts-database` entry, as in:

```
"Dhcp6": { "hosts-databases": [ { "type": "mysql", ... }, ... ], ... }
```

If the same host is configured both in-file and in-database, Kea does not issue a warning, as it would if both were specified in the same data source. Instead, the host configured in-file has priority over the one configured in-database.

10.2.3.2 Using Read-Only Databases for Host Reservations with DHCPv6

In some deployments, the user whose name is specified in the database backend configuration may not have write privileges to the database. This is often required by the policy within a given network to secure the data from being unintentionally modified. In many cases administrators have deployed inventory databases, which contain substantially more information about the hosts than just the static reservations assigned to them. The inventory database can be used to create a view of a Kea hosts database and such a view is often read-only.

Kea host-database backends operate with an implicit configuration to both read from and write to the database. If the user does not have write access to the host database, the backend will fail to start and the server will refuse to start (or reconfigure). However, if access to a read-only host database is required for retrieving reservations for clients and/or assigning specific addresses and options, it is possible to explicitly configure Kea to start in "read-only" mode. This is controlled by the `readonly` boolean parameter as follows:

```
"Dhcp6": { "hosts-database": { "readonly": true, ... }, ... }
```

Setting this parameter to `false` configures the database backend to operate in "read-write" mode, which is also the default configuration if the parameter is not specified.

Note

The `readonly` parameter is only supported for MySQL and PostgreSQL databases.

Since Kea.2.7.4, the `libdhcp_mysql.so` hook library must be loaded in order to store host reservations in the MySQL Host Database Backend. Specify the lease backend hook library location:

```
"Dhcp6": { "hooks-libraries": [
  {
    // the MySQL host backend hook library required for host storage.
    "library": "/opt/lib/kea/hooks/libdhcp_mysql.so"
  }, ... ], ... }
```

Since Kea.2.7.4, the `libdhcp_pgsql.so` hook library must be loaded in order to store host reservations in the PostgreSQL Host Database Backend. Specify the lease backend hook library location.

```
"Dhcp6": { "hooks-libraries": [
  {
    // the PostgreSQL host backend hook library required for host storage.
    "library": "/opt/lib/kea/hooks/libdhcp_pgsql.so"
  }, ... ], ... }
```

10.2.3.3 Tuning Database Timeouts for Hosts Storage

See *Tuning Database Timeouts*.

10.2.4 Interface Configuration

The DHCPv6 server must be configured to listen on specific network interfaces. The simplest network interface configuration tells the server to listen on all available interfaces:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "*" ]
  },
  ...
}
```

The asterisk plays the role of a wildcard and means "listen on all interfaces." However, it is usually a good idea to explicitly specify interface names:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ]
  },
  ...
}
```

It is possible to use an interface wildcard (*) concurrently with explicit interface names:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3", "*" ]
  },
  ...
}
```

This format should only be used when it is desired to temporarily override a list of interface names and listen on all interfaces.

As with the DHCPv4 server, binding to specific addresses and disabling re-detection of interfaces are supported. But `dhcp-socket-type` is not supported, because DHCPv6 uses only UDP/IPv6 sockets. The following example shows how to disable interface detection:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "re-detect": false
  },
  ...
}
```

The loopback interfaces (i.e. the `lo` or `lo0` interface) are not configured by default, unless explicitly mentioned in the configuration. Note that Kea requires a link-local address (which does not exist on all systems) or a specified unicast address, as in:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "enp0s2/2001:db8::1234:abcd" ]
  },
  ...
}
```

Kea binds the service sockets for each interface on startup. If another process is already using a port, then Kea logs the message and suppresses an error. DHCP service runs, but it is unavailable on some interfaces.

The "service-sockets-require-all" option makes Kea require all sockets to be successfully bound. If any opening fails, Kea interrupts the initialization and exits with a non-zero status. (Default is false).

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "service-sockets-require-all": true
  },
  ...
}
```

Sometimes, immediate interruption isn't a good choice. The port can be unavailable only temporary. In this case, retrying the opening may resolve the problem. Kea provides two options to specify the retrying: `service-sockets-max-retries` and `service-sockets-retry-wait-time`.

The first defines a maximal number of retries that Kea makes to open a socket. The zero value (default) means that the Kea doesn't retry the process.

The second defines a wait time (in milliseconds) between attempts. The default value is 5000 (5 seconds).

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1", "eth3" ],
    "service-sockets-max-retries": 5,
    "service-sockets-retry-wait-time": 5000
  },
  ...
}
```

If "service-sockets-max-retries" is non-zero and "service-sockets-require-all" is false, then Kea retries the opening (if needed) but does not fail if any socket is still not opened.

10.2.5 IPv6 Subnet Identifier

The subnet identifier (subnet ID) is a unique number associated with a particular subnet. In principle, it is used to associate clients' leases with their respective subnets. The server configuration must contain unique and stable identifiers for all subnets.

Note

Subnet IDs must be greater than zero and less than 4294967295.

The following configuration assigns the specified subnet identifier to a newly configured subnet:

```
"Dhcp6": {
  "subnet6": [
    {
      "subnet": "2001:db8:1::/64",
      "id": 1024,
      ...
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
]
}
```

10.2.6 IPv6 Subnet Prefix

The subnet prefix is the second way to identify a subnet. Kea can accept non-canonical subnet addresses; for instance, this configuration is accepted:

```
"Dhcp6": {
  "subnet6": [
    {
      "subnet": "2001:db8:1::1/64",
      ...
    }
  ]
}
```

This works even if there is another subnet with the "2001:db8:1::/64" prefix; only the textual form of subnets are compared to avoid duplicates.

Note

Abuse of this feature can lead to incorrect subnet selection (see *IPv6 Subnet Selection*).

10.2.7 Unicast Traffic Support

When the DHCPv6 server starts, by default it listens to the DHCP traffic sent to multicast address ff02::1:2 on each interface that it is configured to listen on (see *Interface Configuration*). In some cases it is useful to configure a server to handle incoming traffic sent to global unicast addresses as well; the most common reason for this is to have relays send their traffic to the server directly. To configure the server to listen on a specific unicast address, add a slash (/) after the interface name, followed by the global unicast address on which the server should listen. The server will listen to this address in addition to normal link-local binding and listening on the ff02::1:2 address. The sample configuration below shows how to listen on 2001:db8::1 (a global address) configured on the eth1 interface.

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "eth1/2001:db8::1" ]
  },
  "option-data": [
    {
      "name": "unicast",
      "data": "2001:db8::1"
    }
  ],
  ...
}
```

This configuration will cause the server to listen on eth1 on the link-local address, the multicast group (ff02::1:2), and 2001:db8::1.

Usually, unicast support is associated with a server unicast option which allows clients to send unicast messages to the server. The example above includes a server unicast option specification which causes the client to send messages to the specified unicast address.

It is possible to mix interface names, wildcards, and interface names/addresses in the list of interfaces. It is not possible, however, to specify more than one unicast address on a given interface.

Care should be taken to specify proper unicast addresses, as the server will attempt to bind to the addresses specified without any additional checks. This approach was selected intentionally, to allow the software to communicate over uncommon addresses if so desired.

10.2.8 Configuration of IPv6 Address Pools

The main role of a DHCPv6 server is address assignment. For this, the server must be configured with at least one subnet and one pool of dynamic addresses to be managed. For example, assume that the server is connected to a network segment that uses the 2001:db8:1::/64 prefix. The administrator of that network decides that addresses from the range 2001:db8:1::1 to 2001:db8:1::ffff are going to be managed by the DHCPv6 server. Such a configuration can be achieved in the following way:

```
"Dhcp6": {
  "subnet6": [
    {
      "subnet": "2001:db8:1::/64",
      "pools": [
        {
          "pool": "2001:db8:1::1-2001:db8:1::ffff"
        }
      ],
      ...
    }
  ]
}
```

Note that `subnet` is defined as a simple string, but the `pools` parameter is actually a list of pools; for this reason, the pool definition is enclosed in square brackets, even though only one range of addresses is specified.

Each `pool` is a structure that contains the parameters that describe a single pool. Currently there is only one parameter, `pool`, which gives the range of addresses in the pool.

It is possible to define more than one pool in a subnet; continuing the previous example, further assume that 2001:db8:1:0:5::/80 should also be managed by the server. It could be written as 2001:db8:1:0:5:: to 2001:db8:1:5:ffff:ffff:ffff, but typing so many `f` characters is cumbersome. The pool can be expressed more simply as 2001:db8:1:0:5::/80. Both formats are supported by `Dhcp6` and they can be mixed in the pool list. For example, the following pools could be defined:

```
"Dhcp6": {
  "subnet6": [
    {
      "subnet": "2001:db8:1::/64",
      "pools": [
        { "pool": "2001:db8:1::1-2001:db8:1::ffff" },
        { "pool": "2001:db8:1:05::/80" }
      ],
      ...
    }
  ]
}
```

White space in pool definitions is ignored, so spaces before and after the hyphen are optional. They can be used to improve readability.

The number of pools is not limited, but for performance reasons it is recommended to use as few as possible.

The server may be configured to serve more than one subnet. To add a second subnet, use a command similar to the following:

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "pools": [
        { "pool": "2001:db8:1::1-2001:db8:1::ffff" }
      ]
    },
    {
      "id": 2,
      "subnet": "2001:db8:2::/64",
      "pools": [
        { "pool": "2001:db8:2::/64" }
      ]
    },
    ...
  ]
}
```

In this example, we allow the server to dynamically assign all addresses available in the whole subnet. Although rather wasteful, it is certainly a valid configuration to dedicate the whole /64 subnet for that purpose. Note that the Kea server does not preallocate the leases, so there is no danger in using gigantic address pools.

When configuring a DHCPv6 server using prefix/length notation, please pay attention to the boundary values. When specifying that the server can use a given pool, it is also able to allocate the first (typically a network address) address from that pool. For example, for pool 2001:db8:2::/64, the 2001:db8:2:: address may be assigned as well. To avoid this, use the min-max notation.

10.2.9 Subnet and Prefix Delegation Pools

Subnets may also be configured to delegate prefixes, as defined in [RFC 8415](#), section 6.3. A subnet may have one or more prefix delegation pools. Each pool has a prefixed address, which is specified as a prefix (`prefix`) and a prefix length (`prefix-len`), as well as a delegated prefix length (`delegated-len`). The delegated length must not be shorter than (i.e. it must be numerically greater than or equal to) the prefix length. If both the delegated and prefix lengths are equal, the server will be able to delegate only one prefix. The delegated prefix does not have to match the subnet prefix.

Below is a sample subnet configuration which enables prefix delegation for the subnet:

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:d8b:1::/64",
      "pd-pools": [
        {
          "prefix": "3000:1::",
          "prefix-len": 64,
          "delegated-len": 96
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ]
  },
  ...
}

```

10.2.10 Prefix Exclude Option

For each delegated prefix, the delegating router may choose to exclude a single prefix out of the delegated prefix as specified in [RFC 6603](#). The requesting router must not assign the excluded prefix to any of its downstream interfaces. The excluded prefix is intended to be used on a link through which the delegating router exchanges DHCPv6 messages with the requesting router. The configuration example below demonstrates how to specify an excluded prefix within a prefix pool definition. The excluded prefix `2001:db8:1:8000:cafe:80::/72` will be sent to a requesting router which includes the Prefix Exclude option in the Option Request option (ORO), and which is delegated a prefix from this pool.

```

"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pd-pools": [
        {
          "prefix": "2001:db8:1:8000::",
          "prefix-len": 56,
          "delegated-len": 64,
          "excluded-prefix": "2001:db8:1:8000:cafe:80::",
          "excluded-prefix-len": 72
        }
      ]
    }
  ]
}

```

Note

Here are some liberties and limits to the values that subnets and pools can take in Kea configurations that are out of the ordinary:

Kea configuration case	allow	Comment
Overlapping subnets	Yes	Administrator consideration needs to be given to how clients are matched to these subnets.
Overlapping address pools in one subnet	No	Startup error: DHCP6_PARSER_FAIL
Overlapping address pools in different subnets	Yes	Specifying the same address pool in different subnets can be used as an equivalent of the global address pool. In that case, the server can assign addresses from the same range regardless of the client's subnet. If an address from such a pool is assigned to a client in one subnet, the same address will be renewed for this client if it moves to another subnet. Another client in a different subnet will not be assigned an address already assigned to the client in any of the subnets.
Address pools that are outside the subnet they are configured under	No	Startup error: DHCP6_PARSER_FAIL
Overlapping prefix delegation pools in one subnet	No	Startup error: DHCP6_PARSER_FAIL
Overlapping prefix delegation pools in different subnets	Yes	Specifying the same prefix delegation pool in different subnets can be used as an equivalent of the global pool. In that case, the server can delegate the same prefixes regardless of the client's subnet. If a prefix from such a pool is delegated to a client in one subnet, the same prefix will be renewed for this client if it moves to another subnet. Another client in a different subnet will not be delegated a prefix already delegated to the client in any of the subnets.
Prefix delegation pools not matching the subnet prefix	Yes	It is common in many deployments to configure the prefix delegation pools not matching the subnet prefix, e.g. a prefix pool of 3000::/96 within the 2001:db8:1::/64 subnet. Such use cases are supported by the Kea DHCPv6 server.

10.2.11 Standard DHCPv6 Options

One of the major features of the DHCPv6 server is the ability to provide configuration options to clients. Although there are several options that require special behavior, most options are sent by the server only if the client explicitly requests them. The following example shows how to configure the addresses of DNS servers, one of the most frequently used options. Options specified in this way are considered global and apply to all configured subnets.

```
"Dhcp6": {
  "option-data": [
    {
      "name": "dns-servers",
```

(continues on next page)

(continued from previous page)

```

        "code": 23,
        "space": "dhcp6",
        "csv-format": true,
        "data": "2001:db8::cafe, 2001:db8::babe"
    },
    ...
]
}

```

The `option-data` line creates a new entry in the option-data table. This table contains information on all global options that the server is supposed to configure in all subnets. The `name` line specifies the option name. (For a complete list of currently supported names, see *List of standard DHCPv6 options configurable by an administrator*.) The next line specifies the option code, which must match one of the values from that list. The line beginning with `space` specifies the option space, which must always be set to `dhcp6` as these are standard DHCPv6 options. For other name spaces, including custom option spaces, see *Nested DHCPv6 Options (Custom Option Spaces)*. The following line specifies the format in which the data will be entered; use of CSV (comma-separated values) is recommended. Finally, the `data` line gives the actual value to be sent to clients. The data parameter is specified as normal text, with values separated by commas if more than one value is allowed.

Options can also be configured as hexadecimal values. If `csv-format` is set to `false`, the option data must be specified as a hexadecimal string. The following commands configure the `dns-servers` option for all subnets with the addresses `2001:db8:1::cafe` and `2001:db8:1::babe`.

```

"Dhcp6": {
  "option-data": [
    {
      "name": "dns-servers",
      "code": 23,
      "space": "dhcp6",
      "csv-format": false,
      "data": "20 01 0D B8 00 01 00 00 00 00 00 00 00 00 CA FE \
              20 01 0D B8 00 01 00 00 00 00 00 00 00 00 BA BE"
    },
    ...
  ]
}

```

Note

The value for the setting of the `data` element is split across two lines in this example for clarity; when entering the command, the whole string should be entered on the same line.

Kea supports the following formats when specifying hexadecimal data:

- **Delimited octets** - one or more octets separated by either colons or spaces (":" or " "). While each octet may contain one or two digits, we strongly recommend always using two digits. Valid examples are "ab:cd:ef" and "ab cd ef".
- **String of digits** - a continuous string of hexadecimal digits with or without a "0x" prefix. Valid examples are "0xabcd" and "abcdef".

Care should be taken to use proper encoding when using hexadecimal format; Kea's ability to validate data correctness in hexadecimal is limited.

It is also possible to specify data for binary options as a single-quoted text string within double quotes, as shown (note that `csv-format` must be set to `false`):

```
"Dhcp6": {
  "option-data": [
    {
      "name": "subscriber-id",
      "code": 38,
      "space": "dhcp6",
      "csv-format": false,
      "data": "'convert this text to binary'"
    },
    ...
  ],
  ...
}
```

Most of the parameters in the `option-data` structure are optional and can be omitted in some circumstances, as discussed in *Unspecified Parameters for DHCPv6 Option Configuration*. Only one of `name` or `code` is required; it is not necessary to specify both. `Space` has a default value of `dhcp6`, so this can be skipped as well if a regular (not encapsulated) DHCPv6 option is defined. Finally, `csv-format` defaults to `true`, so it too can be skipped, unless the option value is specified as hexstring. Therefore, the above example can be simplified to:

```
"Dhcp6": {
  "option-data": [
    {
      "name": "dns-servers",
      "data": "2001:db8::cafe, 2001:db8::babe"
    },
    ...
  ]
}
```

Defined options are added to the response when the client requests them, as well as any options required by a protocol. An administrator can also specify that an option is always sent, even if a client did not specifically request it. To enforce the addition of a particular option, set the `always-send` flag to `true`, as in:

```
"Dhcp6": {
  "option-data": [
    {
      "name": "dns-servers",
      "data": "2001:db8::cafe, 2001:db8::babe",
      "always-send": true
    },
    ...
  ]
}
```

The effect is the same as if the client added the option code in the Option Request Option (or its equivalent for vendor options), as in:

```
"Dhcp6": {
  "option-data": [
    {
```

(continues on next page)

(continued from previous page)

```

        "name": "dns-servers",
        "data": "2001:db8::cafe, 2001:db8::babe",
        "always-send": true
    },
    ...
],
"subnet6": [
    {
        "subnet": "2001:db8:1::/64",
        "option-data": [
            {
                "name": "dns-servers",
                "data": "2001:db8:1::cafe, 2001:db8:1::babe"
            },
            ...
        ],
        ...
    },
    ...
],
...
}

```

In the example above, the `dns-servers` option respects the global `always-send` flag and is always added to responses, but for subnet `2001:db8:1::/64`, the value is taken from the subnet-level option data specification.

Contrary to `always-send`, if the `never-send` flag is set to `true` for a particular option, the server does not add it to the response. The effect is the same as if the client removed the option code in the Option Request Option (or its equivalent for vendor options):

```

"Dhcp6": {
    "option-data": [
        {
            "name": "dns-servers",
            "data": "2001:db8::cafe, 2001:db8::babe"
        },
        ...
    ],
    "subnet6": [
        {
            "subnet": "2001:db8:1::/64",
            "option-data": [
                {
                    "name": "dns-servers",
                    "never-send": true
                },
                ...
            ],
            ...
        },
        ...
    ],
    ...
}

```

(continues on next page)

(continued from previous page)

}

In the example above, the `dns-server` option is never added to responses on subnet `2001:db8:1::/64`. `never-send` has precedence over `always-send`, so if both are `true` the option is not added.

Note

The `always-send` and `never-send` flags are sticky, meaning they do not follow the usual configuration inheritance rules. Instead, if they are enabled at least once along the configuration inheritance chain, they are applied - even if they are disabled in other places which would normally receive a higher priority. For instance, if one of the flags is enabled in the global scope, but disabled at the subnet level, it is enabled, disregarding the subnet-level setting.

Note

The `never-send` flag is less powerful than *`libdhcp_flex_option.so`*; for instance, it has no effect on options managed by the server itself. Both `always-send` and `never-send` have no effect on options which cannot be requested, for instance from a custom space.

It is possible to override options on a per-subnet basis. If clients connected to most subnets are expected to get the same values of a given option, administrators should use global options; it is possible to override specific values for a small number of subnets. On the other hand, if different values are used in each subnet, it does not make sense to specify global option values; rather, only subnet-specific ones should be set.

The following commands override the global `dns-servers` option for a particular subnet, setting a single DNS server with address `2001:db8:1::3`.

```
"Dhcp6": {
  "subnet6": [
    {
      "option-data": [
        {
          "name": "dns-servers",
          "code": 23,
          "space": "dhcp6",
          "csv-format": true,
          "data": "2001:db8:1::3"
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}
```

In some cases it is useful to associate some options with an address or prefix pool from which a client is assigned a lease. Pool-specific option values override subnet-specific and global option values. If the client is assigned multiple leases from different pools, the server assigns options from all pools from which the leases have been obtained. However, if the particular option is specified in multiple pools from which the client obtains the leases, only one instance of this option is handed out to the client. The server's administrator must not try to prioritize assignment of pool-specific

options by trying to order pool declarations in the server configuration.

The following configuration snippet demonstrates how to specify the `dns-servers` option, which will be assigned to a client only if the client obtains an address from the given pool:

```
"Dhcp6": {
  "subnet6": [
    {
      "pools": [
        {
          "pool": "2001:db8:1::100-2001:db8:1::300",
          "option-data": [
            {
              "name": "dns-servers",
              "data": "2001:db8:1::10"
            }
          ]
        }
      ]
    },
    ...
  ],
  ...
}
```

Options can also be specified in class or host-reservation scope. The current Kea options precedence order is (from most important to least): host reservation, pool, subnet, shared network, class, global.

Note

Beginning with Kea 2.7.4, option inclusion can also be controlled through option class-tagging, see [Option Class-Tagging](#)

When a data field is a string and that string contains the comma (,; U+002C) character, the comma must be escaped with two backslashes (\\,; U+005C). This double escape is required because both the routine splitting of CSV data into fields and JSON use the same escape character; a single escape (\\,) would make the JSON invalid. For example, the string "EST5EDT4,M3.2.0/02:00,M11.1.0/02:00" must be represented as:

```
"Dhcp6": {
  "subnet6": [
    {
      "pools": [
        {
          "option-data": [
            {
              "name": "new-posix-timezone",
              "data": "EST5EDT4\\,M3.2.0/02:00\\,M11.1.0/02:00"
            }
          ]
        }
      ],
      ...
    },
    ...
  ],
  ...
}
```

(continues on next page)

(continued from previous page)

```

    },
    ...
],
...
}

```

Some options are designated as arrays, which means that more than one value is allowed. For example, the option `dns-servers` allows the specification of more than one IPv6 address, enabling clients to obtain the addresses of multiple DNS servers.

Custom DHCPv6 Options describes the configuration syntax to create custom option definitions (formats). Creation of custom definitions for standard options is generally not permitted, even if the definition being created matches the actual option format defined in the RFCs. However, there is an exception to this rule for standard options for which Kea currently does not provide a definition. To use such options, a server administrator must create a definition as described in *Custom DHCPv6 Options* in the `dhcp6` option space. This definition should match the option format described in the relevant RFC, but the configuration mechanism allows any option format as there is currently no way to validate it.

The currently supported standard DHCPv6 options are listed in the table below. "Name" and "Code" are the values that should be used as a name/code in the option-data structures. "Type" designates the format of the data; the meanings of the various types are given in *List of standard DHCP option types*.

Table 1: List of standard DHCPv6 options configurable by an administrator

Name	Code	Type	Array?
preference	7	uint8	false
unicast	12	ipv6-address	false
sip-server-dns	21	fqdn	true
sip-server-addr	22	ipv6-address	true
dns-servers	23	ipv6-address	true
domain-search	24	fqdn	true
nis-servers	27	ipv6-address	true
nisp-servers	28	ipv6-address	true
nis-domain-name	29	fqdn	true
nisp-domain-name	30	fqdn	true
sntp-servers	31	ipv6-address	true
information-refresh-time	32	uint32	false
bcmcs-server-dns	33	fqdn	true
bcmcs-server-addr	34	ipv6-address	true
geoconf-civic	36	record (uint8, uint16, binary)	false
remote-id	37	record (uint32, binary)	false
subscriber-id	38	binary	false
client-fqdn	39	record (uint8, fqdn)	false
pana-agent	40	ipv6-address	true
new-posix-timezone	41	string	false
new-tzdb-timezone	42	string	false
ero	43	uint16	true
lq-query (1)	44	record (uint8, ipv6-address)	false
client-data (1)	45	empty	false
clt-time (1)	46	uint32	false
lq-relay-data (1)	47	record (ipv6-address, binary)	false
lq-client-link (1)	48	ipv6-address	true
v6-lost	51	fqdn	false

continues on next page

Table 1 – continued from previous page

Name	Code	Type	Array?
capwap-ac-v6	52	ipv6-address	true
relay-id	53	binary	false
ntp-server	56	empty	false
v6-access-domain	57	fqdn	false
sip-ua-cs-list	58	fqdn	true
bootfile-url	59	string	false
bootfile-param	60	tuple	true
client-arch-type	61	uint16	true
nii	62	record (uint8, uint8, uint8)	false
aftr-name	64	fqdn	false
erp-local-domain-name	65	fqdn	false
rsoo	66	empty	false
pd-exclude	67	binary	false
rdnss-selection	74	record (ipv6-address, uint8, fqdn)	true
client-linklayer-addr	79	binary	false
link-address	80	ipv6-address	false
solmax-rt	82	uint32	false
inf-max-rt	83	uint32	false
dhcp4o6-server-addr	88	ipv6-address	true
s46-rule	89	record (uint8, uint8, uint8, ipv4-address, ipv6-prefix)	false
s46-br	90	ipv6-address	false
s46-dmr	91	ipv6-prefix	false
s46-v4v6bind	92	record (ipv4-address, ipv6-prefix)	false
s46-portparams	93	record(uint8, psid)	false
s46-cont-mape	94	empty	false
s46-cont-mapt	95	empty	false
s46-cont-lw	96	empty	false
v6-captive-portal	103	string	false
v6-sztp-redirect	136	tuple	true
ipv6-address-andsf	143	ipv6-address	true
v6-dnr	144	record (uint16, uint16, fqdn, binary)	false
addr-reg-enable	148	empty	false

Options marked with (1) have option definitions, but the logic behind them is not implemented. That means that, technically, Kea knows how to parse them in incoming messages or how to send them if configured to do so, but not what to do with them. Since the related RFCs require certain processing, the support for those options is non-functional. However, it may be useful in some limited lab testing; hence the definition formats are listed here.

Some options are more complex to configure than others. In particular, the Softwire46 family of options and Discovery of Network-designated Resolvers (DNR) are discussed in separate sections below.

Kea supports more options than those listed above. The following list is mostly useful for readers who want to understand whether Kea is able to support certain options. The following options are returned by the Kea engine itself and in general should not be configured manually.

Table 2: List of standard DHCPv6 options managed by Kea on its own and not directly configurable by an administrator

Name	Code	Description
client-id	1	Sent by the client; Kea uses it to distinguish between clients.
server-id	2	Sent by clients to request action from a specific server and by the server to identify itself. See <i>Server Identifier in DHCPv6</i> for details.
ia-na	3	A container option that conveys IPv6 addresses (<i>ia-addr</i> options). Kea receives and sends those options using its allocation engine.
ia-ta	4	Conveys temporary addresses. Deprecated feature, not supported.
iaaddr	5	Conveys addresses with lifetimes in <i>ia-na</i> and <i>ia-ta</i> options.
oro	6	ORO (or Option Request Option) is used by clients to request a list of options they are interested in. Kea supports it and sends the requested options back if configured with required options.
elapsed time	8	Sent by clients to identify how long they have been trying to obtain a configuration. Kea uses high values sent by clients as an indicator that something is wrong; this is one of the aspects used in HA to determine if the partner is healthy or not.
relay-msg	9	Used by relays to encapsulate the original client message. Kea uses it when sending back relayed responses to the relay agent.
auth	11	Used to pass authentication information between clients and server. The support for this option is very limited.
status-code	13	An option that the server can attach in case of various failures, such as running out of addresses or not being configured to assign prefixes.
rapid-comm	14	Used to signal the client's willingness to support <i>rapid-commit</i> and the server's acceptance for this configuration. See <i>Rapid Commit</i> for details.
user-class	15	Sent by the client to self-identify the device type. Kea can use this for client classification.
vendor-class	16	Similar to <i>user-class</i> , but vendor-specific.
vendor-opts	17	A vendor-specific container that is used by both the client and the server to exchange vendor-specific options. The logic behind those options varies between vendors. Vendor options are explained in <i>DHCPv6 Vendor-Specific Options</i> .
interface-id	18	May be inserted by the relay agent to identify the interface that the original client message was received on. Kea may be told to use this information to select specific subnets. Also, if specified, Kea echoes this option back, so the relay will know which interface to use to reach the client.
ia-pd	25	A container for conveying Prefix Delegations (PDs) that are being delegated to clients. See <i>Subnet and Prefix Delegation Pools</i> for details.
iaprefix	26	Conveys the IPv6 prefix in the <i>ia-pd</i> option. See <i>Subnet and Prefix Delegation Pools</i> for details.

10.2.12 Common Software46 Options

Software46 options are involved in IPv4-over-IPv6 provisioning by means of tunneling or translation, as specified in [RFC 7598](#). The following sections provide configuration examples of these options.

10.2.12.1 Software46 Container Options

Software46 (S46) container options group rules and optional port parameters for a specified domain. There are three container options specified in the "dhcp6" (top-level) option space: the MAP-E Container option, the MAP-T Container option, and the S46 Lightweight 4over6 Container option. These options only contain the encapsulated options specified below; they do not include any data fields.

To configure the server to send a specific container option along with all encapsulated options, the container option must be included in the server configuration as shown below:

```
"Dhcp6": {
  "option-data": [
    {
      "name": "s46-cont-mape"
    } ],
  ...
}
```

This configuration will cause the server to include the MAP-E Container option to the client. Use `s46-cont-mapt` or `s46-cont-lw` for the MAP-T Container and S46 Lightweight 4over6 Container options, respectively.

All remaining Software46 options described below are included in one of the container options. Thus, they must be included in appropriate option spaces by selecting a space name, which specifies the option where they are supposed to be included.

10.2.12.2 S46 Rule Option

The S46 Rule option is used to convey the Basic Mapping Rule (BMR) and Forwarding Mapping Rule (FMR).

```
{
  "space": "s46-cont-mape-options",
  "name": "s46-rule",
  "data": "128, 0, 24, 192.0.2.0, 2001:db8:1::/64"
}
```

Another possible space value is `s46-cont-mapt-options`.

The S46 Rule option conveys a number of parameters:

- `flags` - an unsigned 8-bit integer, with currently only the most-significant bit specified. It denotes whether the rule can be used for forwarding (128) or not (0).
- `ea-len` - an 8-bit-long Embedded Address length. Allowed values range from 0 to 48.
- `IPv4 prefix length` - an 8-bit-long expression of the prefix length of the Rule IPv4 prefix specified in the `ipv4-prefix` field. Allowed values range from 0 to 32.
- `IPv4 prefix` - a fixed-length 32-bit field that specifies the IPv4 prefix for the S46 rule. The bits in the prefix after a specific number of bits (defined in `prefix4-len`) are reserved, and MUST be initialized to zero by the sender and ignored by the receiver.
- `IPv6 prefix` - a field in prefix/length notation that specifies the IPv6 domain prefix for the S46 rule. The field is padded on the right with zero bits up to the nearest octet boundary, when `prefix6-len` is not evenly divisible by 8.

10.2.12.3 S46 BR Option

The S46 BR option is used to convey the IPv6 address of the Border Relay. This option is mandatory in the MAP-E Container option and is not permitted in the MAP-T and S46 Lightweight 4over6 Container options.

```
{
  "space": "s46-cont-mape-options",
  "name": "s46-br",
  "data": "2001:db8:cafe::1"
}
```

Another possible space value is `s46-cont-lw-options`.

10.2.12.4 S46 DMR Option

The S46 DMR option is used to convey values for the Default Mapping Rule (DMR). This option is mandatory in the MAP-T container option and is not permitted in the MAP-E and S46 Lightweight 4over6 Container options.

```
{
  "space": "s46-cont-mapt-options",
  "name": "s46-dmr",
  "data": "2001:db8:cafe::/64"
}
```

This option must not be included in other containers.

10.2.12.5 S46 IPv4/IPv6 Address Binding Option

The S46 IPv4/IPv6 Address Binding option may be used to specify the full or shared IPv4 address of the Customer Edge (CE). The IPv6 prefix field is used by the CE to identify the correct prefix to use for the tunnel source.

```
{
  "space": "s46-cont-lw",
  "name": "s46-v4v6bind",
  "data": "192.0.2.3, 2001:db8:1:cafe::/64"
}
```

This option must not be included in other containers.

10.2.12.6 S46 Port Parameters

The S46 Port Parameters option specifies optional port-set information that may be provided to CEs.

```
{
  "space": "s46-rule-options",
  "name": "s46-portparams",
  "data": "2, 3/4"
}
```

Another possible `space` value is `s46-v4v6bind`, to include this option in the S46 IPv4/IPv6 Address Binding option.

Note that the second value in the example above specifies the PSID and PSID-length fields in the format of PSID/PSID length. This is equivalent to the values of PSID-len=4 and PSID=12288 conveyed in the S46 Port Parameters option.

10.2.13 DNR (Discovery of Network-designated Resolvers) Options for DHCPv6

The Discovery of Network-designated Resolvers, or DNR option, was introduced in [RFC 9463](#) as a way to communicate location of DNS resolvers available over means other than the classic DNS over UDP over port 53. As of spring 2024, the supported technologies are DoT (DNS-over-TLS), DoH (DNS-over-HTTPS), and DoQ (DNS-over-QUIC), but the option was designed to be extensible to accommodate other protocols in the future.

The DNR option may be configured using convenient notation: comma-delimited fields must be provided in the following order:

- Service Priority (mandatory),
- ADN FQDN (mandatory),
- IP address(es) (optional; if more than one, they must be separated by spaces)
- SvcParams as a set of key=value pairs (optional; if more than one, they must be separated by spaces) To provide more than one `alpn-id`, separate them with double backslash-escaped commas as in the example below).

Let's imagine that we want to convey a DoT server operating at `dot1.example.org` (which resolves to two IPv6 addresses: `2001:db8::1` and `2001:db8::2`) on a non-standard port 8530. An example option that would convey this information looks as follows:

```
{
  "name": "v6-dnr", // name of the option

  // The following fields should be specified:
  // - service priority (unsigned 16-bit integer)
  // - authentication-domain-name (FQDN of the encrypted resolver)
  // - a list of one or more IPv6 addresses
  // - list of parameters in key=value format, space separated; any comma
  //   characters in this field must be escaped with double backslashes
  "data": "100, dot1.example.org., 2001:db8::1 2001:db8::2, alpn=dot port=8530"
}
```

The above option will be encoded on-wire as follows:

```
00 64 - service priority (100 in hex as unsigned 16-bit integer)
00 12 - length of the Authentication Domain Name (name of the resolver) FQDN (18 in hex,
→ as unsigned 16-bit integer)
04 64 6f 74 31 07 65 78 61 6d 70 6c 65 03 6f 72 67 00 - 18 octets of the ADN FQDN
00 20 - 32 octets is the length of the following two IPv6 addresses
20 01 0d b8 00 00 00 00 00 00 00 00 00 00 00 01 - 2001:db8::1
20 01 0d b8 00 00 00 00 00 00 00 00 00 00 00 02 - 2001:db8::2
00 01 - SvsParams begin - this is alpn SvcParamKey
00 04 - length of the alpn SvcParamValue field (4 octets)
03 - length of the following alpn-id coded on one octet
64 6f 74 - "dot" - value of the alpn
00 03 - this is port SvcParamKey
00 02 - length of the SvcParamValue field is 2 octets
21 52 - the actual value is 0x2152 or 8530 in decimal
```

The following example shows how to configure more than one ALPN protocol in Service Parameters. The example specifies a resolver known as `resolver.example` that supports:

- DoT on default port 853
- DoQ on default port 853
- DoH at `https://resolver.example/q{?dns}`

```
{
  "name": "v6-dnr", // name of the option

  // Note the double backslash-escaped commas in the alpn-id list.
  "data": "150, resolver.example., 2001:db8::1 2001:db8::2, alpn=dot\\,doq\\,h2\\,h3\\,
→ dohpath=/q{?dns}"
}
```

The above option will be encoded on-wire as follows:

```
00 96 - service priority (150 in hex as unsigned 16-bit integer)
00 12 - length of the Authentication Domain Name (name of the resolver) FQDN (18 in hex,
→ as unsigned 16-bit integer)
08 72 65 73 6f 6c 76 65 72 07 65 78 61 6d 70 6c 65 00 - 18 octets of the ADN FQDN
```

(continues on next page)

(continued from previous page)

```

00 20 - 32 octets is the length of the following two IPv6 addresses
20 01 0d b8 00 00 00 00 00 00 00 00 00 00 01 - 2001:db8::1
20 01 0d b8 00 00 00 00 00 00 00 00 00 00 02 - 2001:db8::2
00 01 - SvcParams begin - this is the alpn SvcParamKey
00 0e - length of the alpn SvcParamValue field (14 octets)
03    - length of the following alpn-id coded on one octet
64 6f 74 - "dot" - value of the alpn
03    - length of the following alpn-id coded on one octet
64 6f 71 - "doq" - value of the alpn
02    - length of the following alpn-id coded on one octet
68 32 - "h2" - value of the alpn "HTTP/2 over TLS"
02    - length of the following alpn-id coded on one octet
68 33 - "h3" - value of the alpn "HTTP/3"
00 07 - this is dohpath SvcParamKey
00 08 - length of the SvcParamValue field is 8 octets
2f 71 7b 3f 64 6e 73 7d - "/q{?dns}" dohpath

```

Note

If "comma" or "pipe" characters are used as text rather than as field delimiters, they must be escaped with double backslashes (\\, or \\|). Escaped commas must be used when configuring more than one ALPN protocol, to separate them. The "pipe" (0x7C) character can be used in the dohpath service parameter, as it is allowed in a URI.

[RFC 9463](#), Section 4.1 encourages the use of the ALPN (Application-Layer Protocol Negotiation) SvcParam, as it is required in most cases. It defines the protocol for reaching the encrypted resolver. The most common values are dot, doq, and h2 (meaning HTTP/2.0 over TLS, used in DoH).

Per [RFC 9461](#) Section 5: if the alpn SvcParam indicates support for HTTP, dohpath **MUST** be present. The URI Template **MUST** contain a "dns" variable. For example, when advertising a DoH resolver available at `https://doh1.example.org/query{?dns}`, the dohpath should be set to relative URI `/query{?dns}`.

Users interested in configuring this option are encouraged to read the following materials:

- A very nice set of examples is available in Section 7 of [RFC 9461](#).
- A list of all currently defined service parameters is maintained in the [IANA registry](#). This specifies records that can be stored in the svcParams field of the DNR option.
- A list of currently allowed protocols in the ALPN parameter is maintained in [another IANA registry](#).
- [RFC 9463](#) provides option definitions. In terms of SvcParams, it states that alpn and port must be supported, and support for dohpath (used for DoH) is recommended.
- Section 2.2 of [RFC 9460](#) defines the on-wire format for SvcParams.
- Sections 7.1 and 7.2 of [RFC 9460](#) define the on-wire format for alpn and port.
- Section 5 of [RFC 9461](#) defines the on-wire format for dohpath.

Kea currently supports the following service parameters:

Name	Code	Description
alpn	1	Specifies comma-separated protocol types (DoT, DoH, etc.)
port	3	Unsigned 16-bit integer. Indicates a non-standard TCP or UDP port.
dohpath	7	Mandatory for DoH. Contains URL path for the DoT resolver.

The other currently defined service parameters mandatory (0), no-default-alpn (2), ipv4hint (4), ech (5), ipv6hint (6), and ohhttp (8) are not usable in the DNR option.

Further examples are provided in Kea sources in the `all-options.json` file in the `doc/examples/kea6` directory. The DHCPv4 option is nearly identical, and is described in *DNR (Discovery of Network-designated Resolvers) Options for DHCPv4*.

10.2.14 NTP Server Suboptions

NTP server option is a container of suboptions: `ntp-server-address` (1), `ntp-server-multicast` (2) carrying an IPv6 address, and `ntp-server-fqdn` (3) carrying a FQDN in wire format defined in the "v6-ntp-server-suboptions" option space. Each option instance carries one and only one suboption as required by [RFC 5908](#).

10.2.15 Custom DHCPv6 Options

Kea supports custom (non-standard) DHCPv6 options. Let's say that we want to define a new DHCPv6 option called `foo`, which will have code 100 and will convey a single, unsigned, 32-bit integer value. Such an option can be defined by putting the following entry in the configuration file:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "foo",
      "code": 100,
      "type": "uint32",
      "array": false,
      "record-types": "",
      "space": "dhcp6",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

The `false` value of the `array` parameter determines that the option does NOT comprise an array of `uint32` values but is, instead, a single value. Two other parameters have been left blank: `record-types` and `encapsulate`. The former specifies the comma-separated list of option data fields, if the option comprises a record of data fields. The `record-types` value should be non-empty if `type` is set to `record`; otherwise it must be left blank. The latter parameter specifies the name of the option space being encapsulated by the particular option. If the particular option does not encapsulate any option space, the parameter should be left blank. Note that the `option-def` configuration statement only defines the format of an option and does not set its value(s).

The `name`, `code`, and `type` parameters are required; all others are optional. The `array` parameter's default value is `false`. The `record-types` and `encapsulate` parameters' default values are blank (`""`). The default `space` is `dhcp6`.

Once the new option format is defined, its value is set in the same way as for a standard option. For example, the following commands set a global value that applies to all subnets.

```
"Dhcp6": {
  "option-data": [
    {
      "name": "foo",
      "code": 100,
      "space": "dhcp6",
      "csv-format": true,
      "data": "12345"
    },
    ...
  ],
  ...
}
```

New options can take more complex forms than the simple use of primitives (uint8, string, ipv6-address, etc.); it is possible to define an option comprising a number of existing primitives.

For example, say we want to define a new option that will consist of an IPv6 address, followed by an unsigned 16-bit integer, followed by a boolean value, followed by a text string. Such an option could be defined in the following way:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "bar",
      "code": 101,
      "space": "dhcp6",
      "type": "record",
      "array": false,
      "record-types": "ipv6-address, uint16, boolean, string",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

The `type` parameter is set to `"record"` to indicate that the option contains multiple values of different types. These types are given as a comma-separated list in the `record-types` field and should be ones from those listed in *List of standard DHCP option types*.

The values of the options are set in an `option-data` statement as follows:

```
"Dhcp6": {
  "option-data": [
    {
      "name": "bar",
      "space": "dhcp6",
      "code": 101,
      "csv-format": true,
      "data": "2001:db8:1::10, 123, false, Hello World"
    }
  ],
  ...
}
```

The `csv-format` parameter is set to `true` to indicate that the data field comprises a comma-separated list of values. The values in data must correspond to the types set in the `record-types` field of the option definition.

When `array` is set to `true` and `type` is set to `"record"`, the last field is an array, i.e. it can contain more than one value, as in:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "bar",
      "code": 101,
      "space": "dhcp6",
      "type": "record",
      "array": true,
      "record-types": "ipv6-address, uint16",
      "encapsulate": ""
    },
    ...
  ],
  ...
}
```

The new option content is one IPv6 address followed by one or more 16-bit unsigned integers.

Note

In general, boolean values are specified as `true` or `false`, without quotes. Some specific boolean parameters may also accept `"true"`, `"false"`, `0`, `1`, `"0"`, and `"1"`.

10.2.16 DHCPv6 Vendor-Specific Options

Vendor options in DHCPv6 are carried in the Vendor-Specific Information option (code 17). The idea behind option 17 is that each vendor has its own unique set of options with their own custom formats. The vendor is identified by a 32-bit unsigned integer called `enterprise-number` or `vendor-id`.

The standard spaces defined in Kea and their options are:

- `vendor-2495`: Internet Systems Consortium, Inc. for 4o6 options:

option code	option name	option description
60000	4o6-interface	the name of the 4o6 server's client-facing interface
60001	4o6-source-address	the address that the 4o6 server uses to send packets to the client
60002	4o6-source-port	the port that the 4o6 server opens to send packets to the client

- `vendor-4491`: Cable Television Laboratories, Inc. for DOCSIS3 options:

option code	option name	option description
1	oro	ORO (or Option Request Option) is used by clients to request a list of options they are interested in.
2	tftp-servers	a list of IPv4 addresses of TFTP servers to be used by the cable modem

The following examples show how to define an option "foo" with code 1 that consists of an IPv6 address, an unsigned 16-bit integer, and a string. The "foo" option is conveyed in a Vendor-Specific Information option, which comprises a single uint32 value that is set to 12345. The sub-option "foo" follows the data field holding this value.

The first step is to define the format of the option:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "foo",
      "code": 1,
      "space": "vendor-12345",
      "type": "record",
      "array": false,
      "record-types": "ipv6-address, uint16, string",
      "encapsulate": ""
    }
  ],
  ...
}
```

Note that the option space is set to "vendor-12345". Once the option format is defined, the next step is to define actual values for that option:

```
"Dhcp6": {
  "option-data": [
    {
      "name": "foo",
      "space": "vendor-12345",
      "data": "2001:db8:1::10, 123, Hello World"
    },
    ...
  ],
  ...
}
```

We should also define a value ("enterprise-number") for the Vendor-Specific Information option, to convey the option foo.

```
"Dhcp6": {
  "option-data": [
    {
      "name": "vendor-opts",
      "data": "12345"
    },
    ...
  ],
  ...
}
```

Alternatively, the option can be specified using its code.

```
"Dhcp6": {
  "option-data": [
    {
```

(continues on next page)

(continued from previous page)

```

        "code": 17,
        "data": "12345"
    },
    ...
],
...
}

```

A common configuration is to set the `always-send` flag to `true`, so the vendor option is sent even when the client did not specify it in the query.

This is also how `kea-dhcp6` can be configured to send multiple vendor options from different vendors, along with each of their specific enterprise numbers. To send these options regardless of whether the client specifies an enterprise number, the server must be configured with `"always-send": true`, including the Vendor-Specific Information option (code 17).

```

{
  "Dhcp6": {
    "option-data": [
      {
        "always-send": true,
        "data": "tagged",
        "name": "tag",
        "space": "vendor-2234"
      },
      {
        "always-send": true,
        "data": "https://example.com:1234/path",
        "name": "url",
        "space": "vendor-3561"
      }
    ],
    "option-def": [
      {
        "code": 22,
        "name": "tag",
        "space": "vendor-2234",
        "type": "string"
      },
      {
        "code": 11,
        "name": "url",
        "space": "vendor-3561",
        "type": "string"
      }
    ]
  }
}

```

Note

The `kea-dhcp6` server is able to recognize multiple Vendor Class options (code 16) with different enterprise num-

bers in the client requests, and to send multiple Vendor-Specific Information options (code 17) in the responses, one for each vendor.

10.2.17 Nested DHCPv6 Options (Custom Option Spaces)

It is sometimes useful to define a completely new option space: for example, a user might create a new option to convey sub-options that use a separate numbering scheme, such as sub-options with codes 1 and 2. Those option codes conflict with standard DHCPv6 options, so a separate option space must be defined.

Note that the creation of a new option space is not required when defining sub-options for a standard option, because one is created by default if the standard option is meant to convey any sub-options (see *DHCPv6 Vendor-Specific Options*).

If we want a DHCPv6 option called `container` with code 102, that conveys two sub-options with codes 1 and 2, we first need to define the new sub-options:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "subopt1",
      "code": 1,
      "space": "isc",
      "type": "ipv6-address",
      "record-types": "",
      "array": false,
      "encapsulate": ""
    },
    {
      "name": "subopt2",
      "code": 2,
      "space": "isc",
      "type": "string",
      "record-types": "",
      "array": false,
      "encapsulate": ""
    }
  ],
  ...
}
```

Note that we have defined the options to belong to a new option space (in this case, "isc").

The next step is to define a regular DHCPv6 option with the desired code and specify that it should include options from the new option space:

```
"Dhcp6": {
  "option-def": [
    {
      "name": "container",
      "code": 102,
      "space": "dhcp6",
      "type": "empty",
      "array": false,
      "record-types": "",
      "encapsulate": "isc"
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    },
    ...
  ],
  ...
}

```

The name of the option space in which the sub-options are defined is set in the `encapsulate` field. The `type` field is set to `"empty"`, to indicate that this option does not carry any data other than sub-options.

Finally, we can set values for the new options:

```

{
  "Dhcp6": {
    "option-data": [
      {
        "name": "subopt1",
        "code": 1,
        "space": "isc",
        "data": "2001:db8::abcd"
      },
      {
        "name": "subopt2",
        "code": 2,
        "space": "isc",
        "data": "Hello world"
      },
      {
        "name": "container",
        "code": 102,
        "space": "dhcp6"
      }
    ]
  }
}

```

It is possible to create an option which carries some data in addition to the sub-options defined in the encapsulated option space. For example, if the `container` option from the previous example were required to carry a `uint16` value as well as the sub-options, the `type` value would have to be set to `"uint16"` in the option definition. (Such an option would then have the following data structure: DHCP header, `uint16` value, sub-options.) The value specified with the `data` parameter — which should be a valid integer enclosed in quotes, e.g. `"123"` — would then be assigned to the `uint16` field in the `container` option.

10.2.18 Unspecified Parameters for DHCPv6 Option Configuration

In many cases it is not required to specify all parameters for an option configuration, and the default values can be used. However, it is important to understand the implications of not specifying some of them, as it may result in configuration errors. The list below explains the behavior of the server when a particular parameter is not explicitly specified:

- **name** - the server requires either an option name or an option code to identify an option. If this parameter is unspecified, the option code must be specified.
- **code** - the server requires either an option name or an option code to identify an option; this parameter may be left unspecified if the **name** parameter is specified. However, this also requires that the particular option have a definition (either as a standard option or an administrator-created definition for the option using an `option-def` structure), as the option definition associates an option with a particular name. It is possible to configure an

option for which there is no definition (unspecified option format). Configuration of such options requires the use of the option code.

- **space** - if the option space is unspecified it defaults to `dhcp6`, which is an option space holding standard DHCPv6 options.
- **data** - if the option data is unspecified it defaults to an empty value. The empty value is mostly used for the options which have no payload (boolean options), but it is legal to specify empty values for some options which carry variable-length data and for which the specification allows a length of 0. For such options, the data parameter may be omitted in the configuration.
- **csv-format** - if this value is not specified, the server assumes that the option data is specified as a list of comma-separated values to be assigned to individual fields of the DHCP option.

10.2.19 Controlling the Values Sent for T1 and T2 Times

According to RFC 8415, section 21.4, the recommended T1 and T2 values are 50% and 80% of the preferred lease time, respectively. Kea can be configured to send values that are specified explicitly or that are calculated as percentages of the preferred lease time. The server's behavior is determined by a combination of configuration parameters, of which T1 and T2 are only two.

The lease's preferred and valid lifetimes are expressed as triplets with minimum, default, and maximum values using configuration entries:

- **min-preferred-lifetime** - specifies the minimum preferred lifetime (optional).
- **preferred-lifetime** - specifies the default preferred lifetime.
- **max-preferred-lifetime** - specifies the maximum preferred lifetime (optional).
- **min-valid-lifetime** - specifies the minimum valid lifetime (optional).
- **valid-lifetime** - specifies the default valid lifetime.
- **max-valid-lifetime** - specifies the maximum valid lifetime (optional).

These values may be specified within client classes.

When the client does not specify lifetimes, the default is used. A specified lifetime - using the `IAADDR` or `IAPREFIX` sub-option with non-zero values - uses these values when they are between the configured minimum and maximum bounds. Values outside the bounds are rounded up or down as needed.

Note

If the `preferred-lifetime` has not been explicitly specified, or if the specified value is larger than the value of `valid-lifetime`, the server uses the value of `valid-lifetime` multiplied by 0.625.

To send specific fixed values, use the following two parameters:

- **renew-timer** - specifies the value of T1 in seconds.
- **rebind-timer** - specifies the value of T2 in seconds.

Any value greater than or equal to zero may be specified for T2. T1, if specified, must be less than T2. This flexibility allows a use case where administrators want to suppress client renewals and rebinds by deferring them beyond the lifespan of the lease. This should cause the lease to expire, rather than get renewed by clients. If T1 is specified as larger than T2, T1 is silently set to zero in the outbound IA.

In the great majority of cases, the values should follow this rule: $T1 < T2 < \text{preferred lifetime} < \text{valid lifetime}$. Alternatively, both T1 and T2 values can be configured to 0, which is a signal to DHCPv6 clients that they may renew at their

own discretion. However, there are known broken client implementations in use that will start renewing immediately. Administrators who plan to use $T1=T2=0$ values should test first and make sure their clients behave rationally.

In some rare cases there may be a need to disable a client's ability to renew addresses. This is undesired from a protocol perspective and should be avoided if possible. However, if necessary, administrators can configure the $T1$ and $T2$ values to be equal or greater to the valid lifetime. Be advised that this will cause clients to occasionally lose their addresses, which is generally perceived as poor service. However, there may be some rare business cases when this is desired (e.g. when it is desirable to intentionally break long-lasting connections).

Calculation of the values is controlled by the following three parameters:

- `calculate-tee-times` - when true, $T1$ and $T2$ are calculated as percentages of the valid lease time. It defaults to true.
- `t1-percent` - the percentage of the valid lease time to use for $T1$. It is expressed as a real number between 0.0 and 1.0 and must be less than `t2-percent`. The default value is 0.5, per RFC 8415.
- `t2-percent` - the percentage of the valid lease time to use for $T2$. It is expressed as a real number between 0.0 and 1.0 and must be greater than `t1-percent`. The default value is 0.8 per RFC 8415.

Note

Unlike DHCPv4 the tee ($T1$, $T2$) times are always present in DHCPv6 address and prefix options. Therefore the default value for `calculate-tee-times` for *kea-dhcp6* is true. This ensures the server's default behavior will result in non-zero tee times being sent to clients. This is to avoid the server being swamped by misbehaving clients that do not calculate it for themselves.

Note

If both explicit values are specified and `calculate-tee-times` is true, the server will use the explicit values. Administrators with a setup where some subnets or shared-networks use explicit values and some use calculated values must not define the explicit values at any level higher than where they will be used. Inheriting them from too high a scope, such as global, will cause them to have values at every level underneath (both shared-networks and subnets), effectively disabling calculated values.

10.2.20 IPv6 Subnet Selection

The DHCPv6 server may receive requests from local (connected to the same subnet as the server) and remote (connected via relays) clients. As the server may have many subnet configurations defined, it must select an appropriate subnet for a given request.

In IPv4, the server can determine which of the configured subnets are local, as there is a reasonable expectation that the server will have a (global) IPv4 address configured on the interface. That assumption is not true in IPv6; the DHCPv6 server must be able to operate while only using link-local addresses. Therefore, an optional `interface` parameter is available within a subnet definition to designate that a given subnet is local, i.e. reachable directly over the specified interface. For example, a server that is intended to serve a local subnet over `eth0` may be configured as follows:

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:beef::/48",
      "pools": [
        {
```

(continues on next page)

(continued from previous page)

```

        "pool": "2001:db8:beef::/48"
      }
    ],
    "interface": "eth0"
  }
  ...
}

```

10.2.21 Rapid Commit

The Rapid Commit option, described in [RFC 8415](#), is supported by the Kea DHCPv6 server. However, support is disabled by default. It can be enabled on a per-subnet basis using the `rapid-commit` parameter as shown below:

```

{
  "Dhcp6": {
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8:beef::/48",
        "rapid-commit": true,
        "pools": [
          {
            "pool": "2001:db8:beef::1-2001:db8:beef::10"
          }
        ]
      }
    ]
  }
}

```

This setting only affects the subnet for which `rapid-commit` is set to `true`. For clients connected to other subnets, the server ignores the Rapid Commit option sent by the client and follows the 4-way exchange procedure, i.e. responds with an Advertise for a Solicit containing a Rapid Commit option.

10.2.22 DHCPv6 Relays

A DHCPv6 server with multiple subnets defined must select the appropriate subnet when it receives a request from a client. For clients connected via relays, two mechanisms are used:

The first uses the `linkaddr` field in the `RELAY_FORW` message. The name of this field is somewhat misleading in that it does not contain a link-layer address; instead, it holds an address (typically a global address) that is used to identify a link. The DHCPv6 server checks to see whether the address belongs to a defined subnet and, if it does, that subnet is selected for the client's request.

The second mechanism is based on `interface-id` options. While forwarding a client's message, relays may insert an `interface-id` option into the message that identifies the interface on the relay that received the message. (Some relays allow configuration of that parameter, but it is sometimes hard-coded and may range from the very simple [e.g. "vlan100"] to the very cryptic; one example seen on real hardware was "ISAM144|299|ipv6|nt:vp:1:110".) The server can use this information to select the appropriate subnet. The information is also returned to the relay, which then knows the interface to use to transmit the response to the client. For this to work successfully, the relay interface IDs must be unique within the network and the server configuration must match those values.

When configuring the DHCPv6 server, two similarly named parameters can be configured for a subnet:

- `interface` - defines which local network interface can be used to access a given subnet.
- `interface-id` - specifies the content of the `interface-id` option used by relays to identify the interface on the relay to which the response packet is sent.

The two are mutually exclusive; a subnet cannot be reachable both locally (direct traffic) and via relays (remote traffic). Specifying both is a configuration error and the DHCPv6 server will refuse such a configuration.

The following example configuration shows how to specify an `interface-id` with a value of "vlan123":

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:beef::/48",
      "pools": [
        {
          "pool": "2001:db8:beef::/48"
        }
      ],
      "interface-id": "vlan123"
    }
  ],
  ...
}
```

10.2.23 Relay-Supplied Options

[RFC 6422](#) defines a mechanism called Relay-Supplied DHCP Options. In certain cases relay agents are the only entities that may have specific information, and they can insert options when relaying messages from the client to the server. The server then does certain checks and copies those options to the response sent to the client.

There are certain conditions that must be met for the option to be included. First, the server must not provide the option itself; in other words, if both relay and server provide an option, the server always takes precedence. Second, the option must be RSOO-enabled. (RSOO is the "Relay Supplied Options option.") IANA maintains a list of RSOO-enabled options [here](#). However, there may be cases when system administrators want to echo other options. Kea can be instructed to treat other options as RSOO-enabled; for example, to mark options 110, 120, and 130 as RSOO-enabled, the following syntax should be used:

```
"Dhcp6": {
  "relay-supplied-options": [ "110", "120", "130" ],
  ...
}
```

At this time, only option 65 is RSOO-enabled by IANA. This option will always be treated as RSOO-enabled, so there is no need to explicitly mark it. When enabling standard options, it is also possible to use their names rather than their option code, e.g. use `dns-servers` instead of 23. See [ref:dhcp6-std-options-list](#) for the names. In certain cases this may also work for custom options, but due to the nature of the parser code this may be unreliable and should be avoided.

10.2.24 Client Classification in DHCPv6

The DHCPv6 server includes support for client classification. For a deeper discussion of the classification process, see [Client Classification](#).

In certain cases it is useful to configure the server to differentiate between DHCP client types and treat them accordingly. Client classification can be used to modify the behavior of almost any part of DHCP message processing. Kea

currently offers three mechanisms that take advantage of client classification in DHCPv6: subnet selection, address pool selection, and DHCP options assignment.

Kea can be instructed to limit access to given subnets based on class information. This is particularly useful for cases where two types of devices share the same link and are expected to be served from two different subnets. The primary use case for such a scenario is cable networks, where there are two classes of devices: the cable modem itself, which should be handed a lease from subnet A; and all other devices behind the modem, which should get leases from subnet B. That segregation is essential to prevent overly curious end-users from playing with their cable modems. For details on how to set up class restrictions on subnets, see *Configuring Subnets With Class Information*.

When subnets belong to a shared network, the classification applies to subnet selection but not to pools; that is, a pool in a subnet limited to a particular class can still be used by clients which do not belong to the class, if the pool they are expected to use is exhausted. The limit on access based on class information is also available at the address/prefix pool level within a subnet: see *Configuring Pools With Class Information*. This is useful when segregating clients belonging to the same subnet into different address ranges.

In a similar way, a pool can be constrained to serve only known clients, i.e. clients which have a reservation, using the built-in `KNOWN` or `UNKNOWN` classes. Addresses can be assigned to registered clients without giving a different address per reservation: for instance, when there are not enough available addresses. The determination whether there is a reservation for a given client is made after a subnet is selected, so it is not possible to use `KNOWN/UNKNOWN` classes to select a shared network or a subnet.

The process of classification is conducted in five steps. The first step is to assess an incoming packet and assign it to zero or more classes. The second step is to choose a subnet, possibly based on the class information. When the incoming packet is in the special class `DROP`, it is dropped and a debug message logged. The next step is to evaluate class expressions depending on the built-in `KNOWN/UNKNOWN` classes after host reservation lookup, using them for pool/pd-pool selection and assigning classes from host reservations. The list of required classes is then built and each class of the list has its expression evaluated; when it returns `true`, the packet is added as a member of the class. The last step is to assign options, again possibly based on the class information. More complete and detailed information is available in *Client Classification*.

There are two main methods of classification. The first is automatic and relies on examining the values in the vendor class options or the existence of a host reservation. Information from these options is extracted, and a class name is constructed from it and added to the class list for the packet. The second method specifies an expression that is evaluated for each packet. If the result is `true`, the packet is a member of the class.

Note

The new `early-global-reservations-lookup` global parameter flag enables a lookup for global reservations before the subnet selection phase. This lookup is similar to the general lookup described above with two differences:

- the lookup is limited to global host reservations
- the `UNKNOWN` class is never set

Note

Care should be taken with client classification, as it is easy for clients that do not meet class criteria to be denied all service.

10.2.24.1 Defining and Using Custom Classes

The following example shows how to configure a class using an expression and a subnet using that class. This configuration defines the class named `Client_enterprise`. It is comprised of all clients whose client identifiers start with the given hex string (which would indicate a DUID based on an enterprise id of `0xAABBCCDD`). Members of this

class will be given an address from 2001:db8:1::0 to 2001:db8:1::FFFF and the addresses of their DNS servers set to 2001:db8:0::1 and 2001:db8:2::1.

```
"Dhcp6": {
  "client-classes": [
    {
      "name": "Client_enterprise",
      "test": "substring(option[1].hex,0,6) == 0x0002AABBCCDD",
      "option-data": [
        {
          "name": "dns-servers",
          "code": 23,
          "space": "dhcp6",
          "csv-format": true,
          "data": "2001:db8:0::1, 2001:db8:2::1"
        }
      ]
    },
    ...
  ],
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "pools": [ { "pool": "2001:db8:1::-2001:db8:1::ffff" } ],
      "client-classes": [ "Client_enterprise" ]
    }
  ],
  ...
}
```

This example shows a configuration using an automatically generated `VENDOR_CLASS_` class. The administrator of the network has decided that addresses in the range 2001:db8:1::1 to 2001:db8:1::ffff are to be managed by the DHCPv6 server and that only clients belonging to the eRouter1.0 client class are allowed to use that pool.

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "pools": [
        {
          "pool": "2001:db8:1::-2001:db8:1::ffff"
        }
      ],
      "client-classes": [ "VENDOR_CLASS_eRouter1.0" ]
    }
  ],
  ...
}
```

10.2.24.2 Additional Classification

In some cases it is useful to limit the scope of a class to a pool, subnet, or shared network. There are two parameters which are used to limit the scope of the class by instructing the server to evaluate test expressions when required.

The `evaluate-additional-classes`, which takes a list of class names and is valid in pool, subnet, and shared network scope. Classes in these lists are marked as additional and evaluated after selection of this specific pool/subnet/shared network and before output-option processing.

The second one is the `per-class only-in-additional-list` flag, which is `false` by default. When it is set to `true`, the test expression of the class is not evaluated at the reception of the incoming packet but later, and only if the class is present in an `evaluate-additional-classes` list.

In this example, a class is assigned to the incoming packet when the specified subnet is used:

```
"Dhcp6": {
  "client-classes": [
    {
      "name": "Client_foo",
      "test": "member('ALL')",
      "only-in-additional-list": true
    },
    ...
  ],
  "subnet6": [
    {
      "subnet": "2001:db8:1::/64",
      "pools": [
        {
          "pool": "2001:db8:1::-2001:db8:1::ffff"
        }
      ],
      "evaluate-additional-classes": [ "Client_foo" ],
      ...
    },
    ...
  ],
  ...
}
```

Additional evaluation can be used to express complex dependencies like subnet membership. It can also be used to reverse the precedence; if `option-data` is set in a subnet, it takes precedence over `option-data` in a class. If `option-data` is moved to a required class and required in the subnet, a class evaluated earlier may take precedence.

Additional evaluation is also available at shared network and pool/pd-pool levels. The order in which additional classes are considered is: (pd-)pool, subnet, and shared network, i.e. in the same order from the way in which `option-data` is processed.

Since Kea version 2.7.4 additional classes configured without a test expression are unconditionally added, i.e. they are considered to always be evaluated to `true`.

Note

Because additional evaluation occurs after lease assignment, parameters that would otherwise impact lease life times (e.g. `valid-lifetime`, `preferred-lifetime`) will have no effect when specified in a class that also sets `only-in-additional-list` `true`.

Note

As of Kea version 2.7.4, `only-if-required` and `require-client-classes` have been renamed to `only-in-additional-list` and `evaluate-additional-classes` respectively. The original names will still be accepted as input to allow users to migrate but will eventually be rejected.

10.2.25 DDNS for DHCPv6

As mentioned earlier, [*kea-dhcp6*](#) can be configured to generate requests to the DHCP-DDNS server, [*kea-dhcp-ddns*](#), (referred to herein as "D2") to update DNS entries. These requests are known as NameChangeRequests or NCRs. Each NCR contains the following information:

1. Whether it is a request to add (update) or remove DNS entries.
2. Whether the change requests forward DNS updates (AAAA records), reverse DNS updates (PTR records), or both.
3. The Fully Qualified Domain Name (FQDN), lease address, and DHCID (information identifying the client associated with the FQDN).

DDNS-related parameters are split into two groups:

1. Connectivity Parameters

These are parameters which specify where and how [*kea-dhcp6*](#) connects to and communicates with D2. These parameters can only be specified within the top-level `dhcp-ddns` section in the [*kea-dhcp6*](#) configuration. The connectivity parameters are listed below:

- `enable-updates`
- `server-ip`
- `server-port`
- `sender-ip`
- `sender-port`
- `max-queue-size`
- `ncr-protocol`
- `ncr-format`

2. Behavioral Parameters

These parameters influence behavior such as how client host names and FQDN options are handled. They have been moved out of the `dhcp-ddns` section so that they may be specified at the global, shared-network, subnet and/or pool levels. Furthermore, they are inherited downward from global to shared-network to subnet to pool. In other words, if a parameter is not specified at a given level, the value for that level comes from the level above it. The behavioral parameters are as follows:

- `ddns-send-updates`
- `ddns-override-no-update`
- `ddns-override-client-update`
- `ddns-replace-client-name`
- `ddns-generated-prefix`
- `ddns-qualifying-suffix`

- `ddns-update-on-renew`
- `ddns-conflict-resolution-mode`
- `ddns-ttl`
- `ddns-ttl-percent`
- `ddns-ttl-min`
- `ddns-ttl-max`
- `hostname-char-set`
- `hostname-char-replacement`

Note

Kea 2.7.6 added support for these parameters to the (address) pool level. They are not supported in pd-pools.

Note

Behavioral parameters that affect the FQDN are in effect even if both `enable-updates` and `ddns-send-updates` are `false`, to support environments in which clients are responsible for their own DNS updates. This applies to `ddns-replace-client-name`, `ddns-generated-prefix`, `ddns-qualifying-suffix`, `hostname-char-set`, and `hostname-char-replacement`.

The default configuration and values would appear as follows:

```
"Dhcp6": {
  "dhcp-ddns": {
    // Connectivity parameters
    "enable-updates": false,
    "server-ip": "127.0.0.1",
    "server-port": 53001,
    "sender-ip": "",
    "sender-port": 0,
    "max-queue-size": 1024,
    "ncr-protocol": "UDP",
    "ncr-format": "JSON"
  },

  // Behavioral parameters (global)
  "ddns-send-updates": true,
  "ddns-override-no-update": false,
  "ddns-override-client-update": false,
  "ddns-replace-client-name": "never",
  "ddns-generated-prefix": "myhost",
  "ddns-qualifying-suffix": "",
  "ddns-update-on-renew": false,
  "ddns-conflict-resolution-mode": "check-with-dhcid",
  "hostname-char-set": "",
  "hostname-char-replacement": "",
  ...
}
```


There are two parameters which determine whether *kea-dhcp6* can generate DDNS requests to D2: the existing `dhcp-ddns:enable-updates` parameter, which now only controls whether *kea-dhcp6* connects to D2; and the new behavioral parameter, `ddns-send-updates`, which determines whether DDNS updates are enabled at a given level (i.e. global, shared-network, or subnet). The following table shows how the two parameters function together:

Table 3: Enabling and disabling DDNS updates

dhcp-ddns: enable-updates	Global ddns-send-updates	Outcome
false (default)	false	no updates at any scope
false	true (default)	no updates at any scope
true	false	updates only at scopes with a local value of true for <code>ddns-enable-updates</code>
true	true	updates at all scopes except those with a local value of false for <code>ddns-enable-updates</code>

Kea 1.9.1 added two new parameters; the first is `ddns-update-on-renew`. Normally, when leases are renewed, the server only updates DNS if the DNS information for the lease (e.g. FQDN, DNS update direction flags) has changed. Setting `ddns-update-on-renew` to `true` instructs the server to always update the DNS information when a lease is renewed, even if its DNS information has not changed. This allows Kea to "self-heal" if it was previously unable to add DNS entries or they were somehow lost by the DNS server.

Note

Setting `ddns-update-on-renew` to `true` may impact performance, especially for servers with numerous clients that renew often.

The second parameter added in Kea 1.9.1 is `ddns-use-conflict-resolution`. This boolean parameter was passed through to D2 and enabled or disabled conflict resolution as described in [RFC 4703](#). Beginning with Kea 2.5.0, it is deprecated and replaced by `ddns-conflict-resolution-mode`, which offers four modes of conflict resolution-related behavior:

- **check-with-dhcid** - This mode, the default, instructs D2 to carry out RFC 4703-compliant conflict resolution. Existing DNS entries may only be overwritten if they have a DHCID record and it matches the client's DHCID. This is equivalent to `ddns-use-conflict-resolution: true`;
- **no-check-with-dhcid** - Existing DNS entries may be overwritten by any client, whether those entries include a DHCID record or not. The new entries will include a DHCID record for the client to whom they belong. This is equivalent to `ddns-use-conflict-resolution: false`;
- **check-exists-with-dhcid** - Existing DNS entries may only be overwritten if they have a DHCID record. The DHCID record need not match the client's DHCID. This mode provides a way to protect static DNS entries (those that do not have a DHCID record) while allowing dynamic entries (those that do have a DHCID record) to be overwritten by any client. This behavior was not supported prior to Kea 2.4.0.
- **no-check-without-dhcid** - Existing DNS entries may be overwritten by any client; new entries will not include DHCID records. This behavior was not supported prior to Kea 2.4.0.

Note

For backward compatibility, `ddns-use-conflict-resolution` is still accepted in JSON configuration. The server replaces the value internally with `ddns-conflict-resolution-mode` and an appropriate value: `check-with-dhcid` for `true` and `no-check-with-dhcid` for `false`.

Note

Setting `ddns-conflict-resolution-mode` to any value other than `check-with-dhcid` disables the overwrite safeguards that the rules of conflict resolution (from [RFC 4703](#)) are intended to prevent. This means that existing entries for an FQDN or an IP address made for Client-A can be deleted or replaced by entries for Client-B. Furthermore, there are two scenarios by which entries for multiple clients for the same key (e.g. FQDN or IP) can be created.

1. Client-B uses the same FQDN as Client-A but a different IP address. In this case, the forward DNS entries (AAAA and DHCID RRs) for Client-A will be deleted as they match the FQDN, and new entries for Client-B will be added. The reverse DNS entries (PTR and DHCID RRs) for Client-A, however, will not be deleted as they belong to a different IP address, while new entries for Client-B will still be added.
2. Client-B uses the same IP address as Client-A but a different FQDN. In this case, the reverse DNS entries (PTR and DHCID RRs) for Client-A will be deleted as they match the IP address, and new entries for Client-B will be added. The forward DNS entries (AAAA and DHCID RRs) for Client-A, however, will not be deleted, as they belong to a different FQDN, while new entries for Client-B will still be added.

Disabling conflict resolution should be done only after careful review of specific use cases. The best way to avoid unwanted DNS entries is to always ensure that lease changes are processed through Kea, whether they are released, expire, or are deleted via the `lease6-del` command, prior to reassigning either FQDNs or IP addresses. Doing so causes `kea-dhcp6` to generate DNS removal requests to D2.

The DNS entries Kea creates contain a value for TTL (time to live). By default, the `kea-dhcp4` server calculates that value based on [RFC 4702, Section 5](#), which suggests that the TTL value be 1/3 of the lease's life time, with a minimum value of 10 minutes. There are four optional parameters which may be used to influence the TTL calculation:

- `ddns-ttl` - If specified and greater than 0 use it as the value of TTL unconditionally. It is specified in seconds. By default it is unspecified.
- `ddns-ttl-percent` - If specified and greater than zero use it as the percentage of the lease life time otherwise use 1/3 (per RFC 4702) of the lease life time. It is specified as a decimal percentage (e.g. 0.75, 0.50) and is unspecified by default.
- `ddns-ttl-min` - If specified and greater than 0 use it the minimum otherwise use a minimum of 600 seconds (per RFC 4702). If the calculated TTL is less than the minimum return the minimum. It is specified in seconds. By default it is unspecified.
- `ddns-ttl-max` - If specified and greater than zero limit the calculated TTL to its value. It is specified in seconds. By default it is unspecified.

Note

A value of zero for any of the TTL parameters can be used to suppress (i.e. "unspecify") the value of that parameter inherited from a higher scope.

10.2.25.1 DHCP-DDNS Server Connectivity

For NCRs to reach the D2 server, `kea-dhcp6` must be able to communicate with it. `kea-dhcp6` uses the following configuration parameters to control this communication:

- `enable-updates` - This enables connectivity to `kea-dhcp-ddns` such that DDNS updates can be constructed and sent. It must be `true` for NCRs to be generated and sent to D2. It defaults to `false`.
- `server-ip` - This is the IP address on which D2 listens for requests. The default is the local loopback interface at address 127.0.0.1. Either an IPv4 or IPv6 address may be specified.

- `server-port` - This is the port on which D2 listens for requests. The default value is 53001.
- `sender-ip` - This is the IP address which `kea-dhcp6` uses to send requests to D2. The default value is blank, which instructs `kea-dhcp6` to select a suitable address.
- `sender-port` - This is the port which `kea-dhcp6` uses to send requests to D2. The default value of 0 instructs `kea-dhcp6` to select a suitable port.
- `max-queue-size` - This is the maximum number of requests allowed to queue while waiting to be sent to D2. This value guards against requests accumulating uncontrollably if they are being generated faster than they can be delivered. If the number of requests queued for transmission reaches this value, DDNS updating is turned off until the queue backlog has been sufficiently reduced. The intent is to allow the `kea-dhcp4` server to continue lease operations without running the risk that its memory usage may grow without limit. The default value is 1024.
- `ncr-protocol` - This specifies the socket protocol to use when sending requests to D2. Currently only UDP is supported.
- `ncr-format` - This specifies the packet format to use when sending requests to D2. Currently only JSON format is supported.

By default, `kea-dhcp-ddns` is assumed to be running on the same machine as `kea-dhcp6`, and all of the default values mentioned above should be sufficient. If, however, D2 has been configured to listen on a different address or port, these values must be altered accordingly. For example, if D2 has been configured to listen on 2001:db8::5 port 900, the following configuration is required:

```
"Dhcp6": {
  "dhcp-ddns": {
    "server-ip": "2001:db8::5",
    "server-port": 900,
    ...
  },
  ...
}
```

10.2.25.2 When Does the `kea-dhcp6` Server Generate a DDNS Request?

The `kea-dhcp6` server follows the behavior prescribed for DHCP servers in RFC 4704. It is important to keep in mind that `kea-dhcp6` makes the initial decision of when and what to update and forwards that information to D2 in the form of NCRs. Carrying out the actual DNS updates and dealing with such things as conflict resolution are within the purview of D2 itself (see *The DHCP-DDNS Server*). This section describes when `kea-dhcp6` generates NCRs and the configuration parameters that can be used to influence this decision. It assumes that both the connectivity parameter `enable-updates` and the behavioral parameter `ddns-send-updates` are true.

Note

Currently the interface between `kea-dhcp6` and D2 only supports requests which update DNS entries for a single IP address. If a lease grants more than one address, `kea-dhcp6` creates the DDNS update request for only the first of these addresses.

In general, `kea-dhcp6` generates DDNS update requests when:

1. A new lease is granted in response to a DHCPREQUEST;
2. An existing lease is renewed but the FQDN associated with it has changed; or
3. An existing lease is released in response to a DHCPRELEASE.

In the second case, lease renewal, two DDNS requests are issued: one request to remove entries for the previous FQDN, and a second request to add entries for the new FQDN. In the third case, a lease release - a single DDNS request - to remove its entries will be made.

As for the first case, the decisions involved when granting a new lease are more complex. When a new lease is granted, *kea-dhcp6* generates a DDNS update request only if the DHCPREQUEST contains the FQDN option (code 39). By default, *kea-dhcp6* respects the FQDN N and S flags specified by the client as shown in the following table:

Table 4: Default FQDN flag behavior

Client Flags:N-S	Client Intent	Server Response	Server Flags:N-S-O
0-0	Client wants to do forward updates, server should do reverse updates	Server generates reverse-only request	1-0-0
0-1	Server should do both forward and reverse updates	Server generates request to update both directions	0-1-0
1-0	Client wants no updates done	Server does not generate a request	1-0-0

The first row in the table above represents "client delegation." Here the DHCP client states that it intends to do the forward DNS updates and the server should do the reverse updates. By default, *kea-dhcp6* honors the client's wishes and generates a DDNS request to the D2 server to update only reverse DNS data. The parameter `ddns-override-client-update` can be used to instruct the server to override client delegation requests. When this parameter is `true`, *kea-dhcp6* disregards requests for client delegation and generates a DDNS request to update both forward and reverse DNS data. In this case, the N-S-O flags in the server's response to the client will be 0-1-1, respectively.

(Note that the flag combination N=1, S=1 is prohibited according to [RFC 4702](#). If such a combination is received from the client, the packet will be dropped by *kea-dhcp6*.)

To override client delegation, set the following values in the configuration file:

```
"Dhcp6": {
  "ddns-override-client-update": true,
  ...
}
```

The third row in the table above describes the case in which the client requests that no DNS updates be done. The parameter `ddns-override-no-update` can be used to instruct the server to disregard the client's wishes. When this parameter is `true`, *kea-dhcp6* generates DDNS update requests to *kea-dhcp-ddns* even if the client requests that no updates be done. The N-S-O flags in the server's response to the client will be 0-1-1.

To override client delegation, issue the following commands:

```
"Dhcp6": {
  "ddns-override-no-update": true,
  ...
}
```

The *kea-dhcp6* server always generates DDNS update requests if the client request only contains the Host Name option. In addition, it includes an FQDN option in the response to the client, with the FQDN N-S-O flags set to 0-1-0, respectively. The domain name portion of the FQDN option is the name submitted to D2 in the DDNS update request.

10.2.25.3 kea-dhcp6 Name Generation for DDNS Update Requests

Each NameChangeRequest must of course include the fully qualified domain name whose DNS entries are to be affected. *kea-dhcp6* can be configured to supply a portion or all of that name, based on what it receives from the client in the DHCPREQUEST.

The default rules for constructing the FQDN that will be used for DNS entries are:

1. If the DHCPREQUEST contains the client FQDN option, take the candidate name from there.
2. If the candidate name is a partial (i.e. unqualified) name, then add a configurable suffix to the name and use the result as the FQDN.
3. If the candidate name provided is empty, generate an FQDN using a configurable prefix and suffix.
4. If the client provides neither option, then take no DNS action.

These rules can be amended by setting the `ddns-replace-client-name` parameter, which provides the following modes of behavior:

- **never** - use the name the client sent. If the client sent no name, do not generate one. This is the default mode.
- **always** - replace the name the client sent. If the client sent no name, generate one for the client.
- **when-present** - replace the name the client sent. If the client sent no name, do not generate one.
- **when-not-present** - use the name the client sent. If the client sent no name, generate one for the client.

Note

In early versions of Kea, this parameter was a boolean and permitted only values of `true` and `false`. Boolean values have been deprecated and are no longer accepted; administrators currently using booleans must replace them with the desired mode name. A value of `true` maps to `when-present`, while `false` maps to `never`.

For example, to instruct *kea-dhcp6* to always generate the FQDN for a client, set the parameter `ddns-replace-client-name` to `always` as follows:

```
"Dhcp6": {
  "ddns-replace-client-name": "always",
  ...
}
```

The prefix used in the generation of an FQDN is specified by the `ddns-generated-prefix` parameter. The default value is `"myhost"`. To alter its value, simply set it to the desired string:

```
"Dhcp6": {
  "ddns-generated-prefix": "another.host",
  ...
}
```

The suffix used when generating an FQDN, or when qualifying a partial name, is specified by the `ddns-qualifying-suffix` parameter. It is strongly recommended that the user supply a value for the qualifying suffix when DDNS updates are enabled. For obvious reasons, we cannot supply a meaningful default.

```
"Dhcp6": {
  "ddns-qualifying-suffix": "foo.example.org",
  ...
}
```

When qualifying a partial name, *kea-dhcp6* constructs the name in the format:

```
[candidate-name].[ddns-qualifying-suffix].
```

where *candidate-name* is the partial name supplied in the DHCPREQUEST. For example, if the FQDN domain name value is "some-computer" and the *ddns-qualifying-suffix* is "example.com", the generated FQDN is:

```
some-computer.example.com.
```

When generating the entire name, *kea-dhcp6* constructs the name in the format:

```
[ddns-generated-prefix]-[address-text].[ddns-qualifying-suffix].
```

where *address-text* is simply the lease IP address converted to a hyphenated string. For example, if the lease address is 3001:1::70E, the qualifying suffix is "example.com", and the default value is used for *ddns-generated-prefix*, the generated FQDN is:

```
myhost-3001-1--70E.example.com.
```

As of Kea 2.7.8, generated prefixes that would have ended in a hyphen will have a zero appended to the end. For example, the generated prefix for the address "2001:db8:1::" would be "myhost-2001-db8-1--0".

10.2.25.4 Sanitizing Client FQDN Names

Some DHCP clients may provide values in the name component of the FQDN option (option code 39) that contain undesirable characters. It is possible to configure *kea-dhcp6* to sanitize these values. The most typical use case is ensuring that only characters that are permitted by RFC 1035 be included: A-Z, a-z, 0-9, and "-". This may be accomplished with the following two parameters:

- *hostname-char-set* - a regular expression describing the invalid character set. This can be any valid, regular expression using POSIX extended expression syntax. Embedded nulls (0x00) are always considered an invalid character to be replaced (or omitted). The default is "[^A-Za-z0-9.-]". This matches any character that is not a letter, digit, dot, hyphen, or null.
- *hostname-char-replacement* - a string of zero or more characters with which to replace each invalid character in the host name. An empty string causes invalid characters to be OMITTED rather than replaced. The default is "".

The following configuration replaces anything other than a letter, digit, dot, or hyphen with the letter "x":

```
"Dhcp6": {
  "hostname-char-set": "[^A-Za-z0-9.-]",
  "hostname-char-replacement": "x",
  ...
}
```

Thus, a client-supplied value of "myhost-\$[123.org" would become "myhost-xx123.org". Sanitizing is performed only on the portion of the name supplied by the client, and it is performed before applying a qualifying suffix (if one is defined and needed).

Note

Name sanitizing is meant to catch the more common cases of invalid characters through a relatively simple character-replacement scheme. It is difficult to devise a scheme that works well in all cases. Administrators who find they have clients with odd corner cases of character combinations that cannot be readily handled with this mechanism should consider writing a hook that can carry out sufficiently complex logic to address their needs.

Make sure that the dot, ".", is considered a valid character by the *hostname-char-set* expression, such as this: "[^A-Za-z0-9.-]". When scrubbing FQDNs, dots are treated as delimiters and used to separate the option value into individual domain labels that are scrubbed and then re-assembled.

If clients are sending values that differ only by characters considered as invalid by the `hostname-char-set`, be aware that scrubbing them will yield identical values. In such cases, DDNS conflict rules will permit only one of them to register the name.

Finally, given the latitude clients have in the values they send, it is virtually impossible to guarantee that a combination of these two parameters will always yield a name that is valid for use in DNS. For example, using an empty value for `hostname-char-replacement` could yield an empty domain label within a name, if that label consists only of invalid characters.

Note

It is possible to specify `hostname-char-set` and/or `hostname-char-replacement` at the global scope.

The Kea hook library *libdhcp_ddns_tuning.so* provides the ability for both *kea-dhcp4* and *kea-dhcp6* to generate host names procedurally based on an expression, to skip DDNS updates on a per-client basis, or to fine-tune various DNS update aspects. Please refer to the *libdhcp_ddns_tuning.so: DDNS Tuning* documentation for the configuration options.

Note

Beginning with Kea 3.0.2, when sanitizing results in an empty string, the FQDN option will be ignored as if the client had not sent it. This can happen if the value sent by the client contains only characters that are deemed invalid by the expression in 'hostname-char-set' and 'hostname-char-replacement' is empty, "".

10.2.26 DHCPv4-over-DHCPv6: DHCPv6 Side

The support of DHCPv4-over-DHCPv6 transport is described in [RFC 7341](#) and is implemented using cooperating DHCPv4 and DHCPv6 servers. This section is about the configuration of the DHCPv6 side (the DHCPv4 side is described in *DHCPv4-over-DHCPv6: DHCPv4 Side*).

Note

The DHCPv4-over-DHCPv6 details of the inter-process communication may change; for instance, the support of port relay (RFC 8357) introduced an incompatible change. Both the DHCPv4 and DHCPv6 sides should be running the same version of Kea.

There is only one specific parameter for the DHCPv6 side: `dhcp4o6-port`, which specifies the first of the two consecutive ports of the UDP sockets used for the communication between the DHCPv6 and DHCPv4 servers. The DHCPv6 server is bound to `::1` on `port` and connected to `::1` on `port + 1`.

Two other configuration entries are generally required: unicast traffic support (see *Unicast Traffic Support*) and the DHCP 4o6 server address option (name "dhcp4o6-server-addr", code 88).

ISC tested the following configuration:

```
{
# DHCPv6 conf
"Dhcp6": {
```

(continues on next page)

(continued from previous page)

```

"interfaces-config": {
  "interfaces": [ "eno33554984/2001:db8:1:1::1" ]
},

"lease-database": {
  "type": "memfile",
  "name": "leases6"
},

"preferred-lifetime": 3000,
"valid-lifetime": 4000,
"renew-timer": 1000,
"rebind-timer": 2000,

"subnet6": [ {
  "id": 1,
  "subnet": "2001:db8:1:1::/64",
  "interface": "eno33554984",
  "pools": [ { "pool": "2001:db8:1:1::1:0/112" } ]
} ],

"dhcp4o6-port": 786,

"option-data": [ {
  "name": "dhcp4o6-server-addr",
  "code": 88,
  "space": "dhcp6",
  "csv-format": true,
  "data": "2001:db8:1:1::1"
} ],

"loggers": [ {
  "name": "kea-dhcp6",
  "output-options": [ {
    "output": "kea-dhcp6.log"
  } ],
  "severity": "DEBUG",
  "debuglevel": 0
} ]
}
}

```

Note

Relayed DHCPv4-QUERY DHCPv6 messages are not supported.

10.2.27 Sanity Checks in DHCPv6

An important aspect of a well-running DHCP system is an assurance that the data remains consistent; however, in some cases it may be convenient to tolerate certain inconsistent data. For example, a network administrator who temporarily removes a subnet from a configuration would not want all the leases associated with it to disappear from the lease database. Kea has a mechanism to implement sanity checks for situations like this.

Kea supports a configuration scope called `sanity-checks`. A parameter, called `lease-checks`, governs the verification carried out when a new lease is loaded from a lease file. This mechanism permits Kea to attempt to correct inconsistent data.

Every subnet has a `subnet-id` value; this is how Kea internally identifies subnets. Each lease has a `subnet-id` parameter as well, which identifies the subnet it belongs to. However, if the configuration has changed, it is possible that a lease could exist with a `subnet-id` but without any subnet that matches it. Also, it is possible that the subnet's configuration has changed and the `subnet-id` now belongs to a subnet that does not match the lease.

Kea's corrective algorithm first checks to see if there is a subnet with the `subnet-id` specified by the lease. If there is, it verifies whether the lease belongs to that subnet. If not, depending on the `lease-checks` setting, the lease is discarded, a warning is displayed, or a new subnet is selected for the lease that matches it topologically.

Since delegated prefixes do not have to belong to a subnet in which they are offered, there is no way to implement such a mechanism for IPv6 prefixes. As such, the mechanism works for IPv6 addresses only.

There are five levels which are supported:

- `none` - do no special checks; accept the lease as is.
- `warn` - if problems are detected display a warning, but accept the lease data anyway. This is the default value.
- `fix` - if a data inconsistency is discovered, try to correct it. If the correction is not successful, insert the incorrect data anyway.
- `fix-del` - if a data inconsistency is discovered, try to correct it. If the correction is not successful, reject the lease. This setting ensures the data's correctness, but some incorrect data may be lost. Use with care.
- `del` - if any inconsistency is detected, reject the lease. This is the strictest mode; use with care.

This feature is currently implemented for the memfile backend. The sanity check applies to the lease database in memory, not to the lease file, i.e. inconsistent leases will stay in the lease file.

An example configuration that sets this parameter looks as follows:

```
"Dhcp6": {
  "sanity-checks": {
    "lease-checks": "fix-del"
  },
  ...
}
```

10.2.28 Storing Extended Lease Information

To support such features as DHCPv6 Reconfigure (RFC 3315) and Leasequery (RFC 5007), additional information must be stored with each lease. Because the amount of information stored for each lease has ramifications in terms of performance and system resource consumption, storage of this additional information is configurable through the `store-extended-info` parameter. It defaults to `false` and may be set at the global, shared-network, and subnet levels.

```
"Dhcp6": {
  "store-extended-info": true,
```

(continues on next page)

(continued from previous page)

```
...
}
```

When set to `true`, information relevant to the DHCPv6 query (e.g. REQUEST, RENEW, or REBIND) asking for the lease is added into the lease's `user-context` as a map element labeled "ISC". Currently, the information contained in the map is a list of relays, one for each relay message layer that encloses the client query. The lease's `user-context` for a two-hop query might look something like this (shown pretty-printed for clarity):

```
{
  "ISC": {
    "relay-info": [
      {
        "hop": 3,
        "link": "2001:db8::1",
        "peer": "2001:db8::2"
      },
      {
        "hop": 2,
        "link": "2001:db8::3",
        "options": "0x00C800080102030405060708",
        "peer": "2001:db8::4"
      },
      {
        "hop": 1,
        "link": "2001:db8::5",
        "options": "0x002500060102030405060035000864646464646464",
        "remote-id": "010203040506",
        "relay-id": "6464646464646464"
      }
    ]
  }
}
```

Note

Prior to Kea version 2.3.2, this entry was named `relays`; remote and relay identifier options were not decoded.

Note

It is possible that other hook libraries are already using `user-context`. Enabling `store-extended-info` should not interfere with any other `user-context` content, as long as it does not also use an element labeled "ISC". In other words, `user-context` is intended to be a flexible container serving multiple purposes. As long as no other purpose also writes an "ISC" element to `user-context` there should not be a conflict.

Extended lease information is also subject to configurable sanity checking. The parameter in the `sanity-checks` scope is named `extended-info-checks` and supports these levels:

- `none` - do no check nor upgrade. This level should be used only when extended info is not used at all or when no badly formatted extended info, including using the old format, is expected.
- `fix` - fix some common inconsistencies and upgrade extended info using the old format to the new one. It is the

default level and is convenient when the Leasequery hook library is not loaded.

- **strict** - fix all inconsistencies which have an impact on the (Bulk) Leasequery hook library.
- **pedantic** - enforce full conformance to the format produced by the Kea code; for instance, no extra entries are allowed with the exception of **comment**.

Note

This feature is currently implemented only for the memfile backend. The sanity check applies to the lease database in memory, not to the lease file, i.e. inconsistent leases stay in the lease file.

10.2.29 Multi-Threading Settings

The Kea server can be configured to process packets in parallel using multiple threads. These settings can be found under the **multi-threading** structure and are represented by:

- **enable-multi-threading** - use multiple threads to process packets in parallel. The default is **true**.
- **thread-pool-size** - specify the number of threads to process packets in parallel. It may be set to **0** (auto-detect), or any positive number that explicitly sets the thread count. The default is **0**.
- **packet-queue-size** - specify the size of the queue used by the thread pool to process packets. It may be set to **0** (unlimited), or any positive number that explicitly sets the queue size. The default is **64**.

An example configuration that sets these parameters looks as follows:

```
"Dhcp6": {
  "multi-threading": {
    "enable-multi-threading": true,
    "thread-pool-size": 4,
    "packet-queue-size": 16
  },
  ...
}
```

10.2.30 Multi-Threading Settings With Different Database Backends

The Kea DHCPv6 server is benchmarked by ISC to determine which settings give the best performance. Although this section describes our results, they are merely recommendations and are very dependent on the particular hardware used for benchmarking. We strongly advise that administrators run their own performance benchmarks.

A full report of performance results for the latest stable Kea version can be found [here](#). This includes hardware and benchmark scenario descriptions, as well as current results.

After enabling multi-threading, the number of threads is set by the **thread-pool-size** parameter. Results from our experiments show that the best settings for *kea-dhcp6* are:

- **thread-pool-size**: 4 when using **memfile** for storing leases.
- **thread-pool-size**: 12 or more when using **mysql** for storing leases.
- **thread-pool-size**: 6 when using **postgresql**.

Another very important parameter is **packet-queue-size**; in our benchmarks we used it as a multiplier of **thread-pool-size**. The actual setting strongly depends on **thread-pool-size**.

We saw the best results in our benchmarks with the following settings:

- `packet-queue-size`: $150 * \text{thread-pool-size}$ when using `memfile` for storing leases; in our case it was $150 * 4 = 600$. This means that at any given time, up to 600 packets could be queued.
- `packet-queue-size`: $200 * \text{thread-pool-size}$ when using `mysql` for storing leases; in our case it was $200 * 12 = 2400$. This means that up to 2400 packets could be queued.
- `packet-queue-size`: $11 * \text{thread-pool-size}$ when using `postgresql` for storing leases; in our case it was $11 * 6 = 66$.

10.2.31 Lease Caching

Clients that attempt multiple renewals in a short period can cause the server to update and write to the database frequently, resulting in a performance impact on the server. The cache parameters instruct the DHCP server to avoid updating leases too frequently, thus avoiding this behavior. Instead, the server assigns the same lease (i.e. reuses it) with no modifications except for CLTT (Client Last Transmission Time), which does not require disk operations.

The two parameters are the `cache-threshold` double and the `cache-max-age` integer. These parameters can be configured at the global, shared-network, and subnet levels. The subnet level has precedence over the shared-network level, while the global level is used as a last resort. For example:

```
{
  "subnet6": [
    {
      "subnet": "2001:db8:1:1::/64",
      "pools": [ { "pool": "2001:db8:1:1::1:0/112" } ],
      "cache-threshold": .25,
      "cache-max-age": 600,
      "valid-lifetime": 2000,
      ...
    }
  ],
  ...
}
```

When an already-assigned lease can fulfill a client query:

- any important change, e.g. for DDNS parameter, hostname, or preferred or valid lifetime reduction, makes the lease not reusable.
- lease age, i.e. the difference between the creation or last modification time and the current time, is computed (elapsed duration).
- if `cache-max-age` is explicitly configured, it is compared with the lease age; leases that are too old are not reusable. This means that the value 0 for `cache-max-age` disables the lease cache feature.
- if `cache-threshold` is explicitly configured and is between 0.0 and 1.0, it expresses the percentage of the lease valid lifetime which is allowed for the lease age. Values below and including 0.0 and values greater than 1.0 disable the lease cache feature.

In our example, a lease with a valid lifetime of 2000 seconds can be reused if it was committed less than 500 seconds ago. With a lifetime of 3000 seconds, a maximum age of 600 seconds applies.

In outbound client responses (e.g. DHCPV6_REPLY messages), the used preferred and valid lifetimes are the reusable values, i.e. the expiration dates do not change.

10.3 Host Reservations in DHCPv6

There are many cases where it is useful to provide a configuration on a per-host basis. The most obvious one is to reserve a specific, static IPv6 address or/and prefix for exclusive use by a given client (host); the returning client receives the same address and/or prefix every time, and other clients will never get that address. Host reservations are also convenient when a host has specific requirements, e.g. a printer that needs additional DHCP options or a cable modem that needs specific parameters. Yet another possible use case is to define unique names for hosts.

There may be cases when a new reservation has been made for a client for an address or prefix currently in use by another client. We call this situation a "conflict." These conflicts get resolved automatically over time, as described in subsequent sections. Once a conflict is resolved, the correct client will receive the reserved configuration when it renews.

Host reservations are defined as parameters for each subnet. Each host must be identified by either DUID or its hardware/MAC address; see *MAC/Hardware Addresses in DHCPv6* for details. There is an optional `reservations` array in the `subnet6` structure; each element in that array is a structure that holds information about reservations for a single host. In particular, the structure has an identifier that uniquely identifies a host. In the DHCPv6 context, the identifier is usually a DUID, but it can also be a hardware or MAC address. One or more addresses or prefixes may also be specified, and it is possible to specify a hostname and DHCPv6 options for a given host.

Note

The reserved address must be within the subnet. This does not apply to reserved prefixes.

The following example shows how to reserve addresses and prefixes for specific hosts:

```
{
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pools": [ { "pool": "2001:db8:1::/80" } ],
      "pd-pools": [
        {
          "prefix": "2001:db8:1:8000::",
          "prefix-len": 56,
          "delegated-len": 64
        }
      ],
      "reservations": [
        {
          "duid": "01:02:03:04:05:0A:0B:0C:0D:0E",
          "ip-addresses": [ "2001:db8:1::100" ]
        },
        {
          "hw-address": "00:01:02:03:04:05",
          "ip-addresses": [ "2001:db8:1::101", "2001:db8:1::102" ]
        },
        {
          "duid": "01:02:03:04:05:06:07:08:09:0A",
          "ip-addresses": [ "2001:db8:1::103" ],
          "prefixes": [ "2001:db8:2:abcd::/64" ],
          "hostname": "foo.example.com"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

        }
    ]
}
...
}

```

This example includes reservations for three different clients. The first reservation is for the address 2001:db8:1::100, for a client using DUID 01:02:03:04:05:0A:0B:0C:0D:0E. The second reservation is for two addresses, 2001:db8:1::101 and 2001:db8:1::102, for a client using MAC address 00:01:02:03:04:05. Lastly, address 2001:db8:1::103 and prefix 2001:db8:2:abcd::/64 are reserved for a client using DUID 01:02:03:04:05:06:07:08:09:0A. The last reservation also assigns a hostname to this client.

DHCPv6 allows a single client to lease multiple addresses and multiple prefixes at the same time. Therefore `ip-addresses` and `prefixes` are plural and are actually arrays. When the client sends multiple IA options (IA_NA or IA_PD), each reserved address or prefix is assigned to an individual IA of the appropriate type. If the number of IAs of a specific type is lower than the number of reservations of that type, the number of reserved addresses or prefixes assigned to the client is equal to the number of IA_NAs or IA_PDs sent by the client; that is, some reserved addresses or prefixes are not assigned. However, they still remain reserved for this client and the server will not assign them to any other client. If the number of IAs of a specific type sent by the client is greater than the number of reserved addresses or prefixes, the server will try to assign all reserved addresses or prefixes to the individual IAs and dynamically allocate addresses or prefixes to the remaining IAs. If the server cannot assign a reserved address or prefix because it is in use, the server will select the next reserved address or prefix and try to assign it to the client. If the server subsequently finds that there are no more reservations that can be assigned to the client at that moment, the server will try to assign leases dynamically.

Making a reservation for a mobile host that may visit multiple subnets requires a separate host definition in each subnet that host is expected to visit. It is not possible to define multiple host definitions with the same hardware address in a single subnet. Multiple host definitions with the same hardware address are valid if each is in a different subnet. The reservation for a given host should include only one identifier, either DUID or hardware address; defining both for the same host is considered a configuration error.

Adding host reservations incurs a performance penalty. In principle, when a server that does not support host reservation responds to a query, it needs to check whether there is a lease for a given address being considered for allocation or renewal. The server that does support host reservation has to perform additional checks: not only whether the address is currently used (i.e., if there is a lease for it), but also whether the address could be used by someone else (i.e., if there is a reservation for it). That additional check incurs extra overhead.

10.3.1 Address/Prefix Reservation Types

In a typical Kea scenario there is an IPv6 subnet defined, with a certain part of it dedicated for dynamic address allocation by the DHCPv6 server. There may be an additional address space defined for prefix delegation. Those dynamic parts are referred to as dynamic pools, address and prefix pools, or simply pools. In principle, a host reservation can reserve any address or prefix that belongs to the subnet. The reservations that specify addresses that belong to configured pools are called "in-pool reservations." In contrast, those that do not belong to dynamic pools are called "out-of-pool reservations." There is no formal difference in the reservation syntax and both reservation types are handled uniformly.

Kea supports global host reservations. These are reservations that are specified at the global level within the configuration and that do not belong to any specific subnet. Kea still matches inbound client packets to a subnet as before, but when the subnet's reservation mode is set to "global", Kea looks for host reservations only among the global reservations defined. Typically, such reservations would be used to reserve hostnames for clients which may move from one subnet to another.

Note

Global reservations, while useful in certain circumstances, have aspects that must be given due consideration when using them. Please see *Conflicts in DHCPv6 Reservations* for more details.

Note

Since Kea 1.9.1, reservation mode has been replaced by three boolean flags, `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool`, which allow the configuration of host reservations both globally and in a subnet. In such cases a subnet host reservation has preference over a global reservation when both exist for the same client.

Note

Beginning with Kea 2.7.3, the host reservation syntax supports a new entry, `excluded-prefixes`. It can be used to specify prefixes the client should exclude from delegated prefixes. When present it must have the same number of elements as the `prefixes` entry. Both entries contain strings representing IPv6 prefixes. Each element of the `excluded-prefixes` must be either an empty string or match the prefix at the same position in `prefixes`. An empty `excluded-prefixes` list or a list with only empty strings can be omitted. An example which excludes `2001:db8:0:1::/64` from `2001:db8::/48` is shown below:

```
{
  "reservations": [
    {
      "duid": "01:02:03:04:05:06:07:08:09:0A",
      "ip-addresses": [ "2001:db8:1::103" ],
      "prefixes": [ "2001:db8::/48" ],
      "excluded-prefixes": [ "2001:db8:0:1::/64" ],
      "hostname": "foo.example.com"
    }
  ],
  ...
}
```

Note

Since host reservations have precedence over prefix pools, a reserved prefix without an excluded prefix will not add a `pd-exclude` option to the prefix option even if the delegated prefix is in a configured prefix pool that does specify an excluded prefix (different from previous behavior).

10.3.2 Conflicts in DHCPv6 Reservations

As reservations and lease information are stored separately, conflicts may arise. Consider the following series of events: the server has configured the dynamic pool of addresses from the range of `2001:db8::10` to `2001:db8::20`. Host A requests an address and gets `2001:db8::10`. Now the system administrator decides to reserve address `2001:db8::10` for Host B. In general, reserving an address that is currently assigned to someone else is not recommended, but there are valid use cases where such an operation is warranted.

The server now has a conflict to resolve. If Host B boots up and requests an address, the server cannot immediately

assign the reserved address 2001:db8::10. A naive approach would be to immediately remove the lease for Host A and create a new one for Host B. That would not solve the problem, though, because as soon as Host B gets the address, it will detect that the address is already in use (by Host A) and will send a DHCPDECLINE message. Therefore, in this situation, the server has to temporarily assign a different address from the dynamic pool (not matching what has been reserved) to Host B.

When Host A renews its address, the server will discover that the address being renewed is now reserved for someone else - Host B. The server will remove the lease for 2001:db8::10, select a new address, and create a new lease for it. It will send two addresses in its response: the old address, with the lifetime set to 0 to explicitly indicate that it is no longer valid; and the new address, with a non-zero lifetime. When Host B tries to renew its temporarily assigned address, the server will detect that the existing lease does not match the reservation, so it will release the current address Host B has and will create a new lease matching the reservation. As before, the server will send two addresses: the temporarily assigned one with a zero lifetime, and the new one that matches the reservation with the proper lifetime set.

This recovery will succeed, even if other hosts attempt to get the reserved address. If Host C requests the address 2001:db8::10 after the reservation is made, the server will propose a different address.

This recovery mechanism allows the server to fully recover from a case where reservations conflict with existing leases; however, this procedure takes roughly as long as the value set for `renew-timer`. The best way to avoid such a recovery is not to define new reservations that conflict with existing leases. Another recommendation is to use out-of-pool reservations; if the reserved address does not belong to a pool, there is no way that other clients can get it.

Note

The conflict-resolution mechanism does not work for global reservations. Although the global address reservations feature may be useful in certain settings, it is generally recommended not to use global reservations for addresses. Administrators who do choose to use global reservations must manually ensure that the reserved addresses are not in dynamic pools.

10.3.3 Reserving a Hostname

When the reservation for a client includes the `hostname`, the server assigns this hostname to the client and sends it back in the Client FQDN option, if the client included the Client FQDN option in its message to the server. The reserved hostname always takes precedence over the hostname supplied by the client (via the FQDN option) or the autogenerated (from the IPv6 address) hostname.

The server qualifies the reserved hostname with the value of the `ddns-qualifying-suffix` parameter. For example, the following subnet configuration:

```
{
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pools": [ { "pool": "2001:db8:1::/80" } ],
      "ddns-qualifying-suffix": "example.isc.org.",
      "reservations": [
        {
          "duid": "01:02:03:04:05:0A:0B:0C:0D:0E",
          "ip-addresses": [ "2001:db8:1::100" ],
          "hostname": "alice-laptop"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

],
"dhcp-ddns": {
    "enable-updates": true
},
...
}

```

will result the "alice-laptop.example.isc.org." hostname being assigned to the client using the DUID "01:02:03:04:05:0A:0B:0C:0D:0E". If the `ddns-qualifying-suffix` is not specified, the default (empty) value will be used, and in this case the value specified as a `hostname` will be treated as a fully qualified name. Thus, by leaving the `ddns-qualifying-suffix` empty it is possible to qualify hostnames for different clients with different domain names:

```

{
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pools": [ { "pool": "2001:db8:1::/80" } ],
      "reservations": [
        {
          "duid": "01:02:03:04:05:0A:0B:0C:0D:0E",
          "ip-addresses": [ "2001:db8:1::100" ],
          "hostname": "mark-desktop.example.org."
        }
      ]
    }
  ],
  "dhcp-ddns": {
    "enable-updates": true
  }
}

```

The above example results in the assignment of the "mark-desktop.example.org." hostname to the client using the DUID "01:02:03:04:05:0A:0B:0C:0D:0E".

10.3.4 Including Specific DHCPv6 Options in Reservations

Kea offers the ability to specify options on a per-host basis. These options follow the same rules as any other options. These can be standard options (see *Standard DHCPv6 Options*), custom options (see *Custom DHCPv6 Options*), or vendor-specific options (see *DHCPv6 Vendor-Specific Options*). The following example demonstrates how standard options can be defined.

```

{
  "reservations": [
    {
      "duid": "01:02:03:05:06:07:08",
      "ip-addresses": [ "2001:db8:1::2" ],
      "option-data": [
        {
          "name": "dns-servers",
          "data": "3000:1::234"
        }
      ],
    }
  ],
}

```

(continues on next page)

(continued from previous page)

```

{
    "name": "nis-servers",
    "data": "3000:1::234"
},
...
],
...
},
...
],
...
}

```

Vendor-specific options can be reserved in a similar manner:

```

{
  "reservations": [
    {
      "duid": "aa:bb:cc:dd:ee:ff",
      "ip-addresses": [ "2001:db8::1" ],
      "option-data": [
        {
          "name": "vendor-opts",
          "data": 4491
        },
        {
          "name": "tftp-servers",
          "space": "vendor-4491",
          "data": "3000:1::234"
        },
        ...
      ],
      ...
    },
    ...
  ],
  ...
}

```

Options defined at the host level have the highest priority. In other words, if there are options defined with the same type on global, subnet, class, and host levels, the host-specific values are used.

10.3.5 Reserving Client Classes in DHCPv6

Using Expressions in Classification explains how to configure the server to assign classes to a client, based on the content of the options that this client sends to the server. Host reservation mechanisms also allow for the static assignment of classes to clients. The definitions of these classes are placed in the Kea configuration file or a database. The following configuration snippet shows how to specify that a client belongs to the classes `reserved-class1` and `reserved-class2`. Those classes are associated with specific options sent to the clients which belong to them.

```

{
  "client-classes": [
    {

```

(continues on next page)

(continued from previous page)

```

    "name": "reserved-class1",
    "option-data": [
      {
        "name": "dns-servers",
        "data": "2001:db8:1::50"
      }
    ]
  },
  {
    "name": "reserved-class2",
    "option-data": [
      {
        "name": "nis-servers",
        "data": "2001:db8:1::100"
      }
    ]
  }
],
"subnet6": [
  {
    "id": 1,
    "pools": [ { "pool": "2001:db8:1::/64" } ],
    "subnet": "2001:db8:1::/48",
    "reservations": [
      {
        "duid": "01:02:03:04:05:06:07:08",

        "client-classes": [ "reserved-class1", "reserved-class2" ]

      }
    ]
  }
]
}

```

In some cases the host reservations can be used in conjunction with client classes specified within the Kea configuration. In particular, when a host reservation exists for a client within a given subnet, the "KNOWN" built-in class is assigned to the client. Conversely, when there is no static assignment for the client, the "UNKNOWN" class is assigned to the client. Class expressions within the Kea configuration file can refer to "KNOWN" or "UNKNOWN" classes using the "member" operator. For example:

```

{
  "client-classes": [
    {
      "name": "dependent-class",
      "test": "member('KNOWN')",
      "only-in-additional-list": true
    }
  ]
}

```

The `only-in-additional-list` parameter is needed here to force evaluation of the class after the lease has been allocated and thus the reserved class has been also assigned.

Note

The classes specified in non-global host reservations are assigned to the processed packet after all classes with the `only-in-additional-list` parameter set to `false` have been evaluated. This means that these classes must not depend on the statically assigned classes from the host reservations. If such a dependency is needed, the `only-in-additional-list` parameter must be set to `true` for the dependent classes. Such classes are evaluated after the static classes have been assigned to the packet. This, however, imposes additional configuration overhead, because all classes marked as `only-in-additional-list` must be listed in the `evaluate-additional-classes` list for every subnet where they are used.

Note

Client classes specified within the Kea configuration file may depend on the classes specified within the global host reservations. In such a case the `only-in-additional-list` parameter is not needed. Refer to the [Pool Selection with Client Class Reservations](#) and [Subnet Selection with Client Class Reservations](#) for specific use cases.

10.3.6 Storing Host Reservations in MySQL or PostgreSQL

Kea can store host reservations in MySQL or PostgreSQL. See [Hosts Storage](#) for information on how to configure Kea to use reservations stored in MySQL or PostgreSQL. Kea provides a dedicated hook for managing reservations in a database; section [libdhcp_host_cmds.so: Host Commands](#) provides detailed information. The [Kea wiki](#) provides some examples of how to conduct common host reservation operations.

Note

In Kea, the maximum length of an option specified per-host is arbitrarily set to 4096 bytes.

10.3.7 Fine-Tuning DHCPv6 Host Reservation

The host reservation capability introduces additional restrictions for the allocation engine (the component of Kea that selects an address for a client) during lease selection and renewal. In particular, three major checks are necessary. First, when selecting a new lease, it is not sufficient for a candidate lease to simply not be in use by another DHCP client; it also must not be reserved for another client. Similarly, when renewing a lease, an additional check must be performed to see whether the address being renewed is reserved for another client. Finally, when a host renews an address or a prefix, the server must check whether there is a reservation for this host, which would mean the existing (dynamically allocated) address should be revoked and the reserved one be used instead.

Some of those checks may be unnecessary in certain deployments, and not performing them may improve performance. The Kea server provides the `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool` configuration parameters to select the types of reservations allowed for a particular subnet. Each reservation type has different constraints for the checks to be performed by the server when allocating or renewing a lease for the client.

Configuration flags are:

- `reservations-in-subnet` - when set to `true`, it enables in-pool host reservation types. This setting is the default value, and is the safest and most flexible. However, as all checks are conducted, it is also the slowest. It does not check against global reservations. This flag defaults to `true`.
- `reservations-out-of-pool` - when set to `true`, it allows only out-of-pool host reservations. In this case the server assumes that all host reservations are for addresses that do not belong to the dynamic pool. Therefore, it can skip the reservation checks when dealing with in-pool addresses, thus improving performance. Do not use this mode if any reservations use in-pool addresses. Caution is advised when using this setting; Kea does not

sanity-check the reservations against `reservations-out-of-pool` and misconfiguration may cause problems. This flag defaults to `false`.

- **reservations-global** - allows global host reservations. With this setting in place, the server searches for reservations for a client among the defined global reservations. If an address is specified, the server skips the reservation checks carried out in other modes, thus improving performance. Caution is advised when using this setting; Kea does not sanity-check the reservations when **reservations-global** is set to **true**, and misconfiguration may cause problems. This flag defaults to **false**.

Note: setting all flags to false disables host reservation support.


As there are no reservations, the server skips all checks. Any reservations defined are completely ignored. As checks are skipped, the server may operate faster in this mode.

Since Kea 1.9.1 the `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool` flags are supported.

The `reservations-global`, `reservations-in-subnet` and `reservations-out-of-pool` parameters can be specified at:


- global level: `.Dhcp6["reservations-global"]` (lowest priority: gets overridden by all others)
- subnet level: `.Dhcp6.subnet6[["reservations-in-subnet"]]` (low priority)
- shared-network level: `.Dhcp6["shared-networks"][["reservations-out-of-pool"]]` (high priority)
- shared-network subnet-level: `.Dhcp6["shared-networks"][].subnet6[["reservations-out-of-pool"]]` (highest priority: overrides all others)

To decide which flags to use, the following decision diagram may be useful:



(continues on next page)

(continued from previous page)



An example configuration that disables reservations looks as follows:

```
{
  "Dhcp6": {
    "subnet6": [
      {
        "id": 1,
        "pools": [
          {
            "pool": "2001:db8:1::-2001:db8:1::100"
          }
        ],
        "reservations-global": false,
        "reservations-in-subnet": false,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}
```

An example configuration using global reservations is shown below:

```
{
  "Dhcp6": {
    "reservations-global": true,
    "reservations": [
      {
        "duid": "00:03:00:01:11:22:33:44:55:66",
        "hostname": "host-one"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    {
      "duid": "00:03:00:01:99:88:77:66:55:44",
      "hostname": "host-two"
    }
  ],
  "subnet6": [
    {
      "id": 1,
      "pools": [
        {
          "pool": "2001:db8:1::-2001:db8:1::100"
        }
      ],
      "subnet": "2001:db8:1::/64"
    }
  ]
}

```

The meaning of the reservation flags are:

- `reservations-global`: fetch global reservations.
- `reservations-in-subnet`: fetch subnet reservations. For a shared network this includes all subnet members of the shared network.
- `reservations-out-of-pool`: this makes sense only when the `reservations-in-subnet` flag is true. When `reservations-out-of-pool` is true, the server assumes that all host reservations are for addresses that do not belong to the dynamic pool. Therefore, it can skip the reservation checks when dealing with in-pool addresses, thus improving performance. The server will not assign reserved addresses that are inside the dynamic pools to the respective clients. This also means that the addresses matching the respective reservations from inside the dynamic pools (if any) can be dynamically assigned to any client.

The disabled configuration corresponds to:

```

{
  "Dhcp6": {
    "reservations-global": false,
    "reservations-in-subnet": false
  }
}

```

The global configuration using `reservations-global` corresponds to:

```

{
  "Dhcp6": {
    "reservations-global": true,
    "reservations-in-subnet": false
  }
}

```

The out-of-pool configuration using `reservations-out-of-pool` corresponds to:

```
{
  "Dhcp6": {
    "reservations-global": false,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": true
  }
}
```

And the in-subnet configuration using `reservations-in-subnet` corresponds to:

```
{
  "Dhcp6": {
    "reservations-global": false,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": false
  }
}
```

To activate both global and in-subnet, the following combination can be used:

```
{
  "Dhcp6": {
    "reservations-global": true,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": false
  }
}
```

To activate both global and out-of-pool, the following combination can be used:

```
{
  "Dhcp6": {
    "reservations-global": true,
    "reservations-in-subnet": true,
    "reservations-out-of-pool": true
  }
}
```

Enabling out-of-pool and disabling in-subnet at the same time is not recommended because out-of-pool applies to host reservations in a subnet, which are fetched only when the in-subnet flag is true.

The parameter can be specified at the global, subnet, and shared-network levels.

An example configuration that disables reservations looks as follows:

```
{
  "Dhcp6": {
    "subnet6": [
      {
        "reservations-global": false,
        "reservations-in-subnet": false,
        "subnet": "2001:db8:1::/64",
        "id": 1
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    ]
  }
}

```

An example configuration using global reservations is shown below:

```

{
  "Dhcp6": {
    "reservations": [
      {
        "duid": "00:03:00:01:11:22:33:44:55:66",
        "hostname": "host-one"
      },
      {
        "duid": "00:03:00:01:99:88:77:66:55:44",
        "hostname": "host-two"
      }
    ],
    "reservations-global": true,
    "reservations-in-subnet": false,
    "subnet6": [
      {
        "pools": [
          {
            "pool": "2001:db8:1::-2001:db8:1::100"
          }
        ],
        "subnet": "2001:db8:1::/64",
        "id": 1
      }
    ]
  }
}

```

For more details regarding global reservations, see *Global Reservations in DHCPv6*.

Another aspect of host reservations is the different types of identifiers. Kea currently supports two types of identifiers in DHCPv6: hardware address and DUID. This is beneficial from a usability perspective; however, there is one drawback. For each incoming packet Kea has to extract each identifier type and then query the database to see if there is a reservation by this particular identifier. If nothing is found, the next identifier is extracted and the next query is issued. This process continues until either a reservation is found or all identifier types have been checked. Over time, with an increasing number of supported identifier types, Kea would become slower and slower.

To address this problem, a parameter called `host-reservation-identifiers` is available. It takes a list of identifier types as a parameter. Kea checks only those identifier types enumerated in `host-reservation-identifiers`. From a performance perspective, the number of identifier types should be kept to a minimum, ideally one. If the deployment uses several reservation types, please enumerate them from most- to least-frequently used, as this increases the chances of Kea finding the reservation using the fewest queries. An example of a `host-reservation-identifiers` configuration looks as follows:

```

{
  "host-reservation-identifiers": [ "duid", "hw-address" ],
  "subnet6": [

```

(continues on next page)

(continued from previous page)

```

    {
        "subnet": "2001:db8:1::/64",
        ...
    }
],
...
}
```

If not specified, the default value is:

```
"host-reservation-identifiers": [ "hw-address", "duid" ]
```

Note

As soon as a host reservation is found, the search is stopped; when a client has two host reservations using different enabled identifier types, the first is always returned and the second ignored. This is usually a configuration error. In those rare cases when having two reservations for the same host makes sense, the one to be used can be specified by ordering the list of identifier types in `host-reservation-identifiers`.

10.3.8 Global Reservations in DHCPv6

In some deployments, such as mobile networks, clients can roam within the network and certain parameters must be specified regardless of the client's current location. To meet such a need, Kea offers a global reservation mechanism. The idea behind it is that regular host reservations are tied to specific subnets, by using a specific subnet ID. Kea can specify a global reservation that can be used in every subnet that has global reservations enabled.

This feature can be used to assign certain parameters, such as hostname or other dedicated, host-specific options. It can also be used to assign addresses or prefixes.

An address assigned via global host reservation must be feasible for the subnet the server selects for the client. In other words, the address must lie within the subnet; otherwise, it is ignored and the server will attempt to dynamically allocate an address. If the selected subnet belongs to a shared network, the server checks for feasibility against the subnet's siblings, selecting the first in-range subnet. If no such subnet exists, the server falls back to dynamically allocating the address. This does not apply to globally reserved prefixes.

Note

Prior to release 2.3.5, the server did not perform feasibility checks on globally reserved addresses, which allowed the server to be configured to hand out nonsensical leases for arbitrary address values. Later versions of Kea perform these checks.

To use global host reservations, a configuration similar to the following can be used:

```

"Dhcp6": {
    # This specifies global reservations.
    # They will apply to all subnets that
    # have global reservations enabled.

    "reservations": [
        {
            "hw-address": "aa:bb:cc:dd:ee:ff",
```

(continues on next page)

(continued from previous page)

```

    "hostname": "hw-host-dynamic"
  },
  {
    "hw-address": "01:02:03:04:05:06",
    "hostname": "hw-host-fixed",

    # Use of IP addresses in global reservations is risky.
    # If used outside of matching subnet, such as 3001::/64,
    # it will result in a broken configuration being handed
    # to the client.
    "ip-address": "2001:db8:ff::77"
  },
  {
    "duid": "01:02:03:04:05",
    "hostname": "duid-host"
  }
],
"valid-lifetime": 600,
"subnet4": [ {
  "subnet": "2001:db8:1::/64",
  # Specify if the server should look up global reservations.
  "reservations-global": true,
  # Specify if the server should look up in-subnet reservations.
  "reservations-in-subnet": false,
  # Specify if the server can assume that all reserved addresses
  # are out-of-pool. It can be ignored because "reservations-in-subnet"
  # is false.
  # "reservations-out-of-pool": false,
  "pools": [ { "pool": "2001:db8:1::-2001:db8:1::100" } ]
} ]
}

```

When using database backends, the global host reservations are distinguished from regular reservations by using a `subnet-id` value of 0.

10.3.9 Pool Selection with Client Class Reservations

Client classes can be specified both in the Kea configuration file and/or via host reservations. The classes specified in the Kea configuration file are evaluated immediately after receiving the DHCP packet and therefore can be used to influence subnet selection using the `client-classes` parameter specified in the subnet scope. The classes specified within the host reservations are fetched and assigned to the packet after the server has already selected a subnet for the client. This means that the client class specified within a host reservation cannot be used to influence subnet assignment for this client, unless the subnet belongs to a shared network. If the subnet belongs to a shared network, the server may dynamically change the subnet assignment while trying to allocate a lease. If the subnet does not belong to a shared network, once selected, the subnet is not changed once selected.

If the subnet does not belong to a shared network, it is possible to use host-reservation-based client classification to select a pool within the subnet as follows:

```

{
  "Dhcp6": {
    "client-classes": [
      {

```

(continues on next page)

(continued from previous page)

```

        "name": "reserved_class"
    },
    {
        "name": "unreserved_class",
        "test": "not member('reserved_class')"
    }
],
"subnet6": [
    {
        "id": 1,
        "subnet": "2001:db8:1::/64",
        "reservations": [
            {
                "hw-address": "aa:bb:cc:dd:ee:fe",
                "client-classes": [ "reserved_class" ]
            }
        ],
        "pools": [
            {
                "pool": "2001:db8:1::10 - 2001:db8:1::20",
                "client-classes": [ "unreserved_class" ]
            },
            {
                "pool": "2001:db8:1::30 - 2001:db8:1::40",
                "client-classes": [ "reserved_class" ]
            },
            {
                "pool": "2001:db8:1::50 - 2001:db8:1::60"
            }
        ]
    }
]
}

```

`reserved_class` is declared without the `test` parameter because it may be only assigned to a client via host reservation mechanism. The second class, `unreserved_class`, is assigned to clients which do not belong to `reserved_class`.

The first pool in the subnet is used for clients not having such a reservation. The second pool is only used for clients having a reservation for `reserved_class`. The third pool is an unrestricted pool for any clients, comprising of both `reserved_class` clients and `unreserved_class`.

The configuration snippet includes one host reservation which causes the client with the MAC address `aa:bb:cc:dd:ee:fe` to be assigned to `reserved_class`. Thus, this client will be given an IP address from the second address pool.

Reservations defined on a subnet that belongs to a shared network are not visible to an otherwise matching client, so they cannot be used to select pools, nor subnets for that matter.

10.3.10 Subnet Selection with Client Class Reservations

There is one specific use case when subnet selection may be influenced by client classes specified within host reservations: when the client belongs to a shared network. In such a case it is possible to use classification to select a subnet within this shared network. Consider the following example:

```
{
  "Dhcp6": {
    "client-classes": [
      {
        "name": "reserved_class"
      },
      {
        "name": "unreserved_class",
        "test": "not member('reserved_class')"
      }
    ],
    "reservations": [
      {
        "hw-address": "aa:bb:cc:dd:ee:fe",
        "client-classes": [ "reserved_class" ]
      }
    ],
    "reservations-global": true,
    "reservations-in-subnet": false,
    "shared-networks": [
      {
        "name": "net",
        "subnet6": [
          {
            "id": 1,
            "subnet": "2001:db8:1::/64",
            "pools": [
              {
                "pool": "2001:db8:1::10 - 2001:db8:1::20",
                "client-classes": [ "unreserved_class" ]
              },
              {
                "pool": "2001:db8:1::30 - 2001:db8:1::40",
                "client-classes": [ "unreserved_class" ]
              }
            ]
          },
          {
            "id": 2,
            "subnet": "2001:db8:2::/64",
            "pools": [
              {
                "pool": "2001:db8:2::10 - 2001:db8:2::20",
                "client-classes": [ "reserved_class" ]
              },
              {
                "pool": "2001:db8:2::30 - 2001:db8:2::40",
                "client-classes": [ "reserved_class" ]
              }
            ]
          }
        ]
      }
    ]
  }
}
```

(continues on next page)

```
    }  
  ],  
  {  
    "id": 3,  
    "subnet": "2001:db8:3::/64",  
    "pools": [  
      {  
        "pool": "2001:db8:3::10 - 2001:db8:3::20"  
      },  
      {  
        "pool": "2001:db8:3::30 - 2001:db8:3::40"  
      }  
    ]  
  }  
]  
}  
}
```

This is similar to the example described in the *Pool Selection with Client Class Reservations*. This time, however, there are three subnets, of which the first two have a pool associated with a different class each.

The clients that do not have a reservation for `reserved_class` are assigned an address from the first subnet and when that is filled from the third subnet. Clients with a reservation for `reserved_class` are assigned an address from the second subnet and when that is filled from the third subnet.

The subnets must belong to the same shared network.

For a subnet to be restricted to a certain class, or skipped, all of the pools inside that subnet must be guarded by `reserved_class` or `unreserved_class` respectively.

In addition, the reservation for the client class must be specified at the global scope (global reservation) and `reservations-global` must be set to `true`.

In the example above, the `client-classes` configuration parameter could also be specified at the subnet level rather than the pool level, and would yield the same effect.

If the subnets were defined outside shared networks, and `client-classes` were specified at the subnet level, then `early-global-reservations-lookup` would also need to be enabled in order for subnet selection to work.

10.3.11 Multiple Reservations for the Same IP

Host reservations were designed to preclude the creation of multiple reservations for the same IP address or delegated prefix within a particular subnet, to avoid having two different clients compete for the same lease. When using the default settings, the server returns a configuration error when it finds two or more reservations for the same lease within a subnet in the Kea configuration file. `libdhcp_host_cmds.so` returns an error in response to the `reservation-add` command when it detects that the reservation exists in the database for the lease for which the new reservation is being added.

Similar to DHCPv4 (see [Multiple Reservations for the Same IP](#)), the DHCPv6 server can also be configured to allow the creation of multiple reservations for the same IPv6 address and/or delegated prefix in a given subnet. This is supported since Kea release 1.9.1 as an optional mode of operation enabled with the `ip-reservations-unique` global parameter.

`ip-reservations-unique` is a boolean parameter that defaults to `true`, which forbids the specification of more than one reservation for the same lease in a given subnet. Setting this parameter to `false` allows such reservations to be created both in the Kea configuration file and in the host database backend, via [libdhcp_host_cmds.so](#).

Setting `ip-reservations-unique` to `false` when using `memfile`, `MySQL`, or `PostgreSQL` is supported. This setting is not supported when using `Host Cache` (see [libdhcp_host_cache.so: Host Cache Reservations for Improved Performance](#)) or the `RADIUS` backend (see [libdhcp_radius.so: RADIUS Server Support](#)). These reservation backends do not support multiple reservations for the same IP; if either of these hooks is loaded and `ip-reservations-unique` is set to `false`, then a configuration error is emitted and the server fails to start.

Note

When `ip-reservations-unique` is set to `true` (the default value), the server ensures that IP reservations are unique for a subnet within a single host backend and/or Kea configuration file. It does not guarantee that the reservations are unique across multiple backends. On server startup, only IP reservations defined in the Kea configuration file are checked for uniqueness.

The following is an example configuration with two reservations for the same IPv6 address but different MAC addresses:

```
"Dhcp6": {
  "ip-reservations-unique": false,
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "reservations": [
        {
          "hw-address": "1a:1b:1c:1d:1e:1f",
          "ip-address": "2001:db8:1::11"
        },
        {
          "hw-address": "2a:2b:2c:2d:2e:2f",
          "ip-address": "2001:db8:1::11"
        }
      ]
    }
  ]
}
```

It is possible to control the `ip-reservations-unique` parameter via the [Configuration Backend in DHCPv6](#). If the new setting of this parameter conflicts with the currently used backends (i.e. backends do not support the new setting), the new setting is ignored and a warning log message is generated. The backends continue to use the default setting, expecting that IP reservations are unique within each subnet. To allow the creation of non-unique IP reservations, the administrator must remove the backends which lack support for them from the configuration file.

Administrators must be careful when they have been using multiple reservations for the same IP address and/or delegated prefix and later decide to return to the default mode in which this is no longer allowed. They must make sure that at most one reservation for a given IP address or delegated prefix exists within a subnet, prior to switching back to the default mode. If such duplicates are left in the configuration file, the server reports a configuration error. Leaving such reservations in the host databases does not cause configuration errors but may lead to lease allocation errors during the server's operation, when it unexpectedly finds multiple reservations for the same IP address or delegated prefix.

Note

Currently, the Kea server does not verify whether multiple reservations for the same IP address and/or delegated prefix exist in MySQL and/or PostgreSQL) host databases when `ip-reservations-unique` is updated from `false` to `true`. This may cause issues with lease allocations. The administrator must ensure that there is at most one reservation for each IP address and/or delegated prefix within each subnet, prior to the configuration update.

`reservations-lookup-first` is a boolean parameter which controls whether host reservations lookup should be performed before lease lookup. This parameter has effect only when multi-threading is disabled. When multi-threading is enabled, host reservations lookup is always performed first to avoid lease-lookup resource locking. The `reservations-lookup-first` parameter defaults to `false` when multi-threading is disabled.

10.3.12 Host Reservations as Basic Access Control

It is possible to define a host reservation that contains just an identifier, without any address, options, or values. In some deployments this is useful, as the hosts that have a reservation belong to the `KNOWN` class while others do not. This can be used as a basic access control mechanism.

The following example demonstrates this concept. It indicates a single IPv6 subnet and all clients will get an address from it. However, only known clients (those that have reservations) will get their default DNS server configured. Empty reservations, i.e. reservations that only have the identification criterion, can be useful as a way of making the clients known.

```
"Dhcp6": {
  "client-classes": [
    {
      "name": "KNOWN",
      "option-data": [
        {
          "name": "dns-servers",
          "data": "2001:db8::1"
        }
      ]
    }
  ],
  "reservations": [
    // Clients on this list will be added to the KNOWN class.
    { "duid": "01:02:03:04:05:0A:0B:0C:0D:0E" },
    { "duid": "02:03:04:05:0A:0B:0C:0D:0E:0F" }
  ],
  "reservations-in-subnet": true,

  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pools": [
        {
          "pool": "2001:db8:1:1::/64"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

}

This concept can be extended further. A good real-life scenario might be a situation where some customers of an ISP have not paid their bills. A new class can be defined to use an alternative default DNS server that, instead of giving access to the Internet, redirects those customers to a captive portal urging them to bring their accounts up to date.

```
"Dhcp6": {
  "client-classes": [
    {
      "name": "blocked",
      "option-data": [
        {
          "name": "dns-servers",
          "data": "2001:db8::2"
        }
      ]
    }
  ],
  "reservations": [
    // Clients on this list will be added to the KNOWN class. Some
    // will also be added to the blocked class.
    { "duid": "01:02:03:04:05:0A:0B:0C:0D:0E",
      "client-classes": [ "blocked" ] },
    { "duid": "02:03:04:05:0A:0B:0C:0D:0E:0F" }
  ],
  "reservations-in-subnet": true,

  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/48",
      "pools": [
        {
          "pool": "2001:db8:1:1::/64"
        }
      ],
      "option-data": [
        {
          "name": "dns-servers",
          "data": "2001:db8::1"
        }
      ]
    }
  ]
}
```

10.4 Shared Networks in DHCPv6

DHCP servers use subnet information in two ways. It is used to both determine the point of attachment, i.e. where the client is connected to the network, and to group information pertaining to a specific location in the network. Sometimes it is useful to have more than one logical IP subnet being deployed on the same physical link. Understanding that two or more subnets are used on the same link requires additional logic in the DHCP server. This capability is called "shared

networks" in Kea, and sometimes also "shared subnets"; in Microsoft's nomenclature it is called "multinet."

There are many cases where the shared networks feature is useful; here we explain just a handful of the most common ones. The first and by far most common use case is an existing IPv4 network that has grown and is running out of available address space. This is less common in IPv6, but shared networks are still useful: for example, with the exhaustion of IPv6- delegated prefixes within a subnet, or the desire to experiment with an addressing scheme. With the advent of IPv6 deployment and a vast address space, many organizations split the address space into subnets, deploy it, and then after a while discover that they want to split it differently. In the transition period, they want both the old and new addressing to be available: thus the need for more than one subnet on the same physical link.

Finally, the case of cable networks is directly applicable in IPv6. There are two types of devices in cable networks: cable modems and the end-user devices behind them. It is a common practice to use different subnets for cable modems to prevent users from tinkering with them. In this case, the distinction is based on the type of device, rather than on address-space exhaustion.

A client connected to a shared network may be assigned a lease (address or prefix) from any of the pools defined within the subnets belonging to the shared network. Internally, the server selects one of the subnets belonging to a shared network and tries to allocate a lease from this subnet. If the server is unable to allocate a lease from the selected subnet (e.g., due to pool exhaustion), it uses another subnet from the same shared network and tries to allocate a lease from this subnet. The server typically allocates all leases available in a given subnet before it starts allocating leases from other subnets belonging to the same shared network. However, in certain situations the client can be allocated a lease from another subnet before the pools in the first subnet get exhausted; this sometimes occurs when the client provides a hint that belongs to another subnet, or the client has reservations in a subnet other than the default.

Note

Deployments should not assume that Kea waits until it has allocated all the addresses from the first subnet in a shared network before allocating addresses from other subnets.

To define a shared network, an additional configuration scope is introduced:

```
{
"Dhcp6": {
  "shared-networks": [
    {
      # Name of the shared network. It may be an arbitrary string
      # and it must be unique among all shared networks.
      "name": "ipv6-lab-1",

      # The subnet selector can be specified on the shared network
      # level. Subnets from this shared network will be selected
      # for clients communicating via relay agent having
      # the specified IP address.
      "relay": {
        "ip-addresses": [ "2001:db8:2:34::1" ]
      },

      # This starts a list of subnets in this shared network.
      # There are two subnets in this example.
      "subnet6": [
        {
          "id": 1,
          "subnet": "2001:db8::/48",
          "pools": [ { "pool": "2001:db8::1 - 2001:db8::ffff" } ]
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    },
    {
        "id": 2,
        "subnet": "3ffe:ffe::/64",
        "pools": [ { "pool": "3ffe:ffe::/64" } ]
    }
]
}
],
# end of shared-networks

# It is likely that in the network there will be a mix of regular,
# "plain" subnets and shared networks. It is perfectly valid
# to mix them in the same configuration file.
#
# This is a regular subnet. It is not part of any shared-network.
"subnet6": [
    {
        "id": 3,
        "subnet": "2001:db9::/48",
        "pools": [ { "pool": "2001:db9::/64" } ],
        "relay": {
            "ip-addresses": [ "2001:db8:1:2::1" ]
        }
    }
]
}
}

```

As demonstrated in the example, it is possible to mix shared and regular ("plain") subnets. Each shared network must have a unique name. This is similar to the ID for subnets, but gives administrators more flexibility. It is used for logging, but also internally for identifying shared networks.

In principle it makes sense to define only shared networks that consist of two or more subnets. However, for testing purposes, an empty subnet or a network with just a single subnet is allowed. This is not a recommended practice in production networks, as the shared network logic requires additional processing and thus lowers the server's performance. To avoid unnecessary performance degradation, shared subnets should only be defined when required by the deployment.

Shared networks provide an ability to specify many parameters in the shared network scope that apply to all subnets within it. If necessary, it is possible to specify a parameter in the shared-network scope and then override its value in the subnet scope. For example:

```

{
  "shared-networks": [
    {
      "name": "lab-network3",
      "relay": {
        "ip-addresses": [ "2001:db8:2:34::1" ]
      },
      # This applies to all subnets in this shared network, unless
      # values are overridden on subnet scope.
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

"valid-lifetime": 600,

# This option is made available to all subnets in this shared
# network.
"option-data": [ {
    "name": "dns-servers",
    "data": "2001:db8::8888"
} ],

"subnet6": [
    {
        "id": 1,
        "subnet": "2001:db8:1::/48",
        "pools": [ { "pool": "2001:db8:1::1 - 2001:db8:1::ffff" } ],

        # This particular subnet uses different values.
        "valid-lifetime": 1200,
        "option-data": [
            {
                "name": "dns-servers",
                "data": "2001:db8::1:2"
            },
            {
                "name": "unicast",
                "data": "2001:abcd::1"
            }
        ]
    },
    {
        "id": 2,
        "subnet": "2001:db8:2::/48",
        "pools": [ { "pool": "2001:db8:2::1 - 2001:db8:2::ffff" } ],

        # This subnet does not specify its own valid-lifetime value,
        # so it is inherited from shared network scope.
        "option-data": [
            {
                "name": "dns-servers",
                "data": "2001:db8:cafe::1"
            }
        ]
    }
]
}
...
}

```

In this example, there is a `dns-servers` option defined that is available to clients in both subnets in this shared network. Also, the valid lifetime is set to 10 minutes (600s). However, the first subnet overrides some of the values (the valid lifetime is 20 minutes, there is a different IP address for `dns-servers`), but also adds its own option (the unicast address). Assuming a client asking for server unicast and `dns-servers` options is assigned a lease from this subnet, it will get a lease for 20 minutes and `dns-servers`, and be allowed to use server unicast at address `2001:abcd::1`. If the same client is assigned to the second subnet, it will get a 10-minute lease, a `dns-servers` value of `2001:db8:cafe::1`, and no server unicast.

Some parameters must be the same in all subnets in the same shared network. This restriction applies to the `interface` and `rapid-commit` settings. The most convenient way is to define them on the shared-network scope, but they can be specified for each subnet. However, each subnet must have the same value.

Note

There is an inherent ambiguity when using clients that send multiple IA options in a single request, and shared-networks whose subnets have different values for options and configuration parameters. The server sequentially processes IA options in the order that they occur in the client's query; if the leases requested in the IA options end up being fulfilled from different subnets, which parameters and options should apply? Currently, the code uses the values from the last subnet of the last IA option fulfilled.

We view this largely as a site configuration issue. A shared network generally means the same physical link, so services configured by options from subnet A should be as easily reachable from subnet B and vice versa. There are a number of ways to avoid this situation:

- Use the same values for options and parameters for subnets within the shared network.
- Use subnet selectors or client class guards that ensure that for a single client's query, the same subnet is used for all IA options in that query.
- Avoid using shared networks with clients that send multiple IA options per query.

10.4.1 Local and Relayed Traffic in Shared Networks

It is possible to specify an interface name at the shared-network level, to tell the server that this specific shared network is reachable directly (not via relays) using the local network interface. As all subnets in a shared network are expected to be used on the same physical link, it is a configuration error to attempt to define a shared network using subnets that are reachable over different interfaces. In other words, all subnets within the shared network must have the same value for the `interface` parameter. The following configuration is an example of what **NOT** to do:

```
{
  "shared-networks": [
    {
      "name": "office-floor-2",
      "subnet6": [
        {
          "id": 1,
          "subnet": "2001:db8::/64",
          "pools": [ { "pool": "2001:db8::1 - 2001:db8::ffff" } ],
          "interface": "eth0"
        },
        {
          "id": 2,
          "subnet": "3ffe:abcd::/64",
          "pools": [ { "pool": "3ffe:abcd::1 - 3ffe:abcd::ffff" } ],
          ...
          # Specifying a different interface name is a configuration
          # error. This value should rather be "eth0" or the interface
          # name in the other subnet should be "eth1".
          # "interface": "eth1"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
],
...
}
```

To minimize the chance of configuration errors, it is often more convenient to simply specify the interface name once, at the shared-network level, as shown in the example below.

```
{
"shared-networks": [
  {
    "name": "office-floor-2",

    # This tells Kea that the whole shared network is reachable over a
    # local interface. This applies to all subnets in this network.
    "interface": "eth0",

    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8::/64",
        "pools": [ { "pool": "2001:db8::1 - 2001:db8::ffff" } ]
      },
      {
        "id": 2,
        "subnet": "3ffe:abcd::/64",
        "pools": [ { "pool": "3ffe:abcd::1 - 3ffe:abcd::ffff" } ]
      }
    ]
  }
],
...
}
```

With relayed traffic, subnets are typically selected using the relay agents' addresses. If the subnets are used independently (not grouped within a shared network), a different relay address can be specified for each of these subnets. When multiple subnets belong to a shared network they must be selected via the same relay address and, similarly to the case of the local traffic described above, it is a configuration error to specify different relay addresses for the respective subnets in the shared network. The following configuration is another example of what **NOT** to do:

```
{
"shared-networks": [
  {
    "name": "kakapo",
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8::/64",
        "relay": {
          "ip-addresses": [ "2001:db8::1234" ]
        },
        "pools": [ { "pool": "2001:db8::1 - 2001:db8::ffff" } ]
      },
    ]
  }
],
...
}
```

(continues on next page)

(continued from previous page)

```

    {
        "id": 2,
        "subnet": "3ffe:abcd::/64",
        "pools": [ { "pool": "3ffe:abcd::1 - 3ffe:abcd::ffff" } ],
        "relay": {
            # Specifying a different relay address for this
            # subnet is a configuration error. In this case
            # it should be 2001:db8::1234 or the relay address
            # in the previous subnet should be 3ffe:abcd::cafe.
            "ip-addresses": [ "3ffe:abcd::cafe" ]
        }
    }
]
}
...
}

```

Again, it is better to specify the relay address at the shared-network level; this value will be inherited by all subnets belonging to the shared network.

```

{
  "shared-networks": [
    {
      "name": "kakapo",
      "relay": {
        # This relay address is inherited by both subnets.
        "ip-addresses": [ "2001:db8::1234" ]
      },
      "subnet6": [
        {
          "id": 1,
          "subnet": "2001:db8::/64",
          "pools": [ { "pool": "2001:db8::1 - 2001:db8::ffff" } ]
        },
        {
          "id": 2,
          "subnet": "3ffe:abcd::/64",
          "pools": [ { "pool": "3ffe:abcd::1 - 3ffe:abcd::ffff" } ]
        }
      ]
    }
  ]
}
...
}

```

Even though it is technically possible to configure two (or more) subnets within the shared network to use different relay addresses, this will almost always lead to a different behavior than what the user would expect. In this case, the Kea server will initially select one of the subnets by matching the relay address in the client's packet with the subnet's configuration. However, it MAY end up using the other subnet (even though it does not match the relay address) if the client already has a lease in this subnet or has a host reservation in this subnet, or simply if the initially selected subnet has no more addresses available. Therefore, it is strongly recommended to always specify subnet selectors (interface or relay address) at the shared-network level if the subnets belong to a shared network, as it is rarely useful to specify

them at the subnet level and may lead to the configuration errors described above.

10.4.2 Client Classification in Shared Networks

Sometimes it is desirable to segregate clients into specific subnets based on certain properties. This mechanism is called client classification and is described in *Client Classification*. Client classification can be applied to subnets belonging to shared networks in the same way as it is used for subnets specified outside of shared networks. It is important to understand how the server selects subnets for clients when client classification is in use, to ensure that the appropriate subnet is selected for a given client type.

If a subnet is associated with a class, only the clients belonging to this class can use this subnet. If there are no classes specified for a subnet, any client connected to a given shared network can use this subnet. A common mistake is to assume that the subnet that includes a client class is preferred over subnets without client classes.

The `client-classes` parameter may be specified at the shared network, subnet, and/or pool scopes. If specified for a shared network, clients must belong to at least one of the classes specified for that network to be considered for subnets within that network. If specified for a subnet, clients must belong to at least one of the classes specified for that subnet to be considered for any of that subnet's pools or host reservations. If specified for a pool, clients must belong to at least one of the classes specified for that pool to be given a lease from that pool.

Consider the following example:

```
{
  "client-classes": [
    {
      "name": "b-devices",
      "test": "option[1234].hex == 0x0002"
    }
  ],
  "shared-networks": [
    {
      "name": "galah",
      "relay": {
        "ip-address": [ "2001:db8:2:34::1" ]
      },
      "subnet6": [
        {
          "id": 1,
          "subnet": "2001:db8:1::/64",
          "pools": [ { "pool": "2001:db8:1::20 - 2001:db8:1::ff" } ]
        },
        {
          "id": 2,
          "subnet": "2001:db8:3::/64",
          "pools": [ { "pool": "2001:db8:3::20 - 2001:db8:3::ff" } ],
          "client-classes": "b-devices"
        }
      ]
    }
  ]
}
```

If the client belongs to the "b-devices" class (because it includes option 1234 with a value of 0x0002), that does not guarantee that the subnet 2001:db8:3::/64 will be used (or preferred) for this client. The server can use either of the two subnets, because the subnet 2001:db8:1::/64 is also allowed for this client. The client classification used in this case should be perceived as a way to restrict access to certain subnets, rather than as a way to express subnet preference.

For example, if the client does not belong to the "b-devices" class, it may only use the subnet 2001:db8:1::/64 and will never use the subnet 2001:db8:3::/64.

A typical use case for client classification is in a cable network, where cable modems should use one subnet and other devices should use another subnet within the same shared network. In this case it is necessary to apply classification on all subnets. The following example defines two classes of devices, and the subnet selection is made based on option 1234 values.

```
{
  "client-classes": [
    {
      "name": "a-devices",
      "test": "option[1234].hex == 0x0001"
    },
    {
      "name": "b-devices",
      "test": "option[1234].hex == 0x0002"
    }
  ],
  "shared-networks": [
    {
      "name": "galah",
      "relay": {
        "ip-addresses": [ "2001:db8:2:34::1" ]
      },
      "subnet6": [
        {
          "id": 1,
          "subnet": "2001:db8:1::/64",
          "pools": [ { "pool": "2001:db8:1:20 - 2001:db8:1::ff" } ],
          "client-classes": [ "a-devices" ]
        },
        {
          "id": 2,
          "subnet": "2001:db8:3::/64",
          "pools": [ { "pool": "2001:db8:3:20 - 2001:db8:3::ff" } ],
          "client-classes": [ "b-devices" ]
        }
      ]
    }
  ]
}
```

In this example each class has its own restriction. Only clients that belong to class "a-devices" are able to use subnet 2001:db8:1::/64 and only clients belonging to "b-devices" are able to use subnet 2001:db8:3::/64. Care should be taken not to define too-restrictive classification rules, as clients that are unable to use any subnets will be refused service. However, this may be a desired outcome if one wishes to provide service only to clients with known properties (e.g. only VoIP phones allowed on a given link).

It is possible to achieve an effect similar to the one presented in this section without the use of shared networks. If the subnets are placed in the global subnets scope, rather than in the shared network, the server will still use classification rules to pick the right subnet for a given class of devices. The major benefit of placing subnets within the shared network is that common parameters for the logically grouped subnets can be specified once, in the shared network scope, e.g. the interface or relay parameter. All subnets belonging to this shared network will inherit those parameters.

10.4.3 Host Reservations in Shared Networks

Subnets that are part of a shared network allow host reservations, similar to regular subnets:

```
{
  "shared-networks": [
    {
      "name": "frog",
      "relay": {
        "ip-addresses": [ "2001:db8:2:34::1" ]
      },
      "subnet6": [
        {
          "subnet": "2001:db8:1::/64",
          "id": 100,
          "pools": [ { "pool": "2001:db8:1::1 - 2001:db8:1::64" } ],
          "reservations": [
            {
              "duid": "00:03:00:01:11:22:33:44:55:66",
              "ip-addresses": [ "2001:db8:1::28" ]
            }
          ]
        },
        {
          "subnet": "2001:db8:3::/64",
          "id": 101,
          "pools": [ { "pool": "2001:db8:3::1 - 2001:db8:3::64" } ],
          "reservations": [
            {
              "duid": "00:03:00:01:aa:bb:cc:dd:ee:ff",
              "ip-addresses": [ "2001:db8:2::28" ]
            }
          ]
        }
      ]
    }
  ]
}
```

It is worth noting that Kea conducts additional checks when processing a packet if shared networks are defined. First, instead of simply checking whether there is a reservation for a given client in its initially selected subnet, Kea looks through all subnets in a shared network for a reservation. This is one of the reasons why defining a shared network may impact performance. If there is a reservation for a client in any subnet, that particular subnet is picked for the client. Although it is technically not an error, it is considered bad practice to define reservations for the same host in multiple subnets belonging to the same shared network.

While not strictly mandatory, it is strongly recommended to use explicit "id" values for subnets if database storage will be used for host reservations. If an ID is not specified, the values for it are auto generated, i.e. Kea assigns increasing integer values starting from 1. Thus, the auto-generated IDs are not stable across configuration changes.

10.5 Server Identifier in DHCPv6

The DHCPv6 protocol uses a "server identifier" (also known as a DUID) to allow clients to discriminate between several servers present on the same link. [RFC 8415](#) currently defines four DUID types: DUID-LLT, DUID-EN, DUID-LL, and DUID-UUID.

The Kea DHCPv6 server generates a server identifier once, upon the first startup, and stores it in a file. This identifier is not modified across restarts of the server and so is a stable identifier.

Kea follows the recommendation from [RFC 8415](#) to use DUID-LLT as the default server identifier. However, ISC has received reports that some deployments require different DUID types, and that there is a need to administratively select both the DUID type and/or its contents.

The server identifier can be configured using parameters within the `server-id` map element in the global scope of the Kea configuration file. The following example demonstrates how to select DUID-EN as a server identifier:

```
"Dhcp6": {
  "server-id": {
    "type": "EN"
  },
  ...
}
```

Currently supported values for the `type` parameter are: "LLT", "EN", and "LL", for DUID-LLT, DUID-EN, and DUID-LL respectively.

When a new DUID type is selected, the server generates its value and replaces any existing DUID in the file. The server then uses the new server identifier in all future interactions with clients.

Note

If the new server identifier is created after some clients have obtained their leases, the clients using the old identifier are not able to renew their leases; the server will ignore messages containing the old server identifier. Clients will continue sending RENEW until they transition to the rebinding state. In this state, they will start sending REBIND messages to the multicast address without a server identifier. The server will respond to the REBIND messages with a new server identifier, and the clients will associate the new server identifier with their leases. Although the clients will be able to keep their leases and will eventually learn the new server identifier, this will be at the cost of an increased number of renewals and multicast traffic due to a need to rebind. Therefore, it is recommended that modification of the server-identifier type and value be avoided if the server has already assigned leases and these leases are still valid.

There are cases when an administrator needs to explicitly specify a DUID value rather than allow the server to generate it. The following example demonstrates how to explicitly set all components of a DUID-LLT.

```
"Dhcp6": {
  "server-id": {
    "type": "LLT",
    "htype": 8,
    "identifier": "A65DC7410F05",
    "time": 2518920166
  },
  ...
}
```

where:

- `htype` is a 16-bit unsigned value specifying hardware type,
- `identifier` is a link-layer address, specified as a string of hexadecimal digits, and
- `time` is a 32-bit unsigned time value.

The hexadecimal representation of the DUID generated as a result of the configuration specified above is:

```
00:01:00:08:96:23:AB:E6:A6:5D:C7:41:0F:05
|type| htype|   time   |   identifier   |
```

A special value of "0" for `htype` and `time` is allowed, which indicates that the server should use ANY value for these components. If the server already uses a DUID-LLT, it will use the values from this DUID; if the server uses a DUID of a different type or does not yet use any DUID, it will generate these values. Similarly, if the `identifier` is assigned an empty string, the value of the `identifier` will be generated. Omitting any of these parameters is equivalent to setting them to those special values.

For example, the following configuration:

```
"Dhcp6": {
  "server-id": {
    "type": "LLT",
    "htype": 0,
    "identifier": "",
    "time": 2518920166
  },
  ...
}
```

indicates that the server should use ANY link-layer address and hardware type. If the server is already using DUID-LLT, it will use the link-layer address and hardware type from the existing DUID. If the server is not yet using any DUID, it will use the link-layer address and hardware type from one of the available network interfaces. The server will use an explicit value of time; if it is different than a time value present in the currently used DUID, that value will be replaced, effectively modifying the current server identifier.

The following example demonstrates an explicit configuration of a DUID-EN:

```
"Dhcp6": {
  "server-id": {
    "type": "EN",
    "enterprise-id": 2495,
    "identifier": "87ABEF7A5BB545"
  },
  ...
}
```

where:

- `enterprise-id` is a 32-bit unsigned value holding an enterprise number, and
- `identifier` is a variable-length identifier within DUID-EN.

The hexadecimal representation of the DUID-EN created according to the configuration above is:

```
00:02:00:00:09:BF:87:AB:EF:7A:5B:B5:45
|type| ent-id |   identifier   |
```

As in the case of the DUID-LLT, special values can be used for the configuration of the DUID-EN. If the `enterprise-id` is "0", the server will use a value from the existing DUID-EN. If the server is not using any DUID

or the existing DUID has a different type, the ISC enterprise ID will be used. When an empty string is entered for `identifier`, the identifier from the existing DUID-EN will be used. If the server is not using any DUID-EN, a new 6-byte-long identifier will be generated.

DUID-LL is configured in the same way as DUID-LLT except that the `time` parameter has no effect for DUID-LL, because this DUID type only comprises a hardware type and link-layer address. The following example demonstrates how to configure DUID-LL:

```
"Dhcp6": {
  "server-id": {
    "type": "LL",
    "htype": 8,
    "identifier": "A65DC7410F05"
  },
  ...
}
```

which will result in the following server identifier:

```
00:03:00:08:A6:5D:C7:41:0F:05
|type|htype|  identifier  |
```

The server stores the generated server identifier in the following location: `[kea-install-dir]/var/lib/kea/kea-dhcp6-serverid`.

In some uncommon deployments where no stable storage is available, the server should be configured not to try to store the server identifier. This choice is controlled by the value of the `persist` boolean parameter:

```
"Dhcp6": {
  "server-id": {
    "type": "EN",
    "enterprise-id": 2495,
    "identifier": "87ABEF7A5BB545",
    "persist": false
  },
  ...
}
```

The default value of the `persist` parameter is `true`, which configures the server to store the server identifier on a disk.

In the example above, the server is configured not to store the generated server identifier on a disk. But if the server identifier is not modified in the configuration, the same value is used after server restart, because the entire server identifier is explicitly specified in the configuration.

10.6 DHCPv6 Data Directory

The Kea DHCPv6 server puts the server identifier file and the default memory lease file into its data directory. By default this directory is `prefix/var/lib/kea` but this location can be changed using the `data-directory` global parameter, as in:

```
"Dhcp6": {
  "data-directory": "/var/lib/kea/kea-server6",
  ...
}
```

Note

As of Kea 2.7.9, `data-directory` is deprecated. The `duid` and `lease` files may only be loaded from the directory determined at compilation: `"[kea-install-dir]/var/lib/kea"`. This path may be overridden at startup by setting the environment variable `KEA_DHCP_DATA_DIR` to the desired path. If a path other than this value is used in `data-directory`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error.

10.7 Stateless DHCPv6 (INFORMATION-REQUEST Message)

Typically DHCPv6 is used to assign both addresses and options. These assignments (leases) have a state that changes over time, hence their description as "stateful." DHCPv6 also supports a "stateless" mode, where clients request only configuration options. This mode is considered lightweight from the server perspective, as it does not require any state tracking.

The Kea server supports stateless mode. When clients send INFORMATION-REQUEST messages, the server sends back answers with the requested options, if they are available in the server configuration. The server attempts to use per-subnet options first; if that fails, it then tries to provide options defined in the global scope.

Stateless and stateful mode can be used together. No special configuration directives are required to handle this; simply use the configuration for stateful clients and the stateless clients will get only the options they requested.

It is possible to run a server that provides only options and no addresses or prefixes. If the options have the same value in each subnet, the configuration can define the required options in the global scope and skip subnet definitions altogether. Here's a simple example of such a configuration:

```
"Dhcp6": {
  "interfaces-config": {
    "interfaces": [ "ethX" ]
  },
  "option-data": [ {
    "name": "dns-servers",
    "data": "2001:db8::1, 2001:db8::2"
  } ],
  "lease-database": {
    "type": "memfile"
  }
}
```

This very simple configuration provides DNS server information to all clients in the network, regardless of their location. The `memfile` lease database must be specified, as Kea requires a lease database to be specified even if it is not used.

10.8 Support for RFC 7550 (now part of RFC 8415)

[RFC 7550](#) introduced some changes to the previous DHCPv6 specifications, [RFC 3315](#) and [RFC 3633](#), to resolve issues with the coexistence of multiple stateful options in the messages sent between clients and servers. Those changes were later included in the most recent DHCPv6 protocol specification, [RFC 8415](#), which obsoleted [RFC 7550](#). Kea supports [RFC 8415](#) along with these protocol changes, which are briefly described below.

When a client, such as a requesting router, requests an allocation of both addresses and prefixes during the 4-way (SARR) exchange with the server, and the server is not configured to allocate any prefixes but can allocate some addresses, it will respond with the `IA_NA(s)` containing allocated addresses and the `IA_PD(s)` containing the `NoPrefix-Avail` status code. According to the updated specifications, if the client can operate without prefixes it should accept

allocated addresses and transition to the "bound" state. When the client subsequently sends RENEW/REBIND messages to the server to extend the lifetimes of the allocated addresses, according to the T1 and T2 times, and if the client is still interested in obtaining prefixes from the server, it may also include an IA_PD in the RENEW/REBIND to request allocation of the prefixes. If the server still cannot allocate the prefixes, it will respond with the IA_PD(s) containing the NoPrefixAvail status code. However, if the server can allocate the prefixes, it allocates and sends them in the IA_PD(s) to the client. A similar situation occurs when the server is unable to allocate addresses for the client but can delegate prefixes: the client may request allocation of the addresses while renewing the delegated prefixes. Allocating leases for other IA types while renewing existing leases is by default supported by the Kea DHCPv6 server, and the server provides no configuration mechanisms to disable this behavior.

The following are the other behaviors first introduced in [RFC 7550](#) (now part of [RFC 8415](#)) and supported by the Kea DHCPv6 server:

- Set T1/T2 timers to the same value for all stateful (IA_NA and IA_PD) options to facilitate renewal of all of a client's leases at the same time (in a single message exchange).
- Place NoAddrsAvail and NoPrefixAvail status codes in the IA_NA and IA_PD options in the ADVERTISE message, rather than as the top-level options.

10.9 Using a Specific Relay Agent for a Subnet

The DHCPv6 server follows the same principles as the DHCPv4 server to select a subnet for the client, with noticeable differences mainly for relays.

Note

When the selected subnet is a member of a shared network, the whole shared network is selected.

A relay must have an interface connected to the link on which the clients are being configured. Typically the relay has a global IPv6 address configured on that interface, which belongs to the subnet from which the server assigns addresses. Normally, the server is able to use the IPv6 address inserted by the relay (in the `link-addr` field in the RELAY-FORW message) to select the appropriate subnet.

However, that is not always the case; the relay address may not match the subnet in certain deployments. This usually means that there is more than one subnet allocated for a given link. The two most common examples of this are long-lasting network renumbering (where both the old and new address spaces are still being used) and a cable network. In a cable network, both cable modems and the devices behind them are physically connected to the same link, yet they use distinct addressing. In such a case, the DHCPv6 server needs additional information (the value of the `interface-id` option or the IPv6 address inserted in the `link-addr` field in the RELAY-FORW message) to properly select an appropriate subnet.

The following example assumes that there is a subnet 2001:db8:1::/64 that is accessible via a relay that uses 3000::1 as its IPv6 address. The server is able to select this subnet for any incoming packets that come from a relay that has an address in the 2001:db8:1::/64 subnet. It also selects that subnet for a relay with address 3000::1.

```
"Dhcp6": {
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "pools": [
        {
          "pool": "2001:db8:1::1-2001:db8:1::ffff"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "relay": {
        "ip-addresses": [ "3000::1" ]
    }
}
]
}

```

If `relay` is specified, the `ip-addresses` parameter within it is mandatory. The `ip-addresses` parameter supports specifying a list of addresses.

10.10 Segregating IPv6 Clients in a Cable Network

In certain cases, it is useful to mix relay address information (introduced in *Using a Specific Relay Agent for a Subnet*) with client classification (explained in *Client Classification*). One specific example is in a cable network, where modems typically get addresses from a different subnet than all the devices connected behind them.

Let us assume that there is one Cable Modem Termination System (CMTS) with one CM MAC (a physical link that modems are connected to). We want the modems to get addresses from the 3000::/64 subnet, while everything connected behind the modems should get addresses from the 2001:db8:1::/64 subnet. The CMTS that acts as a relay uses address 3000::1. The following configuration can serve that situation:

```

"Dhcp6": {
    "subnet6": [
        {
            "id": 1,
            "subnet": "3000::/64",
            "pools": [
                { "pool": "3000::2 - 3000::ffff" }
            ],
            "client-classes": [ "VENDOR_CLASS_docsis3.0" ],
            "relay": {
                "ip-addresses": [ "3000::1" ]
            }
        },
        {
            "id": 2,
            "subnet": "2001:db8:1::/64",
            "pools": [
                {
                    "pool": "2001:db8:1::1-2001:db8:1::ffff"
                }
            ],
            "relay": {
                "ip-addresses": [ "3000::1" ]
            }
        }
    ]
}

```


10.11 MAC/Hardware Addresses in DHCPv6

MAC/hardware addresses are available in DHCPv4 messages from clients, and administrators frequently use that information to perform certain tasks like per-host configuration and address reservation for specific MAC addresses. Unfortunately, the DHCPv6 protocol does not provide any completely reliable way to retrieve that information. To mitigate that issue, a number of mechanisms have been implemented in Kea. Each of these mechanisms works in certain cases, but may not in others. Whether the mechanism works in a particular deployment is somewhat dependent on the network topology and the technologies used.

Kea allows specification of which of the supported methods should be used and in what order, via the `mac-sources` parameter. This configuration may be considered a fine tuning of the DHCP deployment.

Here is an example:

```
"Dhcp6": {
  "mac-sources": [
    "method1",
    "method2",
    "method3",
    ...
  ],

  "subnet6": [
    {
      ...
    },
    ...
  ],
  ...
}
```

When not specified, a value of "any" is used, which instructs the server to attempt to try all the methods in sequence and use the value returned by the first one that succeeds. In a typical deployment the default value of "any" is sufficient and there is no need to select specific methods. Changing the value of this parameter is most useful in cases when an administrator wants to disable certain methods; for example, if the administrator trusts the network infrastructure more than the information provided by the clients themselves, they may prefer information provided by the relays over that provided by clients.

If specified, `mac-sources` must have at least one value.

Supported methods are:

- **any** - this is not an actual method, just a keyword that instructs Kea to try all other methods and use the first one that succeeds. This is the default operation if no `mac-sources` are defined.
- **raw** - in principle, a DHCPv6 server could use raw sockets to receive incoming traffic and extract MAC/hardware address information. This is currently not implemented for DHCPv6 and this value has no effect.
- **duid** - DHCPv6 uses DUID identifiers instead of MAC addresses. There are currently four DUID types defined, and two of them (DUID-LLT, which is the default, and DUID-LL) convey MAC address information. Although [RFC 8415](#) forbids it, it is possible to parse those DUIDs and extract necessary information from them. This method is not completely reliable, as clients may use other DUID types, namely DUID-EN or DUID-UUID.
- **ipv6-link-local** - another possible acquisition method comes from the source IPv6 address. In typical usage, clients are sending their packets from IPv6 link-local addresses. There is a good chance that those addresses are based on EUI-64, which contains a MAC address. This method is not completely reliable, as clients may use other link-local address types. In particular, privacy extensions, defined in [RFC 4941](#), do not use MAC addresses.

Also note that successful extraction requires that the address's u-bit must be set to "1" and its g-bit set to "0", indicating that it is an interface identifier as per [RFC 2373, section 2.5.1](#).

- **client-link-addr-option** - one extension defined to alleviate missing MAC issues is the client link-layer address option, defined in [RFC 6939](#). This is an option that is inserted by a relay and contains information about a client's MAC address. This method requires a relay agent that supports the option and is configured to insert it. This method is useless for directly connected clients. The value `rfc6939` is an alias for `client-link-addr-option`.
- **remote-id** - [RFC 4649](#) defines a `remote-id` option that is inserted by a relay agent. Depending on the relay agent configuration, the inserted option may convey the client's MAC address information. The value `rfc4649` is an alias for `remote-id`.
- **subscriber-id** - Defined in [RFC 4580](#), `subscriber-id` is somewhat similar to `remote-id`; it is also inserted by a relay agent. The value `rfc4580` is an alias for `subscriber-id`. This method is currently not implemented.
- **docsis-cmts** - Yet another possible source of MAC address information are the DOCSIS options inserted by a CMTS that acts as a DHCPv6 relay agent in cable networks. This method attempts to extract MAC address information from sub-option 1026 (cm mac) of the vendor-specific option with `vendor-id=4491`. This vendor option is extracted from the Relay-forward message, not the original client's message.
- **docsis-modem** - The final possible source of MAC address information are the DOCSIS options inserted by the cable modem itself. This method attempts to extract MAC address information from sub-option 36 (`device-id`) of the vendor-specific option with `vendor-id=4491`. This vendor option is extracted from the original client's message, not from any relay options.

An empty `mac-sources` parameter is not allowed. Administrators who do not want to specify it should either simply omit the `mac-sources` definition or specify it with the "any" value, which is the default.

10.12 Duplicate Addresses (DHCPDECLINE Support)

The DHCPv6 server is configured with a certain pool of addresses that it is expected to hand out to DHCPv6 clients. It is assumed that the server is authoritative and has complete jurisdiction over those addresses. However, for various reasons such as misconfiguration or a faulty client implementation that retains its address beyond the valid lifetime, there may be devices connected that use those addresses without the server's approval or knowledge.

Such an unwelcome event can be detected by legitimate clients (using Duplicate Address Detection) and reported to the DHCPv6 server using a DHCPDECLINE message. The server does a sanity check (to see whether the client declining an address really was supposed to use it), then conducts a clean-up operation, and confirms the DHCPDECLINE by sending back a REPLY message. Any DNS entries related to that address are removed, the event is logged, and hooks are triggered. After that is complete, the address is marked as declined (which indicates that it is used by an unknown entity and thus not available for assignment) and a probation time is set on it. Unless otherwise configured, the probation period lasts 24 hours; after that time, the server will recover the lease (i.e. put it back into the available state) and the address will be available for assignment again. It should be noted that if the underlying issue of a misconfigured device is not resolved, the duplicate-address scenario will repeat. If reconfigured correctly, this mechanism provides an opportunity to recover from such an event automatically, without any system administrator intervention.

To configure the decline probation period to a value other than the default, the following syntax can be used:

```
"Dhcp6": {
  "decline-probation-period": 3600,
  "subnet6": [
    {
      ...
    },
    ...
  ],
}
```

(continues on next page)

(continued from previous page)

```
} ...
```

The parameter is expressed in seconds, so the example above instructs the server to recycle declined leases after one hour.

There are several statistics and hook points associated with the decline handling procedure. The `lease6_decline` hook point is triggered after the incoming DHCPDECLINE message has been sanitized and the server is about to decline the lease. The `declined-addresses` statistic is increased after the hook returns (both the global and subnet-specific variants). (See *Statistics in the DHCPv6 Server* and *Hook Libraries* for more details on DHCPv6 statistics and Kea hook points.)

Once the probation time elapses, the declined lease is recovered using the standard expired-lease reclamation procedure, with several additional steps. In particular, both `declined-addresses` statistics (global and subnet-specific) are decreased. At the same time, `reclaimed-declined-addresses` statistics (again in two variants, global and subnet-specific) are increased.

A note about statistics: The Kea server does not decrease the `assigned-nas` statistics when a DHCPDECLINE message is received and processed successfully. While technically a declined address is no longer assigned, the primary usage of the `assigned-nas` statistic is to monitor pool utilization. Most people would forget to include `declined-addresses` in the calculation, and would simply use `assigned-nas/total-nas`. This would cause a bias towards under-representing pool utilization. As this has a potential to cause serious confusion, ISC decided not to decrease `assigned-nas` immediately after receiving DHCPDECLINE, but to do it later when Kea recovers the address back to the available pool.

10.13 Address Registration (RFC 9686 Support)

Kea version 2.7.7 introduces the support of self-generated address registration as defined in [RFC 9686](#) i.e. when a valid ADDR-REG-INFORM (36) message is received a registered lease is added or updated and a ADDR-REG-REPLY (37) is sent back to the client.

Note

Even if they share a common lease database with leases in other states, registered leases are independent: when a lease in another state already exists for an address this address is considered as in use and can't be registered. Similarly a registered lease can't change to another state, e.g. reclamation of expired registered leases removes them.

Note

Kea accepts and handles the Option Request option in the ADDR-REG-INFORM message (the RFC specifies the client MUST NOT put such option in it).

Note

Kea currently only accepts addresses which are "appropriate to the link" but not yet "within a prefix delegated to the client" which is not compatible with the way the topology is described in the Kea configuration.

10.14 Statistics in the DHCPv6 Server

The DHCPv6 server supports the following statistics:

Table 5: DHCPv6 statistics

Statistic	Data Type	Description
pkt6-received	integer	Number of DHCPv6 packets received. This includes all packets: valid, bogus, corrupted, rejected, etc. This statistic is expected to grow rapidly.
pkt6-receive-drop	integer	Number of incoming packets that were dropped. The exact reason for dropping packets is logged, but the most common reasons may be that an unacceptable or not-supported packet type is received, direct responses are forbidden, the server ID sent by the client does not match the server's server ID, or the packet is malformed.
pkt6-parse-failed	integer	Number of incoming packets that could not be parsed. A non-zero value of this statistic indicates that the server received a malformed or truncated packet. This may indicate problems in the network, faulty clients, faulty relay agents, or a bug in the server.
pkt6-solicit-received	integer	Number of SOLICIT packets received. This statistic is expected to grow; its increase means that clients that just booted started their configuration process and their initial packets reached the Kea server.
pkt6-advertise-received	integer	Number of ADVERTISE packets received. ADVERTISE packets are sent by the server and the server is never expected to receive them; a non-zero value of this statistic indicates an error occurring in the network. One likely cause would be a misbehaving relay agent that incorrectly forwards ADVERTISE messages towards the server, rather than back to the clients.
pkt6-request-received	integer	Number of DHCPREQUEST packets received. This statistic is expected to grow. Its increase means that clients that just booted received the server's response (DHCPADVERTISE) and accepted it, and are now requesting an address (DHCPREQUEST).
pkt6-reply-received	integer	Number of REPLY packets received. This statistic is expected to remain zero at all times, as REPLY packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards REPLY messages towards the server, rather than back to the clients.
pkt6-renew-received	integer	Number of RENEW packets received. This statistic is expected to grow; its increase means that clients received their addresses and prefixes and are trying to renew them.
pkt6-rebind-received	integer	Number of REBIND packets received. A non-zero value indicates that clients did not receive responses to their RENEW messages (through the regular lease-renewal mechanism) and are attempting to find any server that is able to take over their leases. It may mean that some servers' REPLY messages never reached the clients.
pkt6-release-received	integer	Number of RELEASE packets received. This statistic is expected to grow when a device is being shut down in the network; it indicates that the address or prefix assigned is reported as no longer needed. Note that many devices, especially wireless, do not send RELEASE packets either because of design choices or due to the client moving out of range.
pkt6-decline-received	integer	Number of DECLINE packets received. This statistic is expected to remain close to zero. Its increase means that a client leased an address, but discovered that the address is currently used by an unknown device in the network. If this statistic is growing, it may indicate a misconfigured server or devices that have statically assigned conflicting addresses.

continues on next page

Table 5 – continued from previous page

Statistic	Data Type	Description
pkt6-infrequest-received	integer	Number of INFORMATION-REQUEST packets received. This statistic is expected to grow if there are devices that are using stateless DHCPv6. INFORMATION-REQUEST messages are used by clients that request stateless configuration, i.e. options and parameters other than addresses or prefixes.
pkt6-dhcpv4-query-received	integer	Number of DHCPv4-QUERY packets received. This statistic is expected to grow if there are devices that are using DHCPv4-over-DHCPv6. DHCPv4-QUERY messages are used by DHCPv4 clients on an IPv6-only line which encapsulates the requests over DHCPv6.
pkt6-dhcpv4-response-received	integer	Number of DHCPv4-RESPONSE packets received. This statistic is expected to remain zero at all times, as DHCPv4-RESPONSE packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards DHCPv4-RESPONSE message towards the server rather than back to the clients.
pkt6-addr-reg-inform-received	integer	Number of ADDR-REG-INFORM packets received. This statistic is expected to grow if there are devices that are using address registration.
pkt6-addr-reg-reply-received	integer	Number of ADDR-REG-REPLY packets received. This statistic is expected to remain zero at all times, as ADDR-REG-REPLY packets are sent by the server and the server is never expected to receive them. A non-zero value indicates an error. One likely cause would be a misbehaving relay agent that incorrectly forwards ADDR-REG-REPLY message towards the server rather than back to the clients.
pkt6-unknown-received	integer	Number of packets received of an unknown type. A non-zero value of this statistic indicates that the server received a packet that it was unable to recognize; either it had an unsupported type or was possibly malformed.
pkt6-sent	integer	Number of DHCPv6 packets sent. This statistic is expected to grow every time the server transmits a packet. In general, it should roughly match <code>pkt6-received</code> , as most incoming packets cause the server to respond. There are exceptions (e.g. a server receiving a REQUEST with server ID matching another server), so do not worry if it is less than <code>pkt6-received</code> .
pkt6-advertise-sent	integer	Number of ADVERTISE packets sent. This statistic is expected to grow in most cases after a SOLICIT is processed. There are certain uncommon but valid cases where incoming SOLICIT packets are dropped, but in general this statistic is expected to be close to <code>pkt6-solicit-received</code> .
pkt6-reply-sent	integer	Number of REPLY packets sent. This statistic is expected to grow in most cases after a SOLICIT (with rapid-commit), REQUEST, RENEW, REBIND, RELEASE, DECLINE, or INFORMATION-REQUEST is processed. There are certain cases where there is no response.
pkt6-dhcpv4-response-sent	integer	Number of DHCPv4-RESPONSE packets sent. This statistic is expected to grow in most cases after a DHCPv4-QUERY is processed. There are certain cases where there is no response.
pkt6-addr-reg-reply-sent	integer	Number of ADDR-REG-REPLY packets sent. This statistic is expected to grow in most cases after a ADDR-REG-INFORM is processed. There are certain cases where there is no response.
subnet[id].total-nas	big integer	Total number of NA addresses available for DHCPv6 management for a given subnet; in other words, this is the count of all addresses in all configured pools. This statistic changes only during configuration changes. It does not take into account any addresses that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.

continues on next page

Table 5 – continued from previous page

Statistic	Data Type	Description
subnet[id].pool[pid].total-nas	big integer	Total number of NA addresses available for DHCPv6 management for a given subnet pool; in other words, this is the count of all addresses in configured subnet pool. This statistic changes only during configuration changes. It does not take into account any addresses that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of a given pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
cumulative-assigned-nas	integer	Cumulative number of NA addresses that have been assigned since server startup. It is incremented each time a NA address is assigned and is not reset when the server is reconfigured.
subnet[id].cumulative-assigned-nas	integer	Cumulative number of NA addresses in a given subnet that were assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is never decreased. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pool[pid].cum-assigned-nas	integer	Cumulative number of NA addresses in a given subnet pool that were assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is never decreased. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of a given pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
subnet[id].assigned-nas	integer	Number of NA addresses in a given subnet that are assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is decreased every time a lease is released (a RELEASE message is received) or expires. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pool[pid].assigned-nas	integer	Number of NA addresses in a given subnet pool that are assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is decreased every time a lease is released (a RELEASE message is received) or expires. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
subnet[id].total-pds	big integer	Total number of PD prefixes available for DHCPv6 management for a given subnet; in other words, this is the count of all prefixes in all configured pools. This statistic changes only during configuration changes. Note that it does not take into account any prefixes that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pd-pool[pid].total-pds	big integer	Total number of PD prefixes available for DHCPv6 management for a given subnet pool; in other words, this is the count of all prefixes in a configured subnet PD pool. This statistic changes only during configuration changes. It does not take into account any prefixes that may be reserved due to host reservation. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of a given pool. This statistic is exposed for each subnet PD pool separately, and is reset during a reconfiguration event.
cumulative-assigned-pds	integer	Cumulative number of PD prefixes that have been assigned since server startup. It is incremented each time a PD prefix is assigned and is not reset when the server is reconfigured.
subnet[id].cumulative-assigned-pds	integer	Cumulative number of PD prefixes in a given subnet that were assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is never decreased. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.

continues on next page

Table 5 – continued from previous page

Statistic	Data Type	Description
subnet[id].pd-pool[pid].cumulative-assigned-pds	integer	Cumulative number of PD prefixes in a given subnet PD pool that were assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is never decreased. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of a given PD pool. This statistic is exposed for each subnet PD pool separately, and is reset during a reconfiguration event.
subnet[id].assigned-pds	integer	Number of PD prefixes in a given subnet that are assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is decreased every time a lease is released (a RELEASE message is received) or expires. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].pd-pool[pid].assigned-pds	integer	Number of PD prefixes in a given subnet pd-pool that are assigned. It increases every time a new lease is allocated (as a result of receiving a REQUEST message) and is decreased every time a lease is released (a RELEASE message is received) or expires. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the PD pool. This statistic is exposed for each subnet PD pool separately, and is reset during a reconfiguration event.
reclaimed-leases	integer	Number of expired leases that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed (counting both NA and PD reclamations). This statistic never decreases. It can be used as a long-term indicator of how many actual leases have been reclaimed. This is a global statistic that covers all subnets.
subnet[id].reclaimed-leases	integer	Number of expired leases associated with a given subnet that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed (counting both NA and PD reclamations). The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
subnet[id].pool[pid].reclaimed-leases	integer	Number of expired NA addresses associated with a given subnet pool that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately, and is reset during a reconfiguration event.
subnet[id].pd-pool[pid].reclaimed-leases	integer	Number of expired PD prefixes associated with a given subnet PD pool that have been reclaimed since server startup. It is incremented each time an expired lease is reclaimed. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the PD pool. This statistic is exposed for each subnet PD pool separately, and is reset during a reconfiguration event.
declined-addresses	integer	Number of IPv6 addresses that are currently declined; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. This is a global statistic that covers all subnets.
subnet[id].declined-addresses	integer	Number of IPv6 addresses that are currently declined in a given subnet; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.

continues on next page

Table 5 – continued from previous page

Statistic	Data Type	Description
subnet[id].pool[pid].declined-addresses	integer	Number of IPv6 addresses that are currently declined in a given subnet pool; a count of the number of leases currently unavailable. Once a lease is recovered, this statistic is decreased; ideally, this statistic should be zero. If this statistic is non-zero or increasing, a network administrator should investigate whether there is a misbehaving device in the network. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately.
reclaimed-declined-addresses	integer	Number of IPv6 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. This is a global statistic that covers all subnets.
subnet[id].reclaimed-declined-addresses	integer	Number of IPv6 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
subnet[id].pool[pid].reclaimed-declined-addresses	integer	Number of IPv6 addresses that were declined, but have now been recovered. Unlike declined-addresses , this statistic never decreases. It can be used as a long-term indicator of how many actual valid declines were processed and recovered from. The <i>id</i> is the subnet ID of a given subnet. The <i>pid</i> is the pool ID of the pool. This statistic is exposed for each subnet pool separately.
v6-allocation-fail	integer	Number of total address allocation failures for a particular client. This consists of the number of lease allocation attempts that the server made before giving up, if it was unable to use any of the address pools. This is a global statistic that covers all subnets.
subnet[id].v6-allocation-fail	integer	Number of total address allocation failures for a particular client. This consists of the number of lease allocation attempts that the server made before giving up, if it was unable to use any of the address pools. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v6-allocation-fail-shared-network	integer	Number of address allocation failures for a particular client connected to a shared network. This is a global statistic that covers all subnets.
subnet[id].v6-allocation-fail-shared-network	integer	Number of address allocation failures for a particular client connected to a shared network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v6-allocation-fail-subnet	integer	Number of address allocation failures for a particular client connected to a subnet that does not belong to a shared network. This is a global statistic that covers all subnets.
subnet[id].v6-allocation-fail-subnet	integer	Number of address allocation failures for a particular client connected to a subnet that does not belong to a shared network. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v6-allocation-fail-no-pools	integer	Number of address allocation failures because the server could not use any configured pools for a particular client. It is also possible that all of the subnets from which the server attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects that some clients should be assigned dynamic addresses. This is a global statistic that covers all subnets.

continues on next page

Table 5 – continued from previous page

Statistic	Data Type	Description
subnet[id].v6-allocation-fail-no-pools	integer	Number of address allocation failures because the server could not use any configured pools for a particular client. It is also possible that all of the subnets from which the server attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects that some clients should be assigned dynamic addresses. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v6-allocation-fail-classes	integer	Number of address allocation failures when the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned: for example, if all pools require packets to belong to certain classes and the incoming packet does not belong to any. Another case where this information may be useful is to indicate that the pool reserved for a given class has run out of addresses. This is a global statistic that covers all subnets.
subnet[id].v6-allocation-fail-classes	integer	Number of address allocation failures when the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned: for example, if all pools require packets to belong to certain classes and the incoming packet does not belong to any. Another case where this information may be useful is to indicate that the pool reserved for a given class has run out of addresses. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately.
v6-ia-na-lease-reuses	integer	Number of times an IA_NA lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is global.
subnet[id].v6-ia-na-lease-reuses	integer	Number of times an IA_NA lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is on a per-subnet basis. The <i>id</i> is the subnet ID of a given subnet.
v6-ia-pd-lease-reuses	integer	Number of times an IA_PD lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is global.
subnet[id].v6-ia-pd-lease-reuses	integer	Number of times an IA_PD lease had its CLTT increased in memory and its expiration time left unchanged in persistent storage, as part of the lease caching feature. This is referred to as a lease reuse. This statistic is on a per-subnet basis. The <i>id</i> is the subnet ID of a given subnet.
cumulative-registered	integer	Cumulative number of NA addresses that have been registered since server startup. It is incremented each time a NA address is registered and is not reset when the server is reconfigured.
subnet[id].cumulative-registered	integer	Cumulative number of NA addresses in a given subnet that were registered. It increases every time a new address is registered (as a result of receiving an ADDR-REG-INFORM message) and is never decreased. The <i>id</i> is the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.
subnet[id].registered	integer	Number of NA addresses in a given subnet that are registered. It increases every time a new address registered (as a result of receiving an ADDR-REG-INFORM message) and is decreased every time a registration expires. The <i>id</i> is the the subnet ID of a given subnet. This statistic is exposed for each subnet separately, and is reset during a reconfiguration event.

Note

The pool ID can be configured on each pool by explicitly setting the `pool-id` parameter in the pool parameter map. If not configured, `pool-id` defaults to 0. The statistics related to pool ID 0 refer to all the statistics of all the pools

that have an unconfigured `pool-id`. The pool ID does not need to be unique within the subnet or across subnets. The statistics regarding a specific pool ID within a subnet are combined with the other statistics of all other pools with the same pool ID in the respective subnet.

Note

This section describes DHCPv6-specific statistics. For a general overview and usage of statistics, see [Statistics](#).

The DHCPv6 server provides two global parameters to control the default sample limits of statistics:

- `statistic-default-sample-count` - determines the default maximum number of samples to be kept. The special value of 0 indicates that a default maximum age should be used.
- `statistic-default-sample-age` - determines the default maximum age, in seconds, of samples to be kept.

For instance, to reduce the statistic-keeping overhead, set the default maximum sample count to 1 so only one sample is kept:

```
"Dhcp6": {
  "statistic-default-sample-count": 1,
  "subnet6": [
    {
      ...
    },
    ...
  ],
  ...
}
```

Statistics can be retrieved periodically to gain more insight into Kea operations. One tool that leverages that capability is ISC Stork. See [Monitoring Kea With Stork](#) for details.

10.15 Management API for the DHCPv6 Server

The management API allows the issuing of specific management commands, such as statistics retrieval, reconfiguration, or shutdown. For more details, see [Management API](#). By default there are no sockets open; to instruct Kea to open a socket, the following entry in the configuration file can be used:

```
"Dhcp6": {
  "control-sockets": [
    {
      "socket-type": "unix",
      "socket-name": "/path/to/the/unix/socket"
    }
  ],
  "subnet6": [
    {
      ...
    },
    ...
  ],
}
```

(continues on next page)

(continued from previous page)

```

    ...
}

```

10.15.1 UNIX Control Socket

Until Kea server 2.7.2 the only supported communication channel type was the UNIX stream socket with `socket-type` set to `unix` and `socket-name` to the file path of the UNIX/LOCAL socket.

The length of the path specified by the `socket-name` parameter is restricted by the maximum length for the UNIX socket name on the administrator's operating system, i.e. the size of the `sun_path` field in the `sockaddr_un` structure, decreased by 1. This value varies on different operating systems, between 91 and 107 characters. Typical values are 107 on Linux and 103 on FreeBSD.

Kea supports only one `unix` control socket in the "control-sockets" list.

Note

As of Kea 2.7.9, control sockets may only reside in the directory determined during compilation as "[`kea-install-dir`]/var/run/kea", which must also have `0750` access rights. This path may be overridden at startup by setting the environment variable `KEA_CONTROL_SOCKET_DIR` to the desired path. If a path other than this value is used in `socket-name`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use simply omit the path component from `socket-name`.

Communication over the control channel is conducted using JSON structures. See the [Control Channel](#) section in the *Kea Developer's Guide* for more details.

The DHCPv6 server supports the following operational commands:

- `build-report`
- `config-get`
- `config-hash-get`
- `config-reload`
- `config-set`
- `config-test`
- `config-write`
- `dhcp-disable`
- `dhcp-enable`
- `leases-reclaim`
- `list-commands`
- `shutdown`
- `status-get`
- `version-get`

as described in *Commands Supported by Both the DHCPv4 and DHCPv6 Servers*. In addition, it supports the following statistics-related commands:

- `statistic-get`

- `statistic-reset`
- `statistic-remove`
- `statistic-get-all`
- `statistic-reset-all`
- `statistic-remove-all`
- `statistic-sample-age-set`
- `statistic-sample-age-set-all`
- `statistic-sample-count-set`
- `statistic-sample-count-set-all`

as described in *Commands for Manipulating Statistics*.

10.15.2 HTTP/HTTPS Control Socket

The `socket-type` must be `http` or `https` (when the type is `https` TLS is required). The `socket-address` (default `:::1`) and `socket-port` (default 8000) specify an IP address and port to which the HTTP service will be bound.

Since Kea 2.7.5 the `http-headers` parameter specifies a list of extra HTTP headers to add to HTTP responses.

The `trust-anchor`, `cert-file`, `key-file`, and `cert-required` parameters specify the TLS setup for HTTP, i.e. HTTPS. If these parameters are not specified, HTTP is used. The TLS/HTTPS support in Kea is described in *TLS/HTTPS Support*.

Basic HTTP authentication protects against unauthorized uses of the control agent by local users. For protection against remote attackers, HTTPS and reverse proxy of *Secure Connections* provide stronger security.

The authentication is described in the `authentication` block with the mandatory `type` parameter, which selects the authentication. Currently only the basic HTTP authentication (type `basic`) is supported.

The `realm` authentication parameter (default `kea-dhcpv6-server`) is used for error messages when the basic HTTP authentication is required but the client is not authorized.

When the `clients` authentication list is configured and not empty, basic HTTP authentication is required. Each element of the list specifies a user ID and a password. The user ID is mandatory, must not be empty, and must not contain the colon (:) character. The password is optional; when it is not specified an empty password is used.

Note

The basic HTTP authentication user ID and password are encoded in UTF-8, but the current Kea JSON syntax only supports the Latin-1 (i.e. 0x00..0xff) Unicode subset.

To avoid exposing the user ID and/or the associated password, these values can be read from files. The syntax is extended by:

- The `directory` authentication parameter, which handles the common part of file paths. The default value is the empty string.
- The `password-file` client parameter, which, alongside the `directory` parameter, specifies the path of a file that can contain the password, or when no user ID is given, the whole basic HTTP authentication secret.
- The `user-file` client parameter, which, with the `directory` parameter, specifies the path of a file where the user ID can be read.

Since Kea-2.7.6 Kea supports multiple HTTP/HTTPS connections. Both IPv4 and IPv6 addresses can be used. The server will issue an error when changing the socket type from HTTP to HTTPS or from HTTPS to HTTP using the same address and port. This action is not allowed as it might introduce a security issue accidentally caused by a user mistake. A different address or port must be specified when using the "config-set" command to switch from HTTP to HTTPS or from HTTPS to HTTP. The same applies when modifying the configuration file and then running "config-reload" command.

When files are used, they are read when the configuration is loaded, to detect configuration errors as soon as possible.

```
"Dhcp6": {
  "control-sockets": [
    {
      "socket-type": "https",
      "socket-address": "2010:30:40::50",
      "socket-port": 8005,
      "http-headers": [
        {
          "name": "Strict-Transport-Security",
          "value": "max-age=31536000"
        }
      ],
      "trust-anchor": "/path/to/the/ca-cert.pem",
      "cert-file": "/path/to/the/agent-cert.pem",
      "key-file": "/path/to/the/agent-key.pem",
      "cert-required": true,
      "authentication": {
        "type": "basic",
        "realm": "kea-control-agent",
        "clients": [
          {
            "user": "admin",
            "password": "1234"
          }
        ]
      }
    },
    {
      "socket-type": "http",
      "socket-address": "10.20.30.40",
      "socket-port": 8005
    }
  ],
  "subnet6": [
    {
      ...
    },
    ...
  ],
  ...
}
```

10.16 User Contexts in IPv6

Kea allows the loading of hook libraries that can sometimes benefit from additional parameters. If such a parameter is specific to the whole library, it is typically defined as a parameter for the hook library. However, sometimes there is a need to specify parameters that are different for each pool.

See *Comments and User Context* for additional background regarding the user-context idea. See *User Contexts in Hooks* for a discussion from the hooks perspective.

User contexts can be specified at global scope; at the shared-network, subnet, pool, client-class, option-data, or definition level; and via host reservation. One other useful feature is the ability to store comments or descriptions.

Let's consider an example deployment of lightweight 4over6, an IPv6 transition technology that allows mapping IPv6 prefixes into full or partial IPv4 addresses. In the DHCP context, these are specific parameters that are supposed to be delivered to clients in the form of additional options. Values of these options are correlated to delegated prefixes, so it is reasonable to keep these parameters together with the prefix delegation (PD) pool. On the other hand, lightweight 4over6 is not a commonly used feature, so it is not a part of the base Kea code. The solution to this problem is to specify a user context. For each PD pool that is expected to be used for lightweight 4over6, a user context with extra parameters is defined. Those extra parameters will be used by a hook library and loaded only when dynamic calculation of the lightweight 4over6 option is actually needed. An example configuration looks as follows:

```
"Dhcp6": {
  "subnet6": [ {
    "pd-pools": [
      {
        "prefix": "2001:db8::",
        "prefix-len": 56,
        "delegated-len": 64,

        # This is a pool-specific context.
        "user-context": {
          "threshold-percent": 85,
          "v4-network": "192.168.0.0/16",
          "v4-overflow": "10.0.0.0/16",
          "lw4over6-sharing-ratio": 64,
          "lw4over6-v4-pool": "192.0.2.0/24",
          "lw4over6-sysports-exclude": true,
          "lw4over6-bind-prefix-len": 56
        }
      }
    ],
    "id": 1,
    "subnet": "2001:db8::/32",

    # This is a subnet-specific context. Any type of
    # information can be entered here as long as it is valid JSON.
    "user-context": {
      "comment": "Those v4-v6 migration technologies are tricky.",
      "experimental": true,
      "billing-department": 42,
      "contacts": [ "Alice", "Bob" ]
    }
  } ]
}
```

Kea does not interpret or use the user-context information; it simply stores it and makes it available to the hook libraries.

It is up to each hook library to extract that information and use it. The parser translates a `comment` entry into a user context with the entry, which allows a comment to be attached inside the configuration itself.

10.17 Supported DHCPv6 Standards

The following standards are currently supported in Kea:

- *Dynamic Host Configuration Protocol for IPv6*, [RFC 3315](#): Supported messages are SOLICIT, ADVERTISE, REQUEST, RELEASE, RENEW, REBIND, INFORMATION-REQUEST, CONFIRM, DECLINE and REPLY. The only unsupported message is RECONFIGURE. Almost all options are supported, except AUTHENTICATION and RECONFIGURE-ACCEPT.
- *Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers*, [RFC 3319](#): All defined options are supported.
- *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*, [RFC 3633](#): Supported options are IA_PD and IA_PREFIX. Also supported is the status code NoPrefixAvail.
- *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, [RFC 3646](#): All defined options are supported.
- *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*, [RFC 3736](#): Server operation in stateless mode is supported. Kea is currently server-only, so the client side is not implemented.
- *Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6*, [RFC 4075](#): The SNTP option is supported.
- *Renumbering Requirements for Stateless Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, [RFC 4076](#): The server supports all the stateless renumbering requirements.
- *Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, [RFC 4242](#): The sole defined option (`information-refresh-time`) is supported.
- *Dynamic Host Configuration Protocol (DHCP) Options for Broadcast and Multicast Control Servers*, [RFC 4280](#): The DHCPv6 options are supported.
- *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Subscriber-ID Option*, [RFC 4580](#): The subscribed-id option is supported and can be used in any expression.
- *The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option*, [RFC 4649](#): The Remote-ID option is supported.
- *A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)*, [RFC 4701](#): The DHCPv6 server supports DHCID records. The DHCP-DDNS server must be running to add, update, and/or delete DHCID records.
- *Resolution of Fully Qualified Domain Name (FQDN) Conflicts among Dynamic Host Configuration Protocol (DHCP) Clients*, [RFC 4703](#): The DHCPv6 server uses the DHCP-DDNS server to resolve conflicts.
- *The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option*, [RFC 4704](#): The supported option is CLIENT_FQDN.
- *Timezone Options for DHCP*: [RFC 4833](#): Both DHCPv6 options are supported.
- *DHCPv6 Leasequery*: [RFC 5007](#): The server functionality (message types, options) is supported. This requires the leasequery hook. See [libdhcp_lease_query.so: Leasequery Support](#) for details.
- *DHCP Options for Protocol for Carrying Authentication for Network Access (PANA) Authentication Agents*: [RFC 5192](#): The PANA option is supported.
- *Discovering Location-to-Service Translation (LoST) Servers Using the Dynamic Host Configuration Protocol (DHCP)*: [RFC 5223](#): The LOST option is supported.

- *Control And Provisioning of Wireless Access Points (CAPWAP) Access Controller DHCP Option*: [RFC 5417](#): The CAPWAP for IPv6 option is supported.
- *DHCPv6 Bulk Leasequery*: [RFC 5460](#): The server functionality (TCP connection, new message types and options, new query types) is supported. This requires the leasequery hook. See [libdhcp_lease_query.so: Leasequery Support](#) for details.
- *Network Time Protocol (NTP) Server Option for DHCPv6*: [RFC 5908](#): The NTP server option and its suboptions are supported. See [NTP Server Suboptions](#) for details.
- *DHCPv6 Options for Network Boot*: [RFC 5970](#): The network boot options are supported.
- *Lightweight DHCPv6 Relay Agent*, [RFC 6221](#): Kea can handle lightweight relay messages and use other methods than link address to perform subnet selection.
- *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*, [RFC 6334](#): The AFTR-Name DHCPv6 Option is supported.
- *Relay-Supplied DHCP Options*, [RFC 6422](#): The full functionality is supported, including OPTION_RSOO; the ability of the server to echo back the options; verification of whether an option is RSOO-enabled; and the ability to mark additional options as RSOO-enabled.
- *The EAP Re-authentication Protocol (ERP) Local Domain Name DHCPv6 Option*, [RFC 6440](#): The option is supported.
- *Prefix Exclude Option for DHCPv6-based Prefix Delegation*, [RFC 6603](#): The Prefix Exclude option is supported.
- *Client Link-Layer Address Option in DHCPv6*, [RFC 6939](#): The supported option is the client link-layer address option.
- *Modification to Default values of SOL_MAX_RT and INF_MAX_RT*, [RFC 7083](#): The new options are supported.
- *Issues and Recommendations with Multiple Stateful DHCPv6 Options*, [RFC 7550](#): All recommendations related to the DHCPv6 server operation are supported.
- *DHCPv6 Options for Configuration of Software Address and Port-Mapped Clients*, [RFC 7598](#): All options indicated in this specification are supported by the DHCPv6 server.
- *Generalized UDP Source Port for DHCP Relay*, [RFC 8357](#): The Kea server is able to handle the Relay Source Port option in a received Relay-forward message, remembers the UDP port, and sends back Relay-reply with a copy of the option to the relay agent using this UDP port.
- *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*, [RFC 8415](#): This new DHCPv6 protocol specification obsoletes RFC 3315, RFC 3633, RFC 3736, RFC 4242, RFC 7083, RFC 7283, and RFC 7550. All features, with the exception of the RECONFIGURE mechanism and the now-deprecated temporary addresses (IA_TA) mechanism, are supported.
- *Captive-Portal Identification in DHCP and Router Advertisements (RAs)*, [RFC 8910](#): The Kea server can configure both v4 and v6 versions of the captive portal options.
- *DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)*, [RFC 9463](#). The Kea server supports the DNR option.
- *Registering Self-Generated IPv6 Addresses Using DHCPv6*, [RFC 9686](#). The Kea server supports self-generated address registration. See [Address Registration \(RFC 9686 Support\)](#) for details.

10.18 DHCPv6 Server Limitations

These are the current known limitations of the Kea DHCPv6 server software. Most of them are reflections of the current stage of development and should be treated as “not yet implemented,” rather than actual limitations.

- The server allocates, renews, or rebinds a maximum of one lease for a particular IA option (IA_NA or IA_PD) sent by a client. [RFC 8415](#) allows for multiple addresses or prefixes to be allocated for a single IA.
- Temporary addresses are not supported. There is no intention to ever implement this feature, as it is deprecated in [RFC 8415](#).
- Client reconfiguration (RECONFIGURE) is not yet supported.

10.19 Kea DHCPv6 Server Examples

A collection of simple-to-use examples for the DHCPv6 component of Kea is available with the source files, located in the `doc/examples/kea6` directory.

10.20 Configuration Backend in DHCPv6

In the *Kea Configuration Backend* section we have described the Configuration Backend (CB) feature, its applicability, and its limitations. This section focuses on the usage of the CB with the DHCPv6 server. It lists the supported parameters, describes limitations, and gives examples of DHCPv6 server configurations to take advantage of the CB. Please also refer to the corresponding section *Configuration Backend in DHCPv4* for DHCPv4-specific usage of the CB.

10.20.1 Supported Parameters

The ultimate goal for the CB is to serve as a central configuration repository for one or multiple Kea servers connected to a database. In currently supported Kea versions, only a subset of the DHCPv6 server parameters can be configured in the database. All other parameters must be specified in the JSON configuration file, if required.

All supported parameters can be configured via `libdhcp_cb_cmds.so`. The general rule is that scalar global parameters are set using `remote-global-parameter6-set`; shared-network-specific parameters are set using `remote-network6-set`; and subnet-level and pool-level parameters are set using `remote-subnet6-set`. Whenever there is an exception to this general rule, it is highlighted in the table. Non-scalar global parameters have dedicated commands; for example, the global DHCPv6 options (`option-data`) are modified using `remote-option6-global-set`. Client classes, together with class-specific option definitions and DHCPv6 options, are configured using the `remote-class6-set` command.

The *Configuration Sharing and Server Tags* section explains the concept of shareable and non-shareable configuration elements and the limitations for sharing them between multiple servers. In the DHCP configuration (both DHCPv4 and DHCPv6), the shareable configuration elements are subnets and shared networks. Thus, they can be explicitly associated with multiple server tags. The global parameters, option definitions, and global options are non-shareable and can be associated with only one server tag. This rule does not apply to the configuration elements associated with all servers. Any configuration element associated with all servers (using the `all` keyword as a server tag) is used by all servers connecting to the configuration database.

The following table lists DHCPv6-specific parameters supported by the configuration backend, with an indication of the level of the hierarchy at which it is currently supported.

Table 6: List of DHCPv6 parameters supported by the configuration backend

Parameter	Global	Client Class	Shared Network	Subnet	Pool	Prefix D
allocator	yes	n/a	yes	yes	n/a	n/a
cache-max-age	yes	n/a	no	no	n/a	n/a
cache-threshold	yes	n/a	no	no	n/a	n/a
calculate-tee-times	yes	n/a	yes	yes	n/a	n/a
client-class	n/a	n/a	yes	yes	yes	yes

continues

Table 6 – continued from previous page

Parameter	Global	Client Class	Shared Network	Subnet	Pool	Prefix D
ddns-send-update	yes	n/a	yes	yes	n/a	n/a
ddns-override-no-update	yes	n/a	yes	yes	n/a	n/a
ddns-override-client-update	yes	n/a	yes	yes	n/a	n/a
ddns-replace-client-name	yes	n/a	yes	yes	n/a	n/a
ddns-generated-prefix	yes	n/a	yes	yes	n/a	n/a
ddns-qualifying-suffix	yes	n/a	yes	yes	n/a	n/a
decline-probation-period	yes	n/a	n/a	n/a	n/a	n/a
delegated-len	n/a	n/a	n/a	n/a	n/a	yes
dhcp4o6-port	yes	n/a	n/a	n/a	n/a	n/a
excluded-prefix	n/a	n/a	n/a	n/a	n/a	yes
excluded-prefix-len	n/a	n/a	n/a	n/a	n/a	yes
hostname-char-set	yes	n/a	yes	yes	n/a	n/a
hostname-char-replacement	yes	n/a	yes	yes	n/a	n/a
interface	n/a	n/a	yes	yes	n/a	n/a
interface-id	n/a	n/a	yes	yes	n/a	n/a
max-preferred-lifetime	yes	yes	yes	yes	n/a	n/a
max-valid-lifetime	yes	yes	yes	yes	n/a	n/a
min-preferred-lifetime	yes	yes	yes	yes	n/a	n/a
min-valid-lifetime	yes	yes	yes	yes	n/a	n/a
option-data	yes (via remote-option6-global-set)	yes	yes	yes	yes	yes
option-def	yes (via remote-option-def6-set)	yes	n/a	n/a	n/a	n/a
pd-allocator	yes	n/a	yes	yes	n/a	n/a
preferred-lifetime	yes	yes	yes	yes	n/a	n/a
prefix	n/a	n/a	n/a	n/a	n/a	yes
prefix-len	n/a	n/a	n/a	n/a	n/a	yes
rapid-commit	yes	n/a	yes	yes	n/a	n/a
rebind-timer	yes	n/a	yes	yes	n/a	n/a
relay	n/a	n/a	yes	yes	n/a	n/a
renew-timer	yes	n/a	yes	yes	n/a	n/a
require-client-classes	no	n/a	yes	yes	yes	yes
evaluate-additional-classes	no	n/a	yes	yes	yes	yes
reservations-global	yes	n/a	yes	yes	n/a	n/a
reservations-in-subnet	yes	n/a	yes	yes	n/a	n/a
reservations-out-of-pool	yes	n/a	yes	yes	n/a	n/a
t1-percent	yes	n/a	yes	yes	n/a	n/a
t2-percent	yes	n/a	yes	yes	n/a	n/a
valid-lifetime	yes	yes	yes	yes	n/a	n/a

- **yes** - indicates that the parameter is supported at the given level of the hierarchy and can be configured via the configuration backend.
- **no** - indicates that a parameter is supported at the given level of the hierarchy but cannot be configured via the configuration backend.
- **n/a** - indicates that a given parameter is not applicable at the particular level of the hierarchy or that the server does not support the parameter at that level.

Some scalar parameters contained by top level global maps are supported by the configuration backend.

Table 7: List of DHCPv6 map parameters supported by the configuration backend

Parameter name (flat naming format)	Global map	Parameter name
compatibility.lenient-option-parsing	compatibility	lenient-option-parsing
dhcp-ddns.enable-updates	dhcp-ddns	enable-updates
dhcp-ddns.max-queue-size	dhcp-ddns	max-queue-size
dhcp-ddns.ncr-format	dhcp-ddns	ncr-format
dhcp-ddns.ncr-protocol	dhcp-ddns	ncr-protocol
dhcp-ddns.sender-ip	dhcp-ddns	sender-ip
dhcp-ddns.sender-port	dhcp-ddns	sender-port
dhcp-ddns.server-ip	dhcp-ddns	server-ip
dhcp-ddns.server-port	dhcp-ddns	server-port
expired-leases-processing.flush-reclaimed-timer-wait-time	expired-leases-processing	flush-reclaimed-timer-wait-time
expired-leases-processing.hold-reclaimed-time	expired-leases-processing	hold-reclaimed-time
expired-leases-processing.max-reclaim-leases	expired-leases-processing	max-reclaim-leases
expired-leases-processing.max-reclaim-time	expired-leases-processing	max-reclaim-time
expired-leases-processing.reclaim-timer-wait-time	expired-leases-processing	reclaim-timer-wait-time
expired-leases-processing.unwarned-reclaim-cycles	expired-leases-processing	unwarned-reclaim-cycles
multi-threading.enable-multi-threading	multi-threading	enable-multi-threading
multi-threading.thread-pool-size	multi-threading	thread-pool-size
multi-threading.packet-queue-size	multi-threading	packet-queue-size
sanity-checks.lease-checks	sanity-checks	lease-checks
sanity-checks.extended-info-checks	sanity-checks	extended-info-checks
server-id.type	server-id	type
server-id.enterprise-id	server-id	enterprise-id
server-id.identifier	server-id	identifier
server-id.persist	server-id	persist
dhcp-queue-control.enable-queue	dhcp-queue-control	enable-queue
dhcp-queue-control.queue-type	dhcp-queue-control	queue-type
dhcp-queue-control.capacity	dhcp-queue-control	capacity

10.20.2 Enabling the Configuration Backend

Consider the following configuration snippet, which uses a MySQL configuration database:

```
{
  "Dhcp6": {
    "server-tag": "my DHCPv6 server",
    "config-control": {
      "config-databases": [
        {
          "type": "mysql",
          "name": "kea",
          "user": "kea",
          "password": "1234",
```

(continues on next page)

(continued from previous page)

```

        "host": "2001:db8:1::1",
        "port": 3302
    },
    "config-fetch-wait-time": 20
},
"hooks-libraries": [
    {
        "library": "/usr/local/lib/kea/hooks/libdhcp_mysql.so"
    },
    {
        "library": "/usr/local/lib/kea/hooks/libdhcp_cb_cmds.so"
    }
]
}
}

```

The following snippet illustrates the use of a PostgreSQL database:

```

{
  "Dhcp6": {
    "server-tag": "my DHCPv6 server",
    "config-control": {
      "config-databases": [
        {
          "type": "postgresql",
          "name": "kea",
          "user": "kea",
          "password": "1234",
          "host": "2001:db8:1::1",
          "port": 3302
        }
      ],
      "config-fetch-wait-time": 20
    },
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_pgsql.so"
      },
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_cb_cmds.so"
      }
    ]
  }
}

```

The configuration structure is almost identical to that of the DHCPv4 server (see *Enabling the Configuration Backend* for the detailed description).

10.21 Kea DHCPv6 Compatibility Configuration Parameters

ISC's intention is for Kea to follow the RFC documents to promote better standards compliance. However, many buggy DHCP implementations already exist that cannot be easily fixed or upgraded. Therefore, Kea provides an easy-to-use compatibility mode for broken or non-compliant clients. For that purpose, the compatibility option must be enabled to permit uncommon practices:

```
{
  "Dhcp6": {
    "compatibility": {
    }
  }
}
```

10.21.1 Lenient Option Parsing

By default, DHCPv6 option 16's vendor-class-data field is parsed as a set of length-value pairs; the same is true for tuple fields defined in custom options.

With "lenient-option-parsing": true, if a length ever exceeds the rest of the option's buffer, previous versions of Kea returned a log message unable to parse the opaque data tuple, the buffer length is x, but the tuple length is y with x < y; this no longer occurs. Instead, the value is considered to be the rest of the buffer, or in terms of the log message above, the tuple length y becomes x.

Enabling this flag is expected to improve compatibility with devices such as RAD MiNID.

```
{
  "Dhcp6": {
    "compatibility": {
      "lenient-option-parsing": true
    }
  }
}
```

Starting with Kea version 2.5.8, this parsing is extended to silently ignore client-fqdn (39) options with some invalid domain names.

10.22 Allocation Strategies in DHCPv6

A DHCP server follows a complicated algorithm to select a DHCPv6 lease for a client. It prefers assigning specific addresses or delegated prefixes requested by the client and the ones for which the client has reservations.

When the client requests a specific delegated prefix, *kea-dhcp6* follows a series of steps to try to satisfy the request, in this order:

1. It searches for a lease that matches the requested prefix and prefix length.
2. It searches for a lease that matches the prefix length.
3. It searches for a lease with a larger address space (smaller prefix length).
4. It searches for a lease with a smaller address space (larger prefix length).

If the client requests no particular lease and has no reservations, or other clients are already using any requested leases, the server must find another available lease within the configured pools. A server function called an "allocator" is responsible in Kea for finding an available lease in such a case.

The Kea DHCPv6 server provides configuration parameters to select different allocators at the global, shared-network, and subnet levels. It also allows different allocation strategies to be selected for address assignments and prefix delegation.

Consider the following example:

```
{
  "Dhcp6": {
    "allocator": "iterative",
    "pd-allocator": "random",
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8:1::/64",
        "allocator": "random"
      },
      {
        "id": 2,
        "subnet": "2001:db8:2::/64",
        "pd-allocator": "iterative"
      }
    ]
  }
}
```

The iterative allocator is globally selected for address assignments, while the random allocator is globally selected for prefix delegation. These settings are selectively overridden at the subnet level.

The following sections describe the supported allocators and their recommended uses.

10.22.1 Allocators Comparison

In the table below, we briefly compare the supported allocators, all of which are described in detail in later sections.

Table 8: Comparison of the lease allocators supported by Kea DHCPv6

Allocator	Low Performance	Utilization	High Performance	Utilization	Lease Ran- domization	Startup/Configurat	Memory Usage
Iterative	very high		low		no	very fast	low
Random	high		low		yes	very fast	high (vary- ing)
Free Lease Queue	high		high		yes	slow (depends on pool sizes)	high (vary- ing)

10.22.2 Iterative Allocator

This is the default allocator used by the Kea DHCPv6 server. It remembers the last offered lease and offers the following sequential lease to the next client. For example, it may offer addresses in this order: 2001:db8:1::10, 2001:db8:1::11, 2001:db8:1::12, and so on. Similarly, it offers the next sequential delegated prefix after the previous one to the next client. The time to find and offer the next lease or delegated prefix is very short; thus, this is the most performant allocator when pool utilization is low and there is a high probability that the next selected lease is available.

The iterative allocation underperforms when multiple DHCP servers share a lease database or are connected to a cluster. The servers tend to offer and allocate the same blocks of addresses to different clients independently, which causes many

allocation conflicts between the servers and retransmissions by clients. A random allocation addresses this issue by dispersing the allocation order.

10.22.3 Random Allocator

The random allocator uses a uniform randomization function to select offered addresses and delegated prefixes from subnet pools. It is suitable in deployments where multiple servers are connected to a shared database or a database cluster. By dispersing the offered leases, the servers minimize the risk of allocating the same lease to two different clients at the same or nearly the same time. In addition, it improves the server's resilience against attacks based on allocation predictability.

The random allocator is, however, slightly slower than the iterative allocator. Moreover, it increases the server's memory consumption because it must remember randomized leases to avoid offering them repeatedly. Memory consumption grows with the number of offered leases; in other words, larger pools and more clients increase memory consumption by random allocation.

10.22.4 Free Lease Queue Allocator (Prefix Delegation Only)

This is a sophisticated allocator whose use should be considered in subnets with highly utilized delegated prefix pools. In such cases, it can take a considerable amount of time for the iterative or random allocator to find an available prefix, because they must repeatedly check whether there is a valid lease for a prefix they will offer. The number of checks can be as high as the number of delegated prefixes in the subnet when the subnet pools are exhausted, which can have a direct negative impact on the DHCP response time for each request.

The Free Lease Queue (FLQ) allocator tracks lease allocations and de-allocations and maintains a running list of available delegated prefixes for each pool. It allows an available lease to be selected within a constant time, regardless of the subnet pools' utilization. The allocator continuously updates the list of free leases by removing any allocated leases and adding released or reclaimed ones.

The following configuration snippet shows how to select the FLQ allocator for prefix delegation in a subnet:

```
{
  "Dhcp6": {
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8:1::/64",
        "pd-allocator": "flq"
      }
    ]
  }
}
```

Note

The Free Lease Queue allocator can only be used for DHCPv6 prefix delegation. An attempt to use this allocator for address assignment (with the `allocator` parameter) will cause a configuration error. DHCPv6 address pools are typically very large and their utilization is low; in this situation, the benefits of using the FLQ allocator diminish. The amount of time required for the allocator to populate the free lease queue would cause the server to freeze upon startup.

There are several considerations that the administrator should take into account before using this allocator for prefix delegation. The FLQ allocator can heavily impact the server's startup and reconfiguration time, because the allocator has to populate the list of free leases for each subnet where it is used. These delays can be observed both during the

configuration reload and when the subnets are created using `libdhcp_subnet_cmds.so`. This allocator increases memory consumption to hold the list of free leases, proportional to the total size of the pools for which this allocator is used. Finally, lease reclamation must be enabled with a low value of the `reclaim-timer-wait-time` parameter, to ensure that the server frequently collects expired leases and makes them available for allocation via the free lease queue. Expired leases are not considered free by the allocator until they are reclaimed by the server. See [Lease Reclamation](#) for more details about the lease reclamation process.

We recommend that the FLQ allocator be selected only after careful consideration. The server puts no restrictions on the delegated-prefix pool sizes used with the FLQ allocator, so we advise users to test how long it takes for the server to load the pools before deploying the configuration using the FLQ allocator in production. We also recommend specifying another allocator type in the global configuration settings and overriding this selection at the subnet or shared-network level, to use the FLQ allocator only for selected subnets. That way, when a new subnet is added without an allocator specification, the global setting is used, thus avoiding unnecessary impact on the server's startup time.

 Warning

The FLQ allocator is not suitable for use with a shared lease database (i.e., when multiple Kea servers store leases in the same database). The servers are unaware of the expired leases reclaimed by the sibling servers and never return them to their local free lease queues. As a result, the servers will not be able to offer some of the available leases to the clients. Only a server reclaiming a particular lease will be able to offer it.

DATABASE CONNECTIVITY

The Kea servers (*kea-dhcp4* and *kea-dhcp6*) can be configured to use a variety of database backends for leases, hosts, and configuration. They can be configured to support automatic recovery when connectivity is lost, via the `on-fail` and `retry-on-startup` parameters. (The `reconnect-wait-time` and `max-reconnect-tries` parameters are described in *Lease Database Configuration* and *Lease Database Configuration*.)

It is important to understand how and when automatic recovery comes into play. Automatic recovery, when configured, only operates after a successful startup or reconfiguration during which connectivity to all backends has been successfully established.

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server will start reconnection attempts even at server startup or on reconfigure events, and will honor the action specified in the `on-fail` parameter.

During dynamic reconfiguration, all backends are disconnected and then reconnected using the new configuration. If connectivity to any of the backends cannot be established, the server logs a fatal error but remains up. It is able to process commands but does not serve clients. This allows the configuration to be corrected via the `config-set` or `remote-*` commands, if required.

During normal operations, if connectivity to any of the backends is lost and automatic recovery for that backend is enabled, the server disconnects from the respective backend and then attempts to reconnect. During the recovery process, the server ceases to serve clients according to the `on-fail` configured option but continues to respond to commands.

The `on-fail` parameter configures the actions the server should take when a connection is lost. It can have one of the following values:

- `stop-retry-exit` - indicates that the server should stop the service while it tries to recover the connection, and exit if recovery is not successful after `max-reconnect-tries`.
- `serve-retry-exit` - indicates that the server should not stop the service while it tries to recover the connection, and exit if recovery is not successful after `max-reconnect-tries`.
- `serve-retry-continue` - indicates that the server should not stop the service while it tries to recover the connection, and not exit if recovery is not successful after `max-reconnect-tries`.

If connectivity to all backends is restored, the server returns to normal operations. If the connection cannot be restored and the server is configured to exit, it issues a fatal error before shutdown.

For Kea DHCP servers to work with database backends, the database schema must be created and must match the version of the Kea server. If the version check fails and the database backend is not configured as read-only, Kea attempts to initialize the schema.

Note

Kea does not automatically attempt to upgrade an existing schema; doing so might deprive users of the opportunity to take precautions, such as backing up the database or temporarily shutting off running Kea servers that are currently operating on the database.

The connection to the database server can optionally be protected by TLS. Corresponding database configuration parameters for Kea servers are:

- The `trust-anchor` specifies the Certification Authority file name or directory path.
- The `cert-file` specifies the client certificate file name.
- The `key-file` specifies the private key file name.
- The `cipher-list` specifies the list of TLS ciphers (the syntax of the content of this parameter is described in the OpenSSL ciphers manual).

These parameters are similar to the parameters of the secure connections with the agent but are interpreted by different backends using database configurations too.

Currently the support for each database is:

- MySQL supports the whole set, additional configuration must be done in the MySQL local setup, for instance certificate revocation list, choice of a specific TLS version, mutual authentication, etc. When a TLS connection was required but the actual connection is in clear text an error log is emitted.
- PostgreSQL only uses the configuration to enable the SSL/TLS support in the client library (libpq). Anything else must be done in the PostgreSQL local configuration.

LEASE EXPIRATION

The primary role of the DHCP server is to assign addresses and/or delegate prefixes to DHCP clients. These addresses and prefixes are often referred to as "leases." Leases are typically assigned to clients for a finite amount of time, known as the "valid lifetime." DHCP clients who wish to continue using their assigned leases periodically renew them by sending the appropriate message to the DHCP server. The DHCP server records the time when these leases are renewed and calculates new expiration times for them.

If the client does not renew a lease before its valid lifetime elapses, the lease is considered expired. There are many situations when the client may cease lease renewals; common scenarios include when the machine running the client shuts down for an extended period of time, or when a mobile device leaves the vicinity of a network.

The process through which the DHCP server makes expired leases available for reassignment is referred to as "lease reclamation," and expired leases returned to availability through this process are referred to as "reclaimed." The DHCP server attempts to reclaim an expired lease as soon as it detects that it has expired. The server has several possible ways to detect expiration: it may attempt to allocate a lease to a client but find this lease already present in the database and expired; or it can periodically query the lease database for expired leases. Regardless of how an expired lease is detected, it must be reclaimed before it can be assigned to a client.

This chapter explains how to configure the server to periodically query for the expired leases, and how to minimize the impact of the periodic lease-reclamation process on the server's responsiveness. Finally, it explains "lease affinity," which provides the means to assign the same lease to a returning client after its lease has expired.

Although all configuration examples in this section are provided for the DHCPv4 server, the same parameters may be used for DHCPv6 server configuration.

12.1 Lease Reclamation

Lease reclamation is the process through which an expired lease becomes available for assignment to the same or a different client. This process involves the following steps for each reclaimed lease:

- Invoke callouts for the `lease4_expire` or `lease6_expire` hook points, if hook libraries supporting those callouts are currently loaded.
- Update the DNS, i.e. remove any DNS entries associated with the expired lease.
- Update lease information in the lease database to indicate that the lease is now available for reassignment.
- Update counters on the server, a process that includes increasing the number of reclaimed leases and decreasing the number of assigned addresses or delegated prefixes.

Please refer to *The DHCP-DDNS Server* to see how to configure DNS updates in Kea, and to *Hook Libraries* for information about using hook libraries.

12.2 Lease Reclamation Configuration Parameters

The following list presents all the configuration parameters pertaining to processing expired leases, with their default values:

- `reclaim-timer-wait-time` - this parameter governs intervals between the completion of the previous reclamation cycle and the start of the next one. Specified in seconds; the default value is 10.
- `flush-reclaimed-timer-wait-time` - this parameter controls how often the server initiates the lease reclamation procedure. Expressed in seconds; the default value is 25. If both `flush-reclaimed-timer-wait-time` and `hold-reclaimed-time` are not 0, when the client sends a release message the lease is expired instead of being deleted from lease storage.
- `hold-reclaimed-time` - this parameter governs how long the lease should be kept after it is reclaimed. This enables lease affinity when set to a non-zero value. Expressed in seconds; the default value is 3600. If both `flush-reclaimed-timer-wait-time` and `hold-reclaimed-time` are not 0, when the client sends a release message the lease is expired instead of being deleted from lease storage.
- `max-reclaim-leases` - this parameter specifies the maximum number of reclaimed leases that can be processed at one time. Zero means unlimited (i.e. process all reclaimed leases). The default value is 100.
- `max-reclaim-time` - this parameter specifies an upper limit to the length of time a lease reclamation procedure can take. Zero means no time limit. Expressed in milliseconds; the default value is 250.
- `unwarned-reclaim-cycles` - if lease reclamation limits are specified (`max-reclaim-leases` and/or `max-reclaim-time`), then under certain circumstances the server may not be able to deal with the leases to be reclaimed fast enough. This parameter specifies how many consecutive clean-up cycles must end with remaining leases to be processed before a warning is printed. The default is 5 cycles.

The parameters are explained in more detail in the rest of this chapter.

The default value for any parameter is used when the parameter is not explicitly specified in the configuration. If the `expired-leases-processing` map is omitted entirely in the configuration, the default values are used for all parameters listed above.

12.3 Configuring Lease Reclamation

Kea can be configured to periodically detect and reclaim expired leases. During this process the lease entries in the database are modified or removed. While this is happening the server does not process incoming DHCP messages, to avoid issues with concurrent access to database information. As a result, the server is unresponsive while lease reclamation is performed and DHCP queries will accumulate; responses will be sent once the lease-reclamation cycle is complete.

In deployments where response time is critical, administrators may wish to minimize the interruptions in service caused by lease reclamation. To this end, Kea provides configuration parameters to control the frequency of lease reclamation cycles, the maximum number of leases processed in a single reclamation cycle, and the maximum amount of time a single reclamation cycle is allowed to run before being interrupted. The following examples demonstrate how these parameters can be used:

```
{
  "Dhcp4": {
    "expired-leases-processing": {
      "reclaim-timer-wait-time": 5,
      "max-reclaim-leases": 0,
      "max-reclaim-time": 0
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

The first parameter is expressed in seconds and specifies an interval between the two consecutive lease reclamation cycles. This is explained by the following diagram:



This diagram shows four lease-reclamation cycles (c1 through c4) of variable duration. The duration of the reclamation cycle depends on the number of expired leases detected and processed in a particular cycle. This duration is usually significantly shorter than the interval between the cycles.


According to the `reclaim-timer-wait-time`, the server keeps fixed intervals of five seconds between the end of one cycle and the start of the next cycle. This guarantees the presence of 5-second-long periods during which the server remains responsive to DHCP queries and does not perform lease reclamation. The `max-reclaim-leases` and `max-reclaim-time` are set to 0, which sets no restriction on the maximum number of leases reclaimed in the particular cycle, or on the maximum duration of each cycle.

In deployments with high lease-pool utilization, relatively short valid lifetimes, and frequently disconnecting clients which allow leases to expire, the number of expired leases requiring reclamation at any given time may rise significantly. In this case, it is often desirable to apply restrictions to the maximum duration of a reclamation cycle or the maximum number of leases reclaimed in a cycle. The following configuration demonstrates how this can be done:

```
{
  "Dhcp4": {
    "expired-leases-processing": {
      "reclaim-timer-wait-time": 3,
      "max-reclaim-leases": 100,
      "max-reclaim-time": 50,
      "unwarned-reclaim-cycles": 10
    }
  }
}
```

In this example, the `max-reclaim-leases` parameter limits the number of leases reclaimed in a single cycle to 100, and the `max-reclaim-time` limits the maximum duration of each cycle to 50ms. The lease-reclamation cycle will be interrupted if either of these limitations is reached. The reclamation of any unreclaimed leases will be attempted in subsequent cycles.

The following diagram illustrates the behavior of the system in the presence of many expired leases, when the limits are applied for the reclamation cycles:



In this case, if any reclamation cycle takes more than 50ms, it is interrupted according to the value of the `max-reclaim-time`. This results in equal durations of all reclamation cycles over time. In this example, the limitation of the maximum 100 leases is not reached. This may be the case when database transactions or callouts in the hook libraries attached to the server are slow. Regardless, the chosen values for either the maximum number of leases

or a maximum cycle time strongly depend on the particular deployment, the lease database backend being used, any hook libraries, etc. Administrators may need to experiment to tune the system to suit the dynamics of their deployment.

It is important to realize that with the use of these limits, there is a risk that expired leases will accumulate faster than the server can reclaim them. This should not be a problem if the server is dealing with a temporary burst of expirations, because it should be able to eventually deal with them over time. However, if leases expire at a high rate for a long period of time, the unreclaimed leases will pile up in the database. To notify the administrator that the current configuration does not satisfy the needs for reclamation of expired leases, the server issues a warning message in the log if it is unable to reclaim all leases within several reclamation cycles. The number of cycles after which such a warning is issued is specified with the `unwarned-reclaim-cycles` configuration parameter.

Setting the `reclaim-timer-wait-time` to 0 disables periodic reclamation of the expired leases.

12.4 Configuring Lease Affinity

Suppose that a laptop goes into sleep mode after a period of user inactivity. While the laptop is in sleep mode, its DHCP client does not renew leases obtained from the server and these leases will eventually expire. When the laptop wakes up, it is often desirable for it to continue using its previous assigned IP addresses. To facilitate this, the server needs to correlate returning clients with their expired leases. When the client returns, the server first checks for those leases and reassigns them if they have not been assigned to another client. The ability of the server to reassign the same lease to a returning client is referred to as "lease affinity."

When lease affinity is enabled (i.e. when `hold-reclaimed-time` is configured to a value greater than zero), the server still reclaims leases according to the parameters described in [Configuring Lease Reclamation](#), but the reclaimed leases are held in the database for a specified amount of time rather than removed. If both `flush-reclaimed-timer-wait-time` and `hold-reclaimed-time` are greater than zero, the lease is expired immediately when the client sends a release message, instead of being deleted from lease storage. When the client returns, the server first verifies whether there are any reclaimed leases associated with this client and then reassigns them if possible. However, it is important to note that any reclaimed lease may be assigned to another client if that client specifically asks for it. Therefore, lease affinity does not guarantee that the reclaimed lease will be available for the client who used it before; it merely increases the chances of the client being assigned the same lease. If the lease pool is small - namely, in DHCPv4, for which address space is limited - there is an increased likelihood that the expired lease will be assigned to another client.

Consider the following configuration:

```
"Dhcp4": {
  "expired-leases-processing": {
    "reclaim-timer-wait-time": 3,
    "hold-reclaimed-time": 1800,
    "flush-reclaimed-timer-wait-time": 5
  },
  ...
}
```

The `hold-reclaimed-time` specifies how many seconds after an expiration a reclaimed lease should be held in the database for reassignment to the same client. In the example given above, reclaimed leases are held for 30 minutes (1800 seconds) after their expiration. During this time, the server will likely be able to reassign the same lease to the returning client, unless another client specifically requests this lease and the server assigns it.

The server must periodically remove reclaimed leases for which the time indicated by `hold-reclaimed-time` has elapsed. The `flush-reclaimed-timer-wait-time` parameter controls how often the server removes such leases. In the example provided above, the server initiates removal of such leases five seconds after the previous removal attempt was completed. Setting this value to 0 disables lease affinity, meaning leases are removed from the lease database when they are reclaimed. If lease affinity is enabled, it is recommended that the `hold-reclaimed-time`

be set to a value significantly higher than the `reclaim-timer-wait-time`, as timely removal of expired-reclaimed leases is less critical than the removal process, which may impact server responsiveness.

There is no guarantee that lease affinity will work every time; if a server is running out of addresses, it will reassign expired addresses to new clients. Also, clients can request specific addresses and the server tries to honor such requests if possible. Administrators who want to ensure a client keeps its address, even after periods of inactivity, should consider using host reservations or leases with very long lifetimes.

12.5 Reclaiming Expired Leases via Command

The `leases-reclaim` command can be used to trigger lease reclamation at any time. Please consult the *The leases-reclaim Command* section for details about using this command.

CONGESTION HANDLING

13.1 What is Congestion?

Congestion occurs when servers are subjected to client queries faster than those queries can be processed. As a result, the servers begin accumulating a backlog of pending queries. The longer the high rate of traffic continues, the farther behind the servers fall. Depending on the client implementations, those that fail to get leases either give up or simply continue to retry forever. In the former case, the server may eventually recover, but the latter case is a vicious cycle from which the server is unable to escape.

Congestion typically occurs when there is a network event that causes overly large numbers of clients to simultaneously need leases, such as recovery after a network outage. In a well-planned deployment, the number and capacity of servers is matched to the maximum expected client load. If the load is routinely too heavy, then the deployment needs to be re-evaluated.

The goal of congestion handling is to help servers mitigate the peak in traffic by fulfilling as many of the most relevant requests as possible until the congestion subsides.

13.2 Configuring Congestion Handling

Congestion handling offers the ability to configure the server to use a separate thread to read packets from the interface socket buffers. As the thread reads packets from the buffers, they are added to an internal packet queue, and the server's main application thread processes packets from this queue rather than from the socket buffers. By structuring it this way, a configurable layer has been introduced which can make decisions on which packets to process, how to store them, and the order in which they are processed by the server.

The default packet queue implementation for both *kea-dhcp4* and *kea-dhcp6* is a simple ring buffer. Once it reaches capacity, new packets get added to the back of the queue by discarding packets from the front of the queue. Rather than always discarding the newest packets, Kea now always discards the oldest packets. The capacity of the buffer, i.e. the maximum number of packets the buffer can contain, is configurable. A reasonable starting point is to match the capacity to the number of leases per second a specific installation of Kea can handle. This figure varies widely depending on the specifics of an individual deployment.

As there is no one algorithm that can best handle the dynamics of all sites, and because over time new approaches will evolve, the packet queue is implemented as a plug-in, which can be replaced by a custom queue implementation via a hook library. This should make it straightforward for interested parties to experiment with their own solutions. (Developers can refer to `isc::dhcp::PacketQueue` and `isc::dhcp::PacketQueueMgr`, described in the [Kea Developer's Guide](#).)

Packet queue behavior is configured in both *kea-dhcp4* and *kea-dhcp6* servers through an optional, top-level, configuration element, `dhcp-queue-control`. Omitting this element disables packet queueing:

```
"dhcp-queue-control": {  
  "enable-queue": true, // true|false
```

(continues on next page)

(continued from previous page)

```

"queue-type": "queue type",
"capacity" : 256 // n packets
}

```

where:

- **enable-queue** - enables or disables packet queueing. When **true**, the server processes packets from the packet queue, which is filled by a separate thread. When **false**, the server processes packets directly from the socket buffers in the main thread. It is disabled (**false**) by default.
- **queue-type** - the name of the queue implementation to use. This value exists so that custom implementations can be registered (via a hook library) and then selected. There is a default packet queue implementation that is pre-registered during server start up: "kea-ring4" for *kea-dhcp4* and "kea-ring6" for *kea-dhcp6*.
- **capacity** - this is the maximum number of packets the queue can hold before packets are discarded. The optimal value for this is extremely site-dependent. The default value is 64 for both "kea-ring4" and "kea-ring6".

The following example enables the default packet queue for *kea-dhcp4*, with a queue capacity of 250 packets:

```

"Dhcp4":
{
  "dhcp-queue-control": {
    "enable-queue": true,
    "queue-type": "kea-ring4",
    "capacity" : 250
  },
  ...
}

```

The following example enables the default packet queue for *kea-dhcp6*, with a queue capacity of 300 packets:

```

"Dhcp6":
{
  "dhcp-queue-control": {
    "enable-queue": true,
    "queue-type": "kea-ring6",
    "capacity" : 300
  },
  ...
}

```

Note

Congestion handling is currently incompatible with multi-threading; when both are enabled, congestion handling is silently disabled.

THE DHCP-DDNS SERVER

14.1 Overview

The DHCP-DDNS Server (*kea-dhcp-ddns*, known informally as D2) conducts the client side of the Dynamic DNS protocol (DDNS, defined in [RFC 2136](#)) on behalf of the DHCPv4 and DHCPv6 servers (*kea-dhcp4* and *kea-dhcp6*, respectively). The DHCP servers construct DDNS update requests, known as NameChangeRequests (NCRs), based on DHCP lease change events and then post them to D2. D2 attempts to match each request to the appropriate DNS server(s) and carries out the necessary conversation with those servers to update the DNS data.

The Kea hook library *libdhcp_ddns_tuning.so* provides the ability for both *kea-dhcp4* and *kea-dhcp6* to generate host names procedurally based on an expression, to skip DDNS updates on a per-client basis, or to fine-tune various DNS update aspects. Please refer to the *libdhcp_ddns_tuning.so: DDNS Tuning* documentation for the configuration options.

14.1.1 DNS Server Selection

To match a request to the appropriate DNS servers, D2 must have a catalog of servers from which to select. In fact, D2 has two such catalogs, one for forward DNS and one for reverse DNS; these catalogs are referred to as "DDNS domain lists." Each list consists of one or more named DDNS domains. Further, each DDNS domain has a list of one or more DNS servers that publish the DNS data for that domain.

When conducting forward-domain matching, D2 compares the fully qualified domain name (FQDN) in the request against the name of each forward DDNS domain in its catalog. The domain whose name matches the longest portion of the FQDN is considered the best match. For example, if the FQDN is "myhost.sample.example.com.", and there are two forward domains in the catalog, "sample.example.com." and "example.com.", the former is regarded as the best match. In some cases, it may not be possible to find a suitable match. Given the same two forward domains there would be no match for the FQDN "bogus.net", so the request would be rejected. Finally, if there are no forward DDNS domains defined, D2 simply disregards the forward-update portion of requests.

When conducting reverse-domain matching, D2 constructs a reverse FQDN from the lease address in the request and compares that against the name of each reverse DDNS domain. Again, the domain whose name matches the longest portion of the FQDN is considered the best match. For instance, if the lease address is "172.16.1.40" and there are two reverse domains in the catalog, "1.16.172.in-addr.arpa." and "16.172.in-addr.arpa", the former is the best match. As with forward matching, D2 may not find a suitable match. Given the same two domains, there would be no match for the lease address, "192.168.1.50", and the request would be rejected. As with forward-domain matching, if there are no reverse DDNS domains defined, D2 simply disregards the reverse-update portion of requests.

14.1.2 Conflict Resolution

D2 implements the conflict resolution strategy prescribed by [RFC 4703](#). Conflict resolution is intended to prevent different clients from mapping to the same FQDN at the same time. To make this possible, the RFC requires that forward DNS entries for a given FQDN must be accompanied by a DHCID resource record (RR). This record contains a client identifier that uniquely identifies the client to whom the name belongs. Furthermore, any DNS updaters that

wishes to update or remove existing forward entries for an FQDN may only do so if their client matches that of the DHCID RR.

In other words, the DHCID RR maps an FQDN to the client to whom it belongs, and thereafter changes to that mapping can only be done by or at the behest of that client.

Conflict resolution can be indirectly enabled or disabled via the configuration parameter `ddns-use-conflict-resolution`, supported by both *kea-dhcp4* and *kea-dhcp6*. These servers use this parameter to set a flag within each NameChangeRequest they send that tells D2 whether conflict resolution should be employed for that request. By default, conflict resolution is enabled. For more details, please refer to discussions of `ddns-use-conflict-resolution` in *DDNS for DHCPv4* and *DDNS for DHCPv6*.

When conflict resolution is disabled, D2 still adds DHCID RRs but does not use them to enforce client ownership of DNS entries. Disabling it should only be used after careful consideration.

14.1.3 Dual-Stack Environments

RFC 4703, section 5.2, describes issues that may arise with dual-stack clients. These are clients that wish to have both IPv4 and IPv6 mappings for the same FQDN. To work properly, clients must embed their IPv6 DUID within their IPv4 client identifier option, as described in RFC 4361. In this way, DNS updates for both IPv4 and IPv6 can be managed under the same DHCID RR. This feature is supported by Kea beginning with release 2.1.2.

14.2 Starting and Stopping the DHCP-DDNS Server

kea-dhcp-ddns is the Kea DHCP-DDNS server and, due to the nature of DDNS, it runs alongside either the DHCPv4 or DHCPv6 component (or both). Like other parts of Kea, it is a separate binary that can be run on its own or through *keactrl* (see *Managing Kea with keactrl*). In normal operation, controlling *kea-dhcp-ddns* with *keactrl* is recommended; however, it is also possible to run the DHCP-DDNS server directly. It accepts the following command-line switches:

- `-c file` - specifies the configuration file. This is the only mandatory switch.
- `-d` - specifies whether server logging should be switched to debug/verbose mode. In verbose mode, the logging severity and debuglevel specified in the configuration file are ignored and "debug" severity and the maximum debuglevel (99) are assumed. The flag is convenient for temporarily switching the server into maximum verbosity, e.g. when debugging.
- `-t file` - specifies the configuration file to be tested. *kea-dhcp-ddns* attempts to load it and conducts sanity checks. Certain checks are possible only while running the actual server. The actual status is reported with an exit code (0 = configuration looks okay, 1 = error encountered). Kea prints out log messages to standard output and errors to standard error when testing the configuration.
- `-v` - displays the Kea version and exits.
- `-V` - displays the extended Kea version and exits.
- `-W` - displays the Kea configuration report and exits. The report is a copy of the `config.report` file produced by `meson setup`; it is embedded in the executable binary.

The contents of the `config.report` file may also be accessed by examining certain libraries in the installation tree or in the source tree.

```
# from installation using libkea-process.so
$ strings ${prefix}/lib/libkea-process.so | sed -n 's/;;; //p'

# from sources using libkea-process.so
$ strings src/lib/process/.libs/libkea-process.so | sed -n 's/;;; //p'
```

(continues on next page)

(continued from previous page)

```
# from sources using libkea-process.a
$ strings src/lib/process/.libs/libkea-process.a | sed -n 's/;;; //p'

# from sources using libcfgrpt.a
$ strings src/lib/process/cfgrpt/.libs/libcfgrpt.a | sed -n 's/;;; //p'
```

- **-X** - As of Kea 3.0, disables path and permissions restrictions. The server will emit a warning at startup that security restrictions have been disabled. Do not use this mode of operation without careful consideration and taking any necessary precautions. Failure to do so can expose deployments to security vulnerabilities. For more information please read section *Securing a Kea Deployment*.

Upon startup, the module loads its configuration and begins listening for NCRs based on that configuration.

During startup, the server attempts to create a PID file of the form: `[pidfile_dir]/[conf name].kea-dhcp-ddns.pid` where:

- **pidfile_dir** - is `[prefix]/[localstatedir]/run/kea` where `prefix` and `localstatedir` are the values passed into meson setup using `--prefix` and `--localstatedir` which default to `/usr/local` and `var` respectively. So the whole `pidfile_dir` defaults to `/usr/local/var`. Note that this value may be overridden at runtime by setting the environment variable `KEA_PIDFILE_DIR` intended primarily for testing purposes.
- **conf name** - is the configuration file name used to start the server, minus all preceding paths and the file extension. For example, given a pathname of `/usr/local/etc/kea/myconf.txt`, the portion used would be `myconf`.

If the file already exists and contains the PID of a live process, the server issues a `DHCP_DDNS_ALREADY_RUNNING` log message and exits. It is possible, though unlikely, that the file is a remnant of a system crash and the process to which the PID belongs is unrelated to Kea. In such a case it is necessary to manually delete the PID file.

14.3 Configuring the DHCP-DDNS Server

Before starting the `kea-dhcp-ddns` module for the first time, a configuration file must be created. The following default configuration is a template that can be customized to individual requirements.

```
"DhcpDdns": {
  "ip-address": "127.0.0.1",
  "port": 53001,
  "dns-server-timeout": 500,
  "ncr-protocol": "UDP",
  "ncr-format": "JSON",
  "tsig-keys": [ ],
  "forward-ddns": {
    "ddns-domains": [ ]
  },
  "reverse-ddns": {
    "ddns-domains": [ ]
  }
}
```

The configuration can be divided into the following sections, each of which is described below:

- *Global Server Parameters* - define values which control connectivity and global server behavior.
- *Control Sockets* - defines the Control Socket list.
- *TSIG Key Info* - defines the TSIG keys used for secure traffic with DNS servers.

- *Forward DDNS* - defines the catalog of forward DDNS domains.
- *Reverse DDNS* - defines the catalog of reverse DDNS domains.

14.3.1 Global Server Parameters

- `ip-address` - the IP address on which D2 listens for requests. The default is the local loopback interface at address 127.0.0.1. Either an IPv4 or IPv6 address may be specified.
- `port` - the port on which D2 listens for requests. The default value is 53001.
- `dns-server-timeout` - the maximum amount of time, in milliseconds, that D2 will wait for a response from a DNS server to a single DNS update message. The default is 500 ms.
- `ncr-protocol` - the socket protocol to use when sending requests to D2. Currently only UDP is supported.
- `ncr-format` - the packet format to use when sending requests to D2. Currently only JSON format is supported.

D2 must listen for change requests on a known address and port. By default it listens at 127.0.0.1 on port 53001. The following example illustrates how to change D2's global parameters so it will listen at 192.168.1.10 port 900:

```
"DhcpDdns": {
  "ip-address": "192.168.1.10",
  "port": 900,
  ...
}
```

Warning

It is possible for a malicious attacker to send bogus NameChangeRequests to the DHCP-DDNS server. Addresses other than the IPv4 or IPv6 loopback addresses (127.0.0.1 or ::1) should only be used for testing purposes; note that local users may still communicate with the DHCP-DDNS server.

Note

If the `ip-address` and `port` are changed, the corresponding values in the DHCP servers' `dhcp-ddns` configuration section must be changed.

14.3.2 Management API for the D2 Server

The management API allows the issuing of specific management commands, such as configuration retrieval or shut-down. For more details, see [Management API](#). By default there are no sockets open; to instruct Kea to open a socket, the following entry in the configuration file can be used:

```
"DhcpDdns": {
  "control-sockets": [
    {
      "socket-type": "unix",
      "socket-name": "/path/to/the/unix/socket"
    }
  ],
  ...
}
```

14.3.2.1 UNIX Control Socket

Until Kea server 2.7.2 the only supported communication channel type was the UNIX stream socket with `socket-type` set to `unix` and `socket-name` to the file path of the UNIX/LOCAL socket.

The length of the path specified by the `socket-name` parameter is restricted by the maximum length for the UNIX socket name on the operating system, i.e. the size of the `sun_path` field in the `sockaddr_un` structure, decreased by 1. This value varies on different operating systems, between 91 and 107 characters. Typical values are 107 on Linux and 103 on FreeBSD.

Kea supports only one `unix` control socket in the "control-sockets" list.

Note

As of Kea 2.7.9, control sockets may only reside in the directory determined during compilation as "[`kea-install-dir`]/var/run/kea", which must also have `0750` access rights. This path may be overridden at startup by setting the environment variable `KEA_CONTROL_SOCKET_DIR` to the desired path. If a path other than this value is used in `socket-name`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use simply omit the path component from `socket-name`.

Communication over the control channel is conducted using JSON structures. See the [Control Channel](#) section in the *Kea Developer's Guide* for more details.

The D2 server supports the following operational commands:

- `build-report`
- `config-get`
- `config-hash-get`
- `config-reload`
- `config-set`
- `config-test`
- `config-write`
- `list-commands`
- `shutdown`
- `status-get`
- `version-get`

Since Kea version 2.0.0, the D2 server also supports the following operational commands for statistics:

- `statistic-get`
- `statistic-get-all`
- `statistic-reset`
- `statistic-reset-all`

The `shutdown` command supports the extra `type` argument, which controls the way the D2 server cleans up on exit. The supported shutdown types are:

- `normal` - stops the queue manager and finishes all current transactions before exiting. This is the default.
- `drain_first` - stops the queue manager but continues processing requests from the queue until it is empty.
- `now` - exits immediately.

An example command may look like this:

```
{
  "command": "shutdown",
  "arguments": {
    "exit-value": 3,
    "type": "drain_first"
  }
}
```

14.3.2.2 HTTP/HTTPS Control Socket

The `socket-type` must be `http` or `https` (when the type is `https` TLS is required). The `socket-address` (default `127.0.0.1`) and `socket-port` (default `8000`) specify an IP address and port to which the HTTP service will be bound.

The `trust-anchor`, `cert-file`, `key-file`, and `cert-required` parameters specify the TLS setup for HTTP, i.e. HTTPS. If these parameters are not specified, HTTP is used. The TLS/HTTPS support in Kea is described in [TLS/HTTPS Support](#).

Basic HTTP authentication protects against unauthorized uses of the control agent by local users. For protection against remote attackers, HTTPS and reverse proxy of [Secure Connections](#) provide stronger security.

The authentication is described in the `authentication` block with the mandatory `type` parameter, which selects the authentication. Currently only the basic HTTP authentication (type `basic`) is supported.

The `realm` authentication parameter (default `kea-dhcp-ddns-server` is used for error messages when the basic HTTP authentication is required but the client is not authorized.

When the `clients` authentication list is configured and not empty, basic HTTP authentication is required. Each element of the list specifies a user ID and a password. The user ID is mandatory, must not be empty, and must not contain the colon (:) character. The password is optional; when it is not specified an empty password is used.

Note

The basic HTTP authentication user ID and password are encoded in UTF-8, but the current Kea JSON syntax only supports the Latin-1 (i.e. 0x00..0xff) Unicode subset.

To avoid exposing the user ID and/or the associated password, these values can be read from files. The syntax is extended by:

- The `directory` authentication parameter, which handles the common part of file paths. The default value is the empty string.
- The `password-file` client parameter, which, alongside the `directory` parameter, specifies the path of a file that can contain the password, or when no user ID is given, the whole basic HTTP authentication secret.
- The `user-file` client parameter, which, with the `directory` parameter, specifies the path of a file where the user ID can be read.

Since Kea-2.7.6 Kea supports multiple HTTP/HTTPS connections. Both IPv4 and IPv6 addresses can be used. The server will issue an error when changing the socket type from HTTP to HTTPS or from HTTPS to HTTP using the same address and port. This action is not allowed as it might introduce a security issue accidentally caused by a user mistake. A different address or port must be specified when using the "config-set" command to switch from HTTP to HTTPS or from HTTPS to HTTP. The same applies when modifying the configuration file and then running "config-reload" command.

When files are used, they are read when the configuration is loaded, to detect configuration errors as soon as possible.

```

"DhcpDdns": {
  "control-sockets": [
    {
      "socket-type": "https",
      "socket-address": "10.20.30.40",
      "socket-port": 8005,
      "trust-anchor": "/path/to/the/ca-cert.pem",
      "cert-file": "/path/to/the/agent-cert.pem",
      "key-file": "/path/to/the/agent-key.pem",
      "cert-required": true,
      "authentication": {
        "type": "basic",
        "realm": "kea-dhcp-ddns-server",
        "clients": [
          {
            "user": "admin",
            "password": "1234"
          }
        ]
      }
    },
    {
      "socket-type": "http",
      "socket-address": "2010:30:40::50",
      "socket-port": 8004
    }
  ],
  ...
}

```

14.3.3 TSIG Key List

A DDNS protocol exchange can be conducted with or without a transaction signature, or TSIG (defined in [RFC 2845](#)). This configuration section allows the administrator to define the set of TSIG keys that may be used in such exchanges.

To use TSIG when updating entries in a DNS domain, a key must be defined in the TSIG key list and referenced by name in that domain's configuration entry. When D2 matches a change request to a domain, it checks whether the domain has a TSIG key associated with it. If so, D2 uses that key to sign DNS update messages sent to and verify responses received from the domain's DNS server(s). For each TSIG key required by the DNS servers that D2 is working with, there must be a corresponding TSIG key in the TSIG key list.

As one might gather from the name, the `tsig-key` section of the D2 configuration lists the TSIG keys. Each entry describes a TSIG key used by one or more DNS servers to authenticate requests and sign responses. Every entry in the list has three parameters:

- **name** - is a unique text label used to identify this key within the list. This value is used to specify which key (if any) should be used when updating a specific domain. As long as the name is unique its content is arbitrary, although for clarity and ease of maintenance it is recommended that it match the name used on the DNS server(s). This field cannot be blank.
- **algorithm** - specifies which hashing algorithm should be used with this key. This value must specify the same algorithm used for the key on the DNS server(s). The supported algorithms are listed below:
 - HMAC-MD5
 - HMAC-SHA1
 - HMAC-SHA224

- HMAC-SHA256
- HMAC-SHA384
- HMAC-SHA512

This value is not case-sensitive.

- **digest-bits** - is used to specify the minimum truncated length in bits. The default value 0 means truncation is forbidden; non-zero values must be an integral number of octets, and be greater than both 80 and half of the full length. (Note that in BIND 9 this parameter is appended to the algorithm name, after a dash.)
- **secret** - is used to specify the shared secret key code for this key. This value is case-sensitive and must exactly match the value specified on the DNS server(s). It is a base64-encoded text value.
- **secret-file** - since Kea 2.5.8, this more secure alternative is supported. This value specifies the file name where the secret can be found, i.e. the base64-encoded secret is the content of the file.

As an example, suppose that a domain D2 will be updating is maintained by a BIND 9 DNS server, which requires dynamic updates to be secured with TSIG. Suppose further that the entry for the TSIG key in BIND 9's `named.conf` file looks like this:

```
:
key "key.four.example.com." {
    algorithm hmac-sha224;
    secret "bZEG70w80gAUPfLWV3aAUQ==";
};
:
```

By default, the TSIG key list is empty:

```
"DhcpDdns": {
    "tsig-keys": [ ],
    ...
}
```

A new key must be added to the list:

```
"DhcpDdns": {
    "tsig-keys": [
        {
            "name": "key.four.example.com.",
            "algorithm": "HMAC-SHA224",
            "secret": "bZEG70w80gAUPfLWV3aAUQ=="
        }
    ],
    ...
}
```

These steps must be repeated for each TSIG key needed, although the same TSIG key can be used with more than one domain.

14.3.4 Forward DDNS

The forward DDNS section is used to configure D2's forward-update behavior. Currently it contains a single parameter, the catalog of forward DDNS domains, which is a list of structures.

```
"DhcpDdns": {
  "forward-ddns": {
    "ddns-domains": [ ]
  },
  ...
}
```

By default, this list is empty, which causes the server to ignore the forward-update portions of requests.

14.3.4.1 Adding Forward DDNS Domains

A forward DDNS domain maps a forward DNS zone to a set of DNS servers which maintain the forward DNS data (i.e. name-to-address mapping) for that zone. Each zone served needs one forward DDNS domain. Some or all of the zones may be maintained by the same servers, but one DDNS domain is still needed for each zone. Remember that matching a request to the appropriate server(s) is done by zone and a DDNS domain only defines a single zone.

This section describes how to add forward DDNS domains; repeat these steps for each forward DDNS domain desired. Each forward DDNS domain has the following parameters:

- **name** - this is the fully qualified domain name (or zone) that this DDNS domain can update. This value is compared against the request FQDN during forward matching. It must be unique within the catalog.
- **key-name** - if TSIG is used with this domain's servers, this value should be the name of the key from the TSIG key list. If the value is blank (the default), TSIG will not be used in DDNS conversations with this domain's servers.
- **dns-servers** - this is a list of one or more DNS servers which can conduct the server side of the DDNS protocol for this domain. The servers are used in a first-to-last preference; in other words, when D2 begins to process a request for this domain, it will pick the first server in this list and attempt to communicate with it. If that attempt fails, D2 will move to the next one in the list and so on, until either it is successful or the list is exhausted.

To create a new forward DDNS domain, add a new domain element and set its parameters:

```
"DhcpDdns": {
  "forward-ddns": {
    "ddns-domains": [
      {
        "name": "other.example.com.",
        "key-name": "",
        "dns-servers": [
          ]
      }
    ]
  }
}
```

It is possible to add a domain without any servers; however, if that domain matches a request, the request will fail. To make the domain useful, at least one DNS server must be added to it.

14.3.4.1.1 Adding Forward DNS Servers

This section describes how to add DNS servers to a forward DDNS domain. Repeat these instructions as needed for all the servers in each domain.

Forward DNS server entries represent actual DNS servers which support the server side of the DDNS protocol. Each forward DNS server has the following parameters:

- `hostname` - the resolvable host name of the DNS server; this parameter is not yet implemented.
- `ip-address` - the IP address at which the server listens for DDNS requests. This may be either an IPv4 or an IPv6 address.
- `port` - the port on which the server listens for DDNS requests. It defaults to the standard DNS service port of 53.

To create a new forward DNS server, a new server element must be added to the domain and its parameters filled in. If, for example, the service is running at "172.88.99.10", set the forward DNS server as follows:

```
"DhcpDdns": {
  "forward-ddns": {
    "ddns-domains": [
      {
        "name": "other.example.com.",
        "key-name": "",
        "dns-servers": [
          {
            "ip-address": "172.88.99.10",
            "port": 53
          }
        ]
      }
    ]
  }
}
```

Note

Since `hostname` is not yet supported, the parameter `ip-address` must be set to the address of the DNS server.

14.3.5 Reverse DDNS

The reverse DDNS section is used to configure D2's reverse update behavior, and the concepts are the same as for the forward DDNS section. Currently it contains a single parameter, the catalog of reverse DDNS domains, which is a list of structures.

```
"DhcpDdns": {
  "reverse-ddns": {
    "ddns-domains": [ ]
  },
  ...
}
```

By default, this list is empty, which causes the server to ignore the reverse-update portions of requests.

14.3.5.1 Adding Reverse DDNS Domains

A reverse DDNS domain maps a reverse DNS zone to a set of DNS servers which maintain the reverse DNS data (address-to-name mapping) for that zone. Each zone served needs one reverse DDNS domain. Some or all of the zones may be maintained by the same servers, but one DDNS domain entry is needed for each zone. Remember that matching a request to the appropriate server(s) is done by zone and a DDNS domain only defines a single zone.

This section describes how to add reverse DDNS domains; repeat these steps for each reverse DDNS domain desired. Each reverse DDNS domain has the following parameters:

- **name** - this is the fully qualified reverse zone that this DDNS domain can update. This is the value used during reverse matching, which compares it with a reversed version of the request's lease address. The zone name should follow the appropriate standards; for example, to support the IPv4 subnet 172.16.1, the name should be "1.16.172.in-addr.arpa.". Similarly, to support an IPv6 subnet of 2001:db8:1, the name should be "1.0.0.8.B.D.0.1.0.0.2.ip6.arpa." The name must be unique within the catalog.
- **key-name** - if TSIG is used with this domain's servers, this value should be the name of the key from the TSIG key list. If the value is blank (the default), TSIG will not be used in DDNS conversations with this domain's servers.
- **dns-servers** - this is a list of one or more DNS servers which can conduct the server side of the DDNS protocol for this domain. Currently, the servers are used in a first-to-last preference; in other words, when D2 begins to process a request for this domain, it will pick the first server in this list and attempt to communicate with it. If that attempt fails, D2 will move to the next one in the list and so on, until either it is successful or the list is exhausted.

To create a new reverse DDNS domain, a new domain element must be added and its parameters set. For example, to support subnet 2001:db8:1::, the following configuration could be used:

```
"DhcpDdns": {
  "reverse-ddns": {
    "ddns-domains": [
      {
        "name": "1.0.0.8.B.D.0.1.0.0.2.ip6.arpa.",
        "key-name": "",
        "dns-servers": [
        ]
      }
    ]
  }
}
```

It is possible to add a domain without any servers; however, if that domain matches a request, the request will fail. To make the domain useful, at least one DNS server must be added to it.

14.3.5.1.1 Adding Reverse DNS Servers

This section describes how to add DNS servers to a reverse DDNS domain. Repeat these instructions as needed for all the servers in each domain.

Reverse DNS server entries represent actual DNS servers which support the server side of the DDNS protocol. Each reverse DNS server has the following parameters:

- **hostname** - the resolvable host name of the DNS server; this value is currently ignored.
- **ip-address** - the IP address at which the server listens for DDNS requests.
- **port** - the port on which the server listens for DDNS requests. It defaults to the standard DNS service port of 53.

To create a new reverse DNS server, a new server element must be added to the domain and its parameters specified. If, for example, the service is running at "172.88.99.10", then set it as follows:

```
"DhcpDdns": {
  "reverse-ddns": {
```

(continues on next page)

(continued from previous page)

```

    "ddns-domains": [
      {
        "name": "1.0.0.0.8.B.D.0.1.0.0.2.ip6.arpa.",
        "key-name": "",
        "dns-servers": [
          {
            "ip-address": "172.88.99.10",
            "port": 53
          }
        ]
      }
    ]
  }
}

```

Note

Since `hostname` is not yet supported, the parameter `ip-address` must be set to the address of the DNS server.

14.3.5.2 Per-DNS-Server TSIG Keys

Since Kea version 2.0.0, a TSIG key can be specified in a DNS server configuration. The priority rule is:

- if a not-empty key name is specified in a DNS server entry, this TSIG key protects DNS updates sent to this server.
- if the DNS server entry is empty, but a not-empty key name is specified in the parent's domain entry, the parent domain's TSIG key protects DNS updates sent to this server.
- if the DNS server entry is empty, and no key name is specified in its parent domain entry, no TSIG protects DNS updates sent to this server.

For instance, in this configuration:

```

"DhcpDdns": {
  "forward-ddns": {
    "ddns-domains": [
      {
        "name": "other.example.com.",
        "key-name": "foo",
        "dns-servers": [
          {
            "ip-address": "172.88.99.10",
            "port": 53
          },
          {
            "ip-address": "172.88.99.11",
            "port": 53,
            "key-name": "bar"
          }
        ]
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    },
    "reverse-ddns": {
      "ddns-domains": [
        {
          "name": "1.0.0.0.8.B.D.0.1.0.0.2.ip6.arpa.",
          "dns-servers": [
            {
              "ip-address": "172.88.99.12",
              "port": 53
            },
            {
              "ip-address": "172.88.99.13",
              "port": 53,
              "key-name": "bar"
            }
          ]
        }
      ]
    },
    "tsig-keys": [
      {
        "name": "foo",
        "algorithm": "HMAC-MD5",
        "secret": "LSWXnfkKZjdPJI5QxlpnfQ=="
      },
      {
        "name": "bar",
        "algorithm": "HMAC-SHA224",
        "secret": "bZEG70w80gAUPfLWV3aAUQ=="
      }
    ]
  }
}

```

The 172.88.99.10 server will use the "foo" TSIG key, the 172.88.99.11 and 172.88.99.13 servers will use the "bar" key, and 172.88.99.12 will not use TSIG.

14.3.6 User Contexts in DDNS

See *Comments and User Context* for additional background regarding the user context idea.

User contexts can be specified on a global scope, a DDNS domain, a DNS server, a TSIG key, and loggers. One other useful usage is the ability to store comments or descriptions; the parser translates a "comment" entry into a user context with the entry, which allows a comment to be attached inside the configuration itself.

14.3.7 Example DHCP-DDNS Server Configuration

This section provides a sample DHCP-DDNS server configuration, based on a small example network. Let's suppose our example network has three domains, each with their own subnet.

Table 1: Our example network

Domain	Subnet	Forward DNS Servers	Reverse DNS Servers
four.example.com	192.0.2.0/24	172.16.1.5, 172.16.2.5	172.16.1.5, 172.16.2.5
six.example.com	2001:db8:1::/64	3001:1::50	3001:1::51
example.com	192.0.0.0/16	172.16.2.5	172.16.2.5

We need to construct three forward DDNS domains:

Table 2: Forward DDNS domains needed

#	DDNS Domain Name	DNS Servers
1.	four.example.com.	172.16.1.5, 172.16.2.5
2.	six.example.com.	3001:1::50
3.	example.com.	172.16.2.5

As discussed earlier, FQDN-to-domain matching is based on the longest match. The FQDN "my-host.four.example.com." matches the first domain ("four.example.com."), while "admin.example.com." matches the third domain ("example.com."). The FQDN "other.example.net." fails to match any domain and is rejected.

The following example configuration specifies the forward DDNS domains.

```
"DhcpDdns": {
  "comment": "example configuration: forward part",
  "forward-ddns": {
    "ddns-domains": [
      {
        "name": "four.example.com.",
        "key-name": "",
        "dns-servers": [
          { "ip-address": "172.16.1.5" },
          { "ip-address": "172.16.2.5" }
        ]
      },
      {
        "name": "six.example.com.",
        "key-name": "",
        "dns-servers": [
          { "ip-address": "2001:db8:1" }
        ]
      },
      {
        "name": "example.com.",
        "key-name": "",
```

(continues on next page)

(continued from previous page)

```

        "dns-servers": [
            { "ip-address": "172.16.2.5" }
        ],
        "user-context": { "backup": false }
    },
    ...
]
}
}

```

Similarly, we need to construct the three reverse DDNS domains:

Table 3: Reverse DDNS domains needed

#	DDNS Domain Name	DNS Servers
1.	2.0.192.in-addr.arpa.	172.16.1.5, 172.16.2.5
2.	1.0.0.0.8.d.b.0.1.0.0.2.ip6.arpa.	3001:1::50
3.	0.182.in-addr.arpa.	172.16.2.5

An address of "192.0.2.150" matches the first domain, "2001:db8:1::10" matches the second domain, and "192.0.50.77" matches the third domain.

These reverse DDNS domains are specified as follows:

```

"DhcpDdns": {
    "comment": "example configuration: reverse part",
    "reverse-ddns": {
        "ddns-domains": [
            {
                "name": "2.0.192.in-addr.arpa.",
                "key-name": "",
                "dns-servers": [
                    { "ip-address": "172.16.1.5" },
                    { "ip-address": "172.16.2.5" }
                ]
            },
            {
                "name": "1.0.0.0.8.B.D.0.1.0.0.2.ip6.arpa.",
                "key-name": "",
                "dns-servers": [
                    { "ip-address": "2001:db8::1" }
                ]
            },
            {
                "name": "0.192.in-addr.arpa.",
                "key-name": "",
                "dns-servers": [

```

(continues on next page)

(continued from previous page)

```

        { "ip-address": "172.16.2.5" }
      ],
      ...
    ]
  }
}

```

14.4 DHCP-DDNS Server Statistics

Kea version 2.0.0 introduced statistics support for DHCP-DDNS.

Statistics are divided into three groups: NameChangeRequests, DNS updates, and per-TSIG-key DNS updates. While the statistics of the first two groups are cumulative, i.e. not affected by configuration change or reload, per-key statistics are reset to 0 when the underlying object is (re)created.

Currently Kea's statistics management has the following limitations:

- only integer samples (i.e. a counter and a timestamp) are used;
- the maximum sample count is 1;
- there is no API to remove one or all statistics;
- there is no API to set the maximum sample count or age.

Note

Hook libraries, such as the ISC subscriber-only GSS-TSIG library, make new statistics available in Kea.

More information about Kea statistics can be found at [Statistics](#).

14.4.1 NCR Statistics

The NameChangeRequest statistics are:

- `ncr-received` - the number of received valid NCRs
- `ncr-invalid` - the number of received invalid NCRs
- `ncr-error` - the number of errors in NCR receptions other than an I/O cancel on shutdown
- `queue-mgr-queue-full` - the number of times the NCR receive queue reached maximum capacity

14.4.2 DNS Update Statistics

The global DNS update statistics are:

- `update-sent` - the number of DNS updates sent
- `update-signed` - the number of DNS updates sent and protected by TSIG
- `update-unsigned` - the number of DNS updates sent and not protected by TSIG
- `update-success` - the number of DNS updates which successfully completed
- `update-timeout` - the number of DNS updates which completed on timeout
- `update-error` - the number of DNS updates which completed with an error other than timeout

14.4.3 Per-TSIG-Key DNS Update Statistics

The per TSIG key DNS update statistics are:

- `update-sent` - the number of DNS updates sent
- `update-success` - the number of DNS updates which successfully completed
- `update-timeout` - the number of DNS updates which completed on timeout
- `update-error` - the number of DNS updates which completed with an error other than timeout

The name format for per-key statistics is `key[<key-DNS-name>].<stat-name>`: for instance, the name of the `update-sent` statistics for the `key.example.com`. TSIG key is `key[key.example.com].update-sent`.

14.5 DHCP-DDNS Server Limitations

The following are the current limitations of the DHCP-DDNS server.

- Requests received from the DHCP servers are placed in a queue until they are processed. Currently, all queued requests are lost if the server shuts down.

14.6 Supported Standards

The following RFCs are supported by the DHCP-DDNS server:

- *Secret Key Transaction Authentication for DNS (TSIG)*, [RFC 2845](#): All DNS update packets sent and received by the DHCP-DDNS server can be protected by TSIG signatures.
- *Dynamic Updates in the Domain Name System (DNS UPDATE)*, [RFC 2136](#): The complete DNS update mechanism is supported.
- *Resolution of Fully Qualified Domain Name (FQDN) Conflicts among Dynamic Host Configuration Protocol (DHCP) Clients*, [RFC 4703](#): DHCP-DDNS takes care of conflict resolution, for both DHCPv4 and DHCPv6 servers.
- *A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)*, [RFC 4701](#): The DHCP-DDNS server uses DHCID records.

THE LFC PROCESS

15.1 Overview

kea-lfc is a service process that removes redundant information from the files used to provide persistent storage for the memfile database backend. This service is written to run as a standalone process.

While *kea-lfc* can be started externally, there is usually no need to do so. *kea-lfc* is run on a periodic basis by the Kea DHCP servers.

The process operates on a set of files, using them to receive input and output of the lease entries and to indicate what stage the process is in, in the event of an interruption. Currently the caller must supply names for all of the files.

15.2 Command-Line Options

kea-lfc is run as follows:

```
kea-lfc [-4 | -6] -c config-file -p pid-file -x previous-file -i copy-file -o output-  
↪file -f finish-file
```

The argument `-4` or `-6` selects the protocol version of the lease files.

The `-c` argument specifies the configuration file. This is required, but is not currently used by the process.

The `-p` argument specifies the PID file. When the *kea-lfc* process starts, it attempts to determine whether another instance of the process is already running by examining the PID file. If one is already running, the new process is terminated; if one is not running, Kea writes its PID into the PID file.

The other filenames specify where the *kea-lfc* process should look for input, write its output, and perform its bookkeeping:

- **previous** — when *kea-lfc* starts, this is the result of any previous run of *kea-lfc*. When *kea-lfc* finishes, it is the result of this run. If *kea-lfc* is interrupted before completing, this file may not exist.
- **input** — before the DHCP server invokes *kea-lfc*, it moves the current lease file here and then calls *kea-lfc* with this file.
- **output** — this is the temporary file where *kea-lfc* writes the leases. Once the file has finished writing, it is moved to the **finish** file (see below).
- **finish** — this is another temporary file *kea-lfc* uses for bookkeeping. When *kea-lfc* completes writing the output file, it moves the contents to the file of this name. After *kea-lfc* finishes deleting the other files (**previous** and **input**), it moves this file to the **previous** lease file. By moving the files in this fashion, *kea-lfc* and the DHCP server processes can determine the correct file to use even if one of the processes is interrupted before completing its task.

There are several additional arguments, mostly for debugging purposes. `-d` sets the logging level to debug. `-v` and `-V` print out version stamps, with `-V` providing a longer form. `-h` prints out the usage string.

CLIENT CLASSIFICATION

16.1 Client Classification Overview

In certain cases it is useful to differentiate among different types of clients and treat them accordingly. Common reasons include:

- The clients represent different pieces of topology, e.g. a cable modem is not the same as the clients behind that modem.
- The clients have different behavior, e.g. a smartphone behaves differently from a laptop.
- The clients require different values for some options, e.g. a docsis3.0 cable modem requires different settings from a docsis2.0 cable modem.

To make management easier, different clients can be grouped into a client class to receive common options.

An incoming packet can be associated with a client class in several ways:

- Implicitly, using a vendor class option or another built-in condition.
- Using an expression which evaluates to `true`.
- Using static host reservations, a shared network, a subnet, etc.
- Using a hook.

Client classification can be used to change the behavior of almost any part of the DHCP message processing. There are currently nine mechanisms that take advantage of client classification:

- Dropping queries.
- Subnet selection.
- Pool selection.
- Lease limiting.
- Rate limiting.
- DDNS tuning.
- Defining DHCPv4 private (codes 224-254) and code 43 options.
- Assigning different options.
- Setting specific options for use with the TFTP server address and the boot file field for DHCPv4 cable modems.

16.1.1 Classification Steps

The classification process is conducted in several steps:

1. The ALL class is associated with the incoming packet.
2. Vendor class options are processed.
3. Classes with matching expressions and not marked for later evaluation ("on request" or depending on the KNOWN/UNKNOWN built-in classes) are processed in the order they are defined in the configuration; the boolean expression is evaluated and, if it returns `true` (a match), the incoming packet is associated with the class.
4. If a private or code 43 DHCPv4 option is received, it is decoded following its client-class or global (or, for option 43, last-resort) definition.
5. When the incoming packet belongs to the special class DROP, it is dropped and an informational message is logged with the packet information.

Note

The `pkt4_receive` and `pkt6_receive` callouts are called here.

6. When the `early-global-reservations-lookup` global parameter is configured to `true`, the process looks up global reservations and partially performs steps 8, 9, and 10. The lookup is limited to global reservations; if one is found the KNOWN class is set, but if none is found the UNKNOWN class is **not** set.
7. A subnet is chosen, possibly based on the class information when some subnets are reserved. More precisely: when choosing a subnet, the server iterates over all of the subnets that are feasible given the information found in the packet (client address, relay address, etc.). It uses the first subnet it finds that either has no class associated with it, or has a class which matches one of the packet's classes.

Note

The `subnet4_select` and `subnet6_select` callouts are called here.

8. The server looks for host reservations. If an identifier from the incoming packet matches a host reservation in the subnet or shared network, the packet is associated with the KNOWN class and all classes of the host reservation. If a reservation is not found, the packet is assigned to the UNKNOWN class.
9. Classes with matching expressions - directly, or indirectly using the KNOWN/UNKNOWN built-in classes and not marked for later evaluation ("on request") - are processed in the order they are defined in the configuration; the boolean expression is evaluated and, if it returns `true` (a match), the incoming packet is associated with the class. After a subnet is selected, the server determines whether there is a reservation for a given client. Therefore, it is not possible to use the UNKNOWN class to select a shared network or a subnet. For the KNOWN class, only global reservations are used and the `early-global-reservations-lookup` parameter must be configured to `true`.
10. When the incoming packet belongs to the special class DROP, it is dropped and an informational message is logged with the packet information. Since Kea version 1.9.8, it is permissible to make the DROP class dependent on the KNOWN/UNKNOWN classes.
11. If needed, addresses and prefixes from pools are assigned, possibly based on the class information when some pools are reserved for class members.

Note

The `lease4_select`, `lease4_renew`, `lease6_select`, `lease6_renew`, and `lease6_rebind` callouts are called here.

12. The `evaluate-additional-classes` lists (if any) are evaluated first for the pool, then the subnet, and the finally shared-network to which the assigned resources belong. Classes are evaluated in the order they appear within each list.
13. Options are assigned, again possibly based on the class information in the order that classes were associated with the incoming packet. The first class matched to the packet which specifies a given option wins. For DHCPv4 private and code 43 options, this includes option definitions specified within classes. The order of precedence for options is as follows:
 1. Host options
 2. Pool options
 3. Subnet options
 4. Shared network options
 5. Options from classes assigned during normal classification:
 - a. ALL class
 - b. Vendor classes
 - c. Host classes
 - d. Evaluated classes that do not depend on KNOWN class
 - e. KNOWN and UNKNOWN classes
 - f. Evaluated classes that depend on KNOWN class
 6. Options from classes in `evaluate-additional-classes` from:
 - a. Pool
 - b. Subnet
 - c. Shared network
 7. Global options

Note

Care should be taken with client classification, as it is easy for clients that do not meet any class criteria to be denied service altogether.

16.2 Built-in Client Classes

Some classes are built-in, so they do not need to be explicitly defined. They can be defined if there is a need to associate lease lifetimes, option data, etc. with them.

Vendor class information is the primary example: the server checks whether an incoming DHCPv4 packet includes the vendor class identifier option (60) or an incoming DHCPv6 packet includes the vendor class option (16). If it does, the content of that option is prepended with `VENDOR_CLASS_` and the result is interpreted as a class for that packet. The content that is considered is the whole class identifier for DHCPv4, and the first vendor class data field for DHCPv6. The enterprise number and subsequent vendor class data fields are not used for the purpose of classification. For example, modern cable modems send such options with value `docsis3.0`, so the packet belongs to class `VENDOR_CLASS_docsis3.0`.

The `HA_` prefix is used by `libdhcp_ha.so` to designate certain servers to process DHCP packets as a result of load balancing. The class name is constructed by prepending the `HA_` prefix to the name of the server which should process the DHCP packet. This server uses an appropriate pool or subnet to allocate IP addresses (and/or prefixes), based on

the assigned client classes. The details can be found in *libdhcp_ha.so: High Availability Outage Resilience for Kea Servers*.

The SPAWN_ prefix is used by template classes to generate spawned class names at runtime. The spawned class name is constructed by prepending the SPAWN_ prefix to the template class name and the evaluated value: SPAWN_<template-class-name>_<evaluated-value>. More details can be found in *Configuring Classes*.

The BOOTP class is used by *libdhcp_bootp.so* to classify and respond to inbound BOOTP queries.

The SKIP_DDNS class is used by the DDNS-tuning hook library to suppress DDNS updates on a per client basis.

Other examples are the ALL class, to which all incoming packets belong, and the KNOWN class, assigned when host reservations exist for a particular client. By convention, the names of built-in classes begin with all capital letters.

Currently recognized built-in class names are ALL, KNOWN and UNKNOWN, and the prefixes VENDOR_CLASS_, HA_, AFTER_, EXTERNAL_, SKIP_DDNS. Although the AFTER_ prefix is a provision for an as-yet-unwritten hook, the EXTERNAL_ prefix can be freely used; built-in classes are implicitly defined so they never raise warnings if they do not appear in the configuration.

16.3 Using Expressions in Classification

The expression portion of a classification definition contains operators and values. All values are currently strings; operators take a string or strings and return another string. When all the operations have completed, the result should be a value of `true` or `false`. The packet belongs to the class (and the class name is added to the list of classes) if the result is `true`. Expressions are written in standard format and can be nested.

Expressions are pre-processed during the parsing of the configuration file and converted to an internal representation. This allows certain types of errors to be caught and logged during parsing. Examples of these errors include an incorrect number or type of argument to an operator. The evaluation code also checks for this class of error and generally throws an exception, though this should not occur in a normally functioning system.

Other issues, such as the starting position of a substring being outside of the substring or an option not existing in the packet, result in the operator returning an empty string.

Dependencies between classes are also checked. For instance, forward dependencies are rejected when the configuration is parsed; an expression can only depend on already-defined classes (including built-in classes) which are evaluated in a previous or the same evaluation phase. This does not apply to the KNOWN or UNKNOWN classes.

Table 1: List of classification values

Name	Example expression	Example value
String literal	'example'	'example'
Hexadecimal string literal	0x5a7d	'Z}'
IP address literal	10.0.0.1	0x0a000001
Integer literal	123	'123'
Binary content of the option	option[123].hex	'(content of the option)'
Option existence	option[123].exists	'true'
Binary content of the sub-option	option[12].option[34].hex	'(content of the sub-option)'
Sub-Option existence	option[12].option[34].exists	'true'
Client class membership	member('foobar')	'true'
Known client	known	member('KNOWN')
Unknown client	unknown	not member('KNOWN')
DHCPv4 relay agent sub-option	relay4[123].hex	'(content of the RAI sub-option)'
DHCPv6 Relay Options	relay6[nest].option[code].hex	(value of the option)
DHCPv6 Relay Peer Address	relay6[nest].peeraddr	2001:DB8::1
DHCPv6 Relay Link Address	relay6[nest].linkaddr	2001:DB8::1

continues on next page

Table 1 – continued from previous page

Name	Example expression	Example value
Interface name of packet	<code>pkt.iface</code>	<code>eth0</code>
Source address of packet	<code>pkt.src</code>	<code>10.1.2.3</code>
Destination address of packet	<code>pkt.dst</code>	<code>10.1.2.3</code>
Length of packet	<code>pkt.len</code>	<code>513</code>
Hardware address in DHCPv4 packet	<code>pkt4.mac</code>	<code>0x010203040506</code>
Hardware length in DHCPv4 packet	<code>pkt4.hlen</code>	<code>6</code>
Hardware type in DHCPv4 packet	<code>pkt4.htype</code>	<code>6</code>
ciaddr field in DHCPv4 packet	<code>pkt4.ciaddr</code>	<code>192.0.2.1</code>
giaddr field in DHCPv4 packet	<code>pkt4.giaddr</code>	<code>192.0.2.1</code>
yiaddr field in DHCPv4 packet	<code>pkt4.yiaddr</code>	<code>192.0.2.1</code>
siaddr field in DHCPv4 packet	<code>pkt4.siaddr</code>	<code>192.0.2.1</code>
Message type in DHCPv4 packet	<code>pkt4.msgtype</code>	<code>1</code>
Transaction ID (xid) in DHCPv4 packet	<code>pkt4.transid</code>	<code>12345</code>
Message type in DHCPv6 packet	<code>pkt6.msgtype</code>	<code>1</code>
Transaction ID in DHCPv6 packet	<code>pkt6.transid</code>	<code>12345</code>
Vendor option existence (any vendor)	<code>vendor[*].exists</code>	<code>'true'</code>
Vendor option existence (specific vendor)	<code>vendor[4491].exists</code>	<code>'true'</code>
Enterprise-id from vendor option	<code>vendor.enterprise</code>	<code>4491</code>
Vendor sub-option existence	<code>vendor[4491].option[1].exists</code>	<code>'true'</code>
Vendor sub-option content	<code>vendor[4491].option[1].hex</code>	<code>docsis3.0</code>
Vendor class option existence (any vendor)	<code>vendor-class[*].exists</code>	<code>'true'</code>
Vendor class option existence (specific vendor)	<code>vendor-class[4491].exists</code>	<code>'true'</code>
Enterprise-id from vendor class option	<code>vendor-class.enterprise</code>	<code>4491</code>
First data chunk from vendor class option	<code>vendor-class[4491].data</code>	<code>docsis3.0</code>
Specific data chunk from vendor class option	<code>vendor-class[4491].data[3]</code>	<code>docsis3.0</code>

Notes:

- Hexadecimal strings are converted into a string as expected. The starting `0X` or `0x` is removed, and if the string is an odd number of characters a "0" is prepended to it.
- IP addresses are converted into strings of length 4 or 16. IPv4, IPv6, and IPv4-embedded IPv6 (e.g. IPv4-mapped IPv6) addresses are supported.
- Integers in an expression are converted to 32-bit unsigned integers and are represented as four-byte strings; for example, 123 is represented as `0x0000007b`. All expressions that return numeric values use 32-bit unsigned integers, even if the field in the packet is smaller. In general, it is easier to use decimal notation to represent integers, but it is also possible to use hexadecimal notation. When writing an integer in hexadecimal, care should be taken to make sure the value is represented as 32 bits, e.g. use `0x00000001` instead of `0x1` or `0x01`. Also, make sure the value is specified in network order, e.g. 1 is represented as `0x00000001`.
- `option[code].hex` extracts the value of the option with the code `code` from the incoming packet. If the packet does not contain the option, it returns an empty string. The string is presented as a byte string of the option payload, without the type code or length fields.
- `option[code].exists` checks whether an option with the code `code` is present in the incoming packet. It can be used with empty options.
- `member('foobar')` checks whether the packet belongs to the client class `foobar`. To avoid dependency loops, the configuration file parser verifies whether client classes were already defined or are built-in, i.e., beginning with `VENDOR_CLASS_`, `AFTER_` (for the to-come "after" hook) and `EXTERNAL_` or equal to `ALL`, `KNOWN`, `UNKNOWN`, etc.

`known` and `unknown` are shorthand for `member('KNOWN')` and `not member('KNOWN')`. Note that the evaluation

of any expression using the KNOWN class (directly or indirectly) is deferred after the host reservation lookup (i.e. when the KNOWN or UNKNOWN partition is determined).

- `relay4[code].hex` attempts to extract the value of the sub-option code from the option inserted as the DHCPv4 Relay Agent Information (82) option. If the packet does not contain a RAI option, or the RAI option does not contain the requested sub-option, the expression returns an empty string. The string is presented as a byte string of the option payload without the type code or length fields. This expression is allowed in DHCPv4 only.
- `relay4` shares the same representation types as `option`; for instance, `relay4[code].exists` is supported.
- `relay6[nest]` allows access to the encapsulations used by any DHCPv6 relays that forwarded the packet. The `nest` level specifies the relay from which to extract the information, with a value of 0 indicating the relay closest to the DHCPv6 server. Negative values allow relays to be specified counting from the DHCPv6 client, with -1 indicating the relay closest to the client. If the requested encapsulation does not exist, an empty string "" is returned. This expression is allowed in DHCPv6 only.
- `relay6[nest].option[code]` shares the same representation types as `option`; for instance, `relay6[nest].option[code].exists` is supported.
- Expressions starting with `pkt4` can be used only in DHCPv4. They allow access to DHCPv4 message fields.
- `pkt6` refers to information from the client request. To access any information from an intermediate relay, use `relay6`. `pkt6.msgtype` and `pkt6.transid` output a 4-byte binary string for the message type or transaction ID. For example, the message type SOLICIT is `0x00000001` or simply 1, as in `pkt6.msgtype == 1`.
- "Vendor option" means the Vendor-Identifying Vendor-Specific Information option in DHCPv4 (code 125; see [Section 4 of RFC 3925](#)) and the Vendor-Specific Information Option in DHCPv6 (code 17, defined in [Section 21.17 of RFC 8415](#)). "Vendor class option" means the Vendor-Identifying Vendor Class Option in DHCPv4 (code 124; see [Section 3 of RFC 3925](#)) in DHCPv4 and the Class Option in DHCPv6 (code 16; see [Section 21.16 of RFC 8415](#)). Vendor options may have sub-options that are referenced by their codes. Vendor class options do not have sub-options, but rather data chunks, which are referenced by index value. Index 0 means the first data chunk, index 1 is for the second data chunk (if present), etc.
- In the vendor and vendor-class constructs an asterisk (*) or 0 can be used to specify a wildcard `enterprise-id` value, i.e. it will match any `enterprise-id` value.
- Vendor Class Identifier (option 60 in DHCPv4) can be accessed using the `option[60]` expression.
- [RFC 3925](#) and [RFC 8415](#) allow for multiple instances of vendor options to appear in a single message. The client classification code currently examines the first instance if more than one appear. For the `vendor.enterprise` and `vendor-class.enterprise` expressions, the value from the first instance is returned. Please submit a feature request on the [Kea GitLab site](#) to request support for multiple instances.

Table 2: List of classification expressions

Name	Example	Description
Equal	'foo' == 'bar'	Compare the two values and return true or false
Match	match('foo.*', 'foobar')	Match a regular expression with a value.
Not	not ('foo' == 'bar')	Logical negation
And	('foo' == 'bar') and ('bar' == 'foo')	Logical and
Or	('foo' == 'bar') or ('bar' == 'foo')	Logical or
Substring	substring('foobar',0,3)	Return the requested substring
Concat	concat('foo','bar')	Return the concatenation of the strings
Concat (operator +)	'foo' + 'bar'	Return the concatenation of the strings
Ifelse	ifelse('foo' == 'bar','us','them')	Return the branch value according to the condition
Hexstring	hexstring('foo', '-')	Converts the value to a hexadecimal string, e.g. 66-6F-6F
Lcase	lcase('LoWeR')	Converts the value of a string expression to lower case, e.g. 'lower'
Ucase	ucase('uPpEr')	Converts the value of a string expression to upper case, e.g. 'UPPER'
Split	split('foo.bar', '.', 2)	Return the second field, splitting on dots.

Table 3: List of conversion-to-text expressions

Name	Example	Description
AddressToText	addrtotext (192.10.0.1) addrtotext (2003:db8::)	Represent the 4 bytes of an IPv4 address or the 16 bytes of an IPv6 address in human readable format
Int8ToText	int8totext (-1)	Represents the 8-bit signed integer in text format
Int16ToText	int16totext (-1)	Represents the 16-bit signed integer in text format
Int32ToText	int32totext (-1)	Represents the 32-bit signed integer in text format
UInt8ToText	uint8totext (255)	Represents the 8-bit unsigned integer in text format
UInt16ToText	uint16totext (65535)	Represents the 16-bit unsigned integer in text format
UInt32ToText	uint32totext (4294967295)	Represents the 32-bit unsigned integer in text format

Notes:

The conversion operators can be used to transform data from binary to the text representation. The only requirement is that the input data type length matches an expected value.

The AddressToText token expects 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses. The Int8ToText and UInt8ToText tokens expect 1 byte, the Int16ToText and UInt16ToText tokens expect 2 bytes, and Int32ToText and UInt32ToText expect 4 bytes. For all conversion tokens, if the data length is 0, the result string is empty.

16.3.1 Predicates

The two predicates are Equal and Match. They can be used to build other common predicates, for instance:

```
not (substring('foobar', 3, 3) == 'bar')
match('foo.*', lcase('FooBar'))
match('.*foo.*', 'is it foo or bar')
match('^.foo.*$', 'is it foo or bar')
```

So inequality, case insensitive pattern matching or pattern search.

Note

Detected invalid regular expressions are considered as syntax errors, runtime exceptions during match are handled as no match.

Warning

Be careful with the match operator as it can result in very poor performance leading to regular expression denial of service (ReDoS).

16.3.2 Logical Operators

The Not, And, and Or logical operators are the common operators. Not has the highest precedence and Or the lowest. And and Or are (left) associative. Parentheses around a logical expression can be used to enforce a specific grouping; for instance, in "A and (B or C)". Without parentheses, "A and B or C" means "(A and B) or C".

16.3.3 Substring

The substring operator `substring(value, start, length)` accepts both positive and negative values for the starting position and the length. For `start`, a value of 0 is the first byte in the string while -1 is the last byte. If the starting point is outside of the original string an empty string is returned. `length` is the number of bytes to extract. A negative number means to count towards the beginning of the string but does not include the byte pointed to by `start`. The special value `all` means to return all bytes from `start` to the end of the string. If the length is longer than the remaining portion of the string, then the entire remaining portion is returned. Some examples may be helpful:

```
substring('foobar', 0, 6) == 'foobar'
substring('foobar', 3, 3) == 'bar'
substring('foobar', 3, all) == 'bar'
substring('foobar', 1, 4) == 'ooba'
substring('foobar', -5, 4) == 'ooba'
substring('foobar', -1, -3) == 'oba'
substring('foobar', 4, -2) == 'ob'
substring('foobar', 10, 2) == ''
```

16.3.4 Concat

The concat function `concat(string1, string2)` returns the concatenation of its two arguments. For instance:

```
concat('foo', 'bar') == 'foobar'
```

For user convenience, Kea version 1.9.8 added an associative operator version of the concat function. For instance:

```
'abc' + 'def' + 'ghi' + 'jkl' + '...'
```

is the same as:

```
concat(concat(concat(concat('abc', 'def'), 'ghi'), 'jkl'), '...')
```

or:

```
concat('abc', concat('def', concat('ghi', concat('jkl', '...'))))
```

or:

```
'abcdefghijkl...'
```

16.3.5 Split

The split operator `split(value, delimiters, field-number)` accepts a list of characters to use as delimiters and a positive field number of the desired field when the value is split into fields separated by the delimiters. Adjacent delimiters are not compressed out, rather they result in an empty string for that field number. If value is an empty string, the result will be an empty string. If the delimiters list is empty, the result will be the original value. If the field-number is less than one or larger than the number of fields, the result will be an empty string. Some examples follow:

```
split ('one.two..four', '.', 1) == 'one'
split ('one.two..four', '.', 2) == 'two'
split ('one.two..four', '.', 3) == ''
split ('one.two..four', '.', 4) == 'four'
split ('one.two..four', '.', 5) == ''
```

Note

To use a hard-to-escape character as a delimiter, use its ASCII hex value. For example, split by single quote using `0x27`: `split(option[39].text, 0x27, 1)`

16.3.6 Ifelse

The ifelse function `ifelse(cond, iftrue, ifelse)` returns the iftrue or ifelse branch value following the boolean condition `cond`. For instance:

```
ifelse(option[230].exists, option[230].hex, 'none')
```

16.3.7 Hexstring

The hexstring function `hexstring(binary, separator)` returns the binary value as its hexadecimal string representation: pairs of hexadecimal digits separated by the separator, e.g. `':'`, `'-'`, `''` (empty separator).

```
hexstring(pkt4.mac, ':')
```

Note

The expression for each class is executed on each packet received. If the expressions are overly complex, the time taken to execute them may impact the performance of the server. Administrators who need complex or time-consuming expressions should consider writing a *hook* to perform the necessary work.

16.4 Configuring Classes

A client class definition can contain the following properties:

- The name parameter is mandatory and must be unique among all classes.

- The `test` expression is not mandatory and represents a string containing the logical expression used to determine membership in the class. The entire expression is included in double quotes ("). The result should evaluate to a boolean value (`true` or `false`).
- The `template-test` expression is not mandatory and represents a string containing the logical expression used to generate a spawning class. The entire expression is included in double quotes ("). The result should evaluate to a string value representing the variable part of the spawned class name. If the resulting string is empty, no spawning class is generated. The resulting spawned class has the following generated name format: `SPAWN_<template-class-name>_<evaluated-value>`. After classes are evaluated and a spawned class is generated, the corresponding template class name is also associated with the packet. An Option specified in a spawned class will take precedence over the same option if set in its template class.
- The `option-data` list is not mandatory and contains options that should be assigned to members of this class. In the case of a template class, these options are assigned to the generated spawned class.
- The `option-def` list is not mandatory and is used to define custom options.
- The `only-if-required` has been replaced with `only-in-additional-list` and is now deprecated. It will still be accepted as input for a time to allow users to migrate but will eventually be unsupported.
- The `only-in-additional-list` flag is not mandatory; when its value is set to `false` (the default), membership is determined during classification and is available for subnet selection, for instance. When the value is set to `true`, membership is evaluated only if the class appears in an `evaluate-additional-classes` list and is usable only for option configuration.
- The `user-context` is not mandatory and represents a map with user-defined data and possibly configuration options for hook libraries.
- The `next-server` parameter is not mandatory and configures the `siaddr` field in packets associated with this class. It is used in DHCPv4 only.
- The `server-hostname` is not mandatory and configures the `sname` field in packets associated with this class. It is used in DHCPv4 only.
- The `boot-file-name` is not mandatory and configures the `file` field in packets associated with this class. It is used in DHCPv4 only.
- The `valid-lifetime`, `min-valid-lifetime`, and `max-valid-lifetime` are not mandatory and configure the valid lifetime fields for this client class.
- The `preferred-lifetime`, `min-preferred-lifetime`, and `max-preferred-lifetime` are not mandatory and configure the preferred lifetime fields for this client class. It is used in DHCPv6 only.

Note

`test` and `template-test` are mutually exclusive in a client class definition. Use either one, or neither, but not both. If both are provided, the configuration is rejected.

In the following example, the class named `Client_foo` is defined. It is comprised of all clients whose client IDs (option 61) start with the string `foo`. Members of this class will be given 192.0.2.1 and 192.0.2.2 as their domain name servers.

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "Client_foo",
      "test": "substring(option[61].hex,0,3) == 'foo'",
      "option-data": [
```

(continues on next page)

(continued from previous page)

```

        {
            "name": "domain-name-servers",
            "code": 6,
            "space": "dhcp4",
            "csv-format": true,
            "data": "192.0.2.1, 192.0.2.2"
        }
    ],
    ...
],
...
}

```

The next example shows a client class named "Client_enterprise" being defined for use by the DHCPv6 server. It is comprised of all clients whose client identifiers start with the given hex string, which would indicate a DUID based on an enterprise ID of 0x0002AABBCCDD. Members of this class will be given 2001:db8:0::1 and 2001:db8:2::1 as their domain name servers.

```

"Dhcp6": {
    "client-classes": [
        {
            "name": "Client_enterprise",
            "test": "substring(option[1].hex,0,6) == 0x0002AABBCCDD",
            "option-data": [
                {
                    "name": "dns-servers",
                    "code": 23,
                    "space": "dhcp6",
                    "csv-format": true,
                    "data": "2001:db8:0::1, 2001:db8:2::1"
                }
            ]
        },
        ...
    ],
    ...
}

```

It is also possible to have both left and right operands of the evaluated expression processed at runtime. Expressions related to packets can appear in the expression as many times as needed; there is no limit. However, each token has a small impact on performance and excessively complex expressions may cause a bottleneck.

```

"Dhcp4": {
    "client-classes": [
        {
            "name": "Infrastructure",
            "test": "option[82].option[2].hex == pkt4.mac",
            ...
        },
        ...
    ],
    ...
}

```

(continues on next page)

(continued from previous page)

```

    ...
}

```

16.4.1 Template Classes

The `template-test` parameter indicates that the class is a template class.

```

"Dhcp4": {
  "client-classes": [
    {
      "name": "Client-ID",
      "template-test": "substring(option[61].hex,0,3)",
      ...
    },
    ...
  ],
  ...
}

```

If the received DHCPv4 packet contains option 61, then the first three bytes represent the value `foo` in ASCII, and the spawned class uses the `SPAWN_Client-ID_foo` name. Both the `SPAWN_Client-ID_foo` and `Client-ID` classes are associated with the packet.

Note

Template classes can also be used to spawn classes which match regular classes, effectively associating the regular class to the packet. To achieve this, the regular class must also contain the fixed part of the spawned class name:

`SPAWN_<template-class-name-used-to-activate-this-regular-class>_<evaluated-value-filtering-this-regular-class>`

```

"Dhcp6": {
  "client-classes": [
    {
      "name": "SPAWN_Client-ID_foobar",
      "test": "substring(option[1].hex,0,6) == 0x0002AABBCCDD",
      ...
    },
    {
      "name": "Client-ID",
      "template-test": "substring(option[1].hex,0,6)",
      ...
    },
    ...
  ],
  ...
}

```

If the received DHCPv6 packet contains option 1 (client identifier) with hex value `0x0002AABBCCDD`, then the `SPAWN_Client-ID_foobar` is associated with the packet. Moreover, if the first six bytes represent value `foobar` in ASCII, then the spawned class uses the `SPAWN_Client-ID_foobar` name, effectively associating the regular class to the packet. In this second case, both the `SPAWN_Client-ID_foobar` and `Client-ID` classes are associated with the

packet. The test expression on the regular class `SPAWN_Client-ID_foobar` is not mandatory and can be omitted, but it is used here with a different match expression for example purposes.

Usually the `test` and `template-test` expressions are evaluated before subnet selection, but in some cases it is useful to evaluate it later when the subnet, shared network, or pools are known but output-option processing has not yet been done. For this purpose, the `only-in-additional-list` flag, which is `false` by default, allows the evaluation of the `test` expression or the `template-test` expression only when it is required by the class's presence in the `evaluate-additional-classes` list of the selected pool, subnet, or shared network.

The `evaluate-additional-classes` list, which is valid for pool, subnet, and shared-network scope, specifies the classes which are evaluated in the second pass before output-option processing. The list is built in same precedence order as the option data, i.e. an option data item in a subnet takes precedence over one in a shared network. An additional class in a subnet is added before one in a shared network. The mechanism is related to the `only-in-additional-list` flag but it is not mandatory that the flag be set to `true`.

Note

The `template-test` expression can also be used to filter generated spawned classes, so that they are created only when needed by using the `ifelse` instruction.

```
"Dhcp4": {
  "client-classes": [
    {
      "name": "Client-ID",
      "template-test": "ifelse(substring(option[61].hex,4,3) == 'foo',
↪substring(option[12].hex,0,12), '')",
      ...
    },
    ...
  ],
  ...
}
```

Note

The template classes can be used to configure limits which, just like options, are associated with the spawned class. This permits the configuration of limits that apply to all packets associated with a class spawned at runtime, according to the `template-test` expression in the parent template class. For a more detailed description of how to configure limits using the limits hook library, see [Configuration](#). For example, using the configuration below, ingress DHCPv6 packets that have client ID values (in the format expressed by the Kea evaluator) `foobar` and `foofoo` both amount to the same limit of 60 packets per day, while other packets that have the first three hexets different than `foo` are put in separate rate-limiting buckets.

```
"Dhcp6": {
  "client-classes": [
    {
      "name": "Client-ID",
      "template-test": "substring(option[1].hex,0,3)",
      "user-context" : {
        "limits": {
          "rate-limit": "60 packets per day"
        }
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

        },
        ...
    },
    ...
],
...
}

```

16.5 Using Static Host Reservations in Classification

Classes can be statically assigned to clients using techniques described in *Reserving Client Classes in DHCPv4* and *Reserving Client Classes in DHCPv6*.

Subnet host reservations are searched after subnet selection. Global host reservations are searched at the same time by default but the `early-global-reservations-lookup` allows to change this behavior into searching them before the subnet selection.

Pool selection is performed after all host reservations lookups.

16.6 Configuring Subnets With Class Information

In certain cases it is beneficial to restrict access to certain subnets only to clients that belong to a given class, using the `client-classes` parameter when defining the subnet. This parameter may be used to specify a list of one or more classes to which clients must belong in order to use the subnet. This can be referred to as a class guard for the subnet, or in other words the subnet is guarded by a client class.

Let's assume that the server is connected to a network segment that uses the 192.0.2.0/24 prefix. The administrator of that network has decided that addresses from the range 192.0.2.10 to 192.0.2.20 will be managed by the DHCPv4 server. Only clients belonging to client class `Client_foo` are allowed to use this subnet. Such a configuration can be achieved in the following way:

```

"Dhcp4": {
  "client-classes": [
    {
      "name": "Client_foo",
      "test": "substring(option[61].hex,0,3) == 'foo'",
      "option-data": [
        {
          "name": "domain-name-servers",
          "code": 6,
          "space": "dhcp4",
          "csv-format": true,
          "data": "192.0.2.1, 192.0.2.2"
        }
      ]
    },
    ...
  ],
  "subnet4": [
    {
      "id": 1,

```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.2.0/24",
        "pools": [ { "pool": "192.0.2.10 - 192.0.2.20" } ],
        "client-classes": [ "Client_foo" ]
    },
    ...
],
...
}

```

The following example shows how to restrict access to a DHCPv6 subnet. This configuration restricts use of the addresses in the range 2001:db8:1::1 to 2001:db8:1::FFFF to members of the "Client_enterprise" class.

```

"Dhcp6": {
    "client-classes": [
        {
            "name": "Client_enterprise",
            "test": "substring(option[1].hex,0,6) == 0x0002AABBCCDD",
            "option-data": [
                {
                    "name": "dns-servers",
                    "code": 23,
                    "space": "dhcp6",
                    "csv-format": true,
                    "data": "2001:db8:0::1, 2001:db8:2::1"
                }
            ]
        }
    ],
    ...
],
"subnet6": [
    {
        "id": 1,
        "subnet": "2001:db8:1::/64",
        "pools": [ { "pool": "2001:db8:1::-2001:db8:1::ffff" } ],
        "client-classes": "Client_enterprise"
    }
],
...
}

```

16.7 Configuring Pools With Class Information

Similar to subnets, in certain cases access to certain address or prefix pools must be restricted to only clients that belong to at least one of a list of one or more classes, using the `client-classes` when defining the pool. This can be referred to as a class guard for the pool, or in other words the pool is guarded by a client class.

Let's assume that the server is connected to a network segment that uses the 192.0.2.0/24 prefix. The administrator of that network has decided that addresses from the range 192.0.2.10 to 192.0.2.20 are going to be managed by the DHCPv4 server. Only clients belonging to client class `Client_foo` are allowed to use this pool. Such a configuration can be achieved in the following way:

```

"Dhcp4": {
  "client-classes": [
    {
      "name": "Client_foo",
      "test": "substring(option[61].hex,0,3) == 'foo'",
      "option-data": [
        {
          "name": "domain-name-servers",
          "code": 6,
          "space": "dhcp4",
          "csv-format": true,
          "data": "192.0.2.1, 192.0.2.2"
        }
      ]
    },
    ...
  ],
  "subnet4": [
    {
      "id": 1,
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          "pool": "192.0.2.10 - 192.0.2.20",
          "client-classes": [ "Client_foo" ]
        }
      ]
    },
    ...
  ],
  ...
}

```

The following example shows how to restrict access to an address pool. This configuration restricts use of the addresses in the range 2001:db8:1::1 to 2001:db8:1::FFFF to members of the "Client_enterprise" class.

```

"Dhcp6": {
  "client-classes": [
    {
      "name": "Client_enterprise_",
      "test": "substring(option[1].hex,0,6) == 0x0002AABBCCDD",
      "option-data": [
        {
          "name": "dns-servers",
          "code": 23,
          "space": "dhcp6",
          "csv-format": true,
          "data": "2001:db8:0::1, 2001:db8:2::1"
        }
      ]
    },
    ...
  ],
  ...
}

```

(continues on next page)

(continued from previous page)

```

"subnet6": [
  {
    "id": 1,
    "subnet": "2001:db8:1::/64",
    "pools": [
      {
        "pool": "2001:db8:1::-2001:db8:1::ffff",
        "client-classes": [ "Client_foo" ]
      }
    ]
  },
  ...
],
...
}

```

16.8 Class Priority

Client classes in Kea follow the order in which they are specified in the configuration (vs. alphabetical order). Additional classes are ordered by pool, subnet, and then shared-network and within each scope by the order in which they appear in `evaluate-additional-classes`.

When determining which client-class information (comprised of options, lease lifetimes, or DHCPv4 field values) is part of the class definitions to be included in the response, the server examines the union of options from all of the assigned classes. If two or more classes include the same class information, the value from the first assigned class is used. ALL is always the first class, i.e. the class with the highest priority, and matching required classes are last, so they have the lowest priority.

Options defined in classes override any global options, and in turn are overridden by options defined for an individual subnet, shared network, pool, or reservation.

On the other hand, lease lifetimes and DHCPv4 field values defined at class scope override any values defined globally, in a subnet scope, or in a shared-network scope.

Note

Because additional evaluation occurs after lease assignment, parameters that would otherwise impact lease lifetimes (e.g. `valid-lifetime`, `offer-lifetime`) will have no effect when specified in a class that also sets `only-in-additional-list` true.

As an example, imagine that an incoming packet matches two classes. Class `foo` defines values for an NTP server (option 42 in DHCPv4) and an SMTP server (option 69 in DHCPv4), while class `bar` defines values for an NTP server and a POP3 server (option 70 in DHCPv4). The server examines the three options - NTP, SMTP, and POP3 - and returns any that the client requested. As the NTP server was defined twice, the server chooses only one of the values for the reply; the class from which the value is obtained is determined as explained in the previous paragraphs.

16.9 Option Class-Tagging

Option class-tagging allows an option value to be conditionally applied to the response based on the client's class membership. The effect is similar to using an if-block in ISC DHCP to conditionally include options at a given scope. Class-tagging is done by specifying a list of one or more class names in the option's `client-classes` entry.

Consider a case where members of a given class need the same value for one option but a subnet-specific value for another option. The following example shows class-tagging used to give clients who belong to "GROUP1" a subnet-specific value for option "bar", while giving all members of "GROUP1" the same value for option "foo":

```
{
  "client-classes": [
    {
      "name": "GROUP1",
      "test": "substring(option[123].hex,0,4) == 'ONE'",
      "option-data": [{
        "name": "foo",
        "data": "somefile.txt"
      }]
    },
    {
      "subnet4": [
        {
          "id": 1,
          "subnet": "178.16.1.0/24",
          "option-data": [{
            "client-classes": ["GROUP1"],
            "name": "bar",
            "data": 123
          }]
        },
        {
          "id": 2,
          "subnet": "178.16.2.0/24",
          "option-data": [
            {
              "client-classes": ["GROUP1"],
              "name": "bar",
              "data": 789
            }
          ]
        }
      ]
    }
  ]
}
```

The `client-classes` list is allowed in an option specification at any scope. Option class-tagging is enforced at the time options are being added to the response which occurs after lease assignment just before the response is to be sent to the client.

When `never-send` for an option is true at any scope, all `client-classes` entries for that option are ignored. The option will not be included.

When `always-send` is true at any scope, the option will be included unless, the option determined by scope specifies a `client-classes` list that does not contain any of the client's classes.

Otherwise, an option requested by the client will be included in the response if the option either does not specify `client-classes` or the client belongs to at least one of the classes in `client-classes`.

When an option's class-tag does not match, it is as though the option was not specified at that scope. In the following

example the option "foo" is specified for both the subnet and the pool. The pool specification includes a class-tag that limits the option to members of class "melon":

```
{
  "id": 100,
  "subnet": "178.16.1.0/24",
  "option-data": [{
    "name": "foo",
    "data": 456
  }],
  "pools": [{
    "pool": "178.16.1.100 - 178.16.1.200",
    "option-data": [{
      "name": "foo",
      "data": 123,
      "client-classes": [ "melon" ]
    }]
  }]
}
```

Clients that match class "melon" will have a value of 123 for option "foo", while clients that do not match "melon" will have a value of 456 for option "foo".

It is possible to achieve an if-elseif-else effect but specifying an option more than once with different class tags and values. Consider the following configuration for a subnet which provides three possible values for the "server-str" option:

```
{
  "id": 1,
  "subnet": "192.0.1.0/24",
  "pools": [{ "pool": "192.0.1.1 - 192.0.1.255" }],
  "option-data": [{
    "client-classes": [ "class-one" ],
    "name": "server-str",
    "data": "string.one"
  }, {
    "client-classes": [ "class-two" ],
    "name": "server-str",
    "data": "string.two"
  }, {
    "name": "server-str",
    "data": "string.other"
  }]
}
```

Clients belonging to "class-one" will get a value of "string.one", clients belonging to "class-two" will get a value of "string.two", and clients that belong to neither "class-one" nor "class-two" will get a value of "string.other".

Note

The order that the options are tested is not guaranteed other than an option with an empty "client-classes" list is checked last.

Note

Though examples above are for DHCPv4, class-tagging syntax and behavior is the same for DHCPv6.

16.10 Classes and Hooks

Hooks may be used to classify packets. This may be useful if the expression would be complex or time-consuming to write, and could be better or more easily written as code. Once the hook has added the proper class name to the packet, the rest of the classification system will work as expected in choosing a subnet and selecting options. For a description of hooks, see *Hook Libraries*; for information on configuring classes, see *Configuring Classes* and *Configuring Subnets With Class Information*.

16.11 Debugging Expressions

While constructing classification expressions, administrators may find it useful to enable logging; see *Logging* for a more complete description of the logging facility.

To enable the debug statements in the classification system, the severity must be set to `DEBUG` and the debug level to at least 55. The specific loggers are `kea-dhcp4.eval` and `kea-dhcp6.eval`.

To understand the logging statements, it is essential to understand a bit about how expressions are evaluated; for a more complete description, refer to [the design document](<https://gitlab.isc.org/isc-projects/kea/-/wikis/designs/client-classification-design>). In brief, there are two structures used during the evaluation of an expression: a list of tokens which represent the expressions, and a value stack which represents the values being manipulated.

The list of tokens is created when the configuration file is processed, with most expressions and values being converted to a token. The list is organized in reverse Polish notation. During execution, the list is traversed in order; as each token is executed, it is able to pop values from the top of the stack and eventually push its result on the top of the stack. Imagine the following expression:

```
"test": "substring(option[61].hex,0,3) == 'foo',
```

This will result in the following tokens:

```
option, number (0), number (3), substring, text ('foo'), equals
```

In this example, the first three tokens will simply push values onto the stack. The substring token will then remove those three values and compute a result that it places on the stack. The text option also places a value on the stack, and finally the equals token removes the two tokens on the stack and places its result on the stack.

When debug logging is enabled, each time a token is evaluated it emits a log message indicating the values of any objects that were popped off of the value stack, and any objects that were pushed onto the value stack.

The values are displayed as either text, if the command is known to use text values, or hexadecimal, if the command either uses binary values or can manipulate either text or binary values. For expressions that pop multiple values off the stack, the values are displayed in the order they were popped. For most expressions this will not matter, but for the concat expression the values are displayed in reverse order from their written order in the expression.

Let us assume that the following test has been entered into the configuration. This example skips most of the configuration to concentrate on the test.

```
"test": "substring(option[61].hex,0,3) == 'foo',
```

The logging might then resemble this:

```
2016-05-19 13:35:04.163 DEBUG [kea.eval/44478] EVAL_DEBUG_OPTION Pushing option 61 with ↵
↵value 0x666F6F626172
2016-05-19 13:35:04.164 DEBUG [kea.eval/44478] EVAL_DEBUG_STRING Pushing text string '0'
2016-05-19 13:35:04.165 DEBUG [kea.eval/44478] EVAL_DEBUG_STRING Pushing text string '3'
2016-05-19 13:35:04.166 DEBUG [kea.eval/44478] EVAL_DEBUG_SUBSTRING Popping length 3, ↵
↵start 0, string 0x666F6F626172 pushing result 0x666F6F
2016-05-19 13:35:04.167 DEBUG [kea.eval/44478] EVAL_DEBUG_STRING Pushing text string 'foo
↵'
2016-05-19 13:35:04.168 DEBUG [kea.eval/44478] EVAL_DEBUG_EQUAL Popping 0x666F6F and ↵
↵0x666F6F pushing result 'true'
```

Note

The debug logging may be quite verbose if there are multiple expressions to evaluate; it is intended as an aid in helping create and debug expressions. Administrators should plan to disable debug logging when expressions are working correctly. Users may also wish to include only one set of expressions at a time in the configuration file while debugging them, to limit the log statements. For example, when adding a new set of expressions, an administrator might find it more convenient to create a configuration file that only includes the new expressions until they are working correctly, and then add the new set to the main configuration file.

HOOK LIBRARIES

17.1 Introduction

Kea is both flexible and customizable, via the use of "hooks." This feature lets Kea load one or more dynamically linked libraries (known as "hook libraries") and call functions in them at various points in its processing ("hook points"). Those functions perform whatever custom processing is required.

The hooks concept allows the core Kea code to remain reasonably small by moving features that only some, but not all, users find useful to external libraries. Those with no need for certain functions can simply choose not to load those libraries.

Hook libraries are loaded by individual Kea processes, not by Kea as a whole. This means, among other things, that it is possible to associate one set of libraries with the DHCPv4 server and a different set with the DHCPv6 server.

It is also possible for a process to load multiple libraries. When processing reaches a hook point, Kea calls the hook library functions attached to it. If multiple libraries have attached a function to a given hook point, Kea calls all of them, in the order in which the libraries are specified in the configuration file. The order may be important; consult the documentation of the libraries for specifics.

When a Kea process unloads a library, it expects the `dllclose` function to remove all library symbols, as well as the library code, from address space. Although most OSes implement the `dllclose` function, this behavior is not required by the POSIX standard and not all systems support it; for example, the musl library, used by default by Alpine Linux, implements the `dllclose` function as a no operation. On such systems a library actually remains loaded for the lifetime of the process, which means that it must be restarted to update libraries with newer versions; it is not sufficient to simply reconfigure or reload the Kea process.

The next sections describe how to install and configure hook libraries. Users who are interested in writing their own hook library can find information in the [Hooks Developer's Guide](#) section of the [Kea Developer's Guide](#).

Note that some libraries are available under different licenses.

Please also note that some libraries may require additional dependencies and/or compilation switches to be enabled.

17.2 Installing Hook Packages

Note

For more details about installing the Kea Subscriber Hooks, please read [this Knowledgebase article](#).

Most hook packages are now included in the base Kea sources. There is no need to do anything special to compile or install them, as they are covered by the usual building and installation procedures. Please refer to [Installation](#) for a general overview of the installation process.

ISC provides several additional subscriber-only hooks in the form of packages, which follow a similar installation procedure but with several additional steps. For our users' convenience, the subscriber hooks' installation procedure is described in this section.

1. Download the package; detailed instructions are provided in the KB article above. The package will be a file with a name similar to `kea-subscriber-|release|.tar.xz`.

2. Administrators who have the sources for the corresponding version of the open-source Kea package on their system from the initial Kea installation should skip this step. Otherwise, extract the Kea source from the original tarball that was downloaded. For example, with a download of Kea 3.0.2, there should be a tarball called `kea-|release|.tar.xz` on the system. Unpack this tarball:

```
$ tar -Jxvf kea-3.0.2.tar.xz
```

This will unpack the tarball into the `kea-|release|` subdirectory of the current working directory.

3. Unpack the Kea subscriber hooks tarball into the same directory where the original Kea source is located. Once Kea 3.0.2 has been unpacked into a `kea-|release|` subdirectory and the Kea subscriber tarball is in the current directory, the following steps will unpack the subscriber tarball into the correct location:

```
$ cd kea-3.0.2
$ tar -Jxvf ../kea-subscriber-3.0.2.tar.xz
```

Note that unpacking the Kea subscriber package puts the files into a directory named **premium**. Regardless of the name of the package, the directory is always called **premium**.

4. Set up the build. Make sure **meson** and **ninja** are installed on your system. The first section of the output should look something like this:

```
$ meson setup build
```

Package:

Name:	kea
Version:	3.0.2
Extended version:	3.0.2 (tarball)
Version type:	development
OS Family:	Linux
Prefix:	/opt/kea
Hooks directory:	/opt/kea/lib/kea/hooks
Premium hooks:	yes

6. Compile Kea.

```
$ meson compile -C build
```

7. Install Kea sources along with the hooks:

```
$ sudo make install
```

The installation location of the hook libraries depends on whether the `--prefix` parameter was specified at the **meson setup** step. If not, the default location is `/usr/local/lib/kea/hooks`. The proper installation of the libraries can be verified with this command:

```
$ ls -l /usr/local/lib/kea/hooks/*.so
/usr/local/lib/kea/hooks/libddns_gss_tsig.so
/usr/local/lib/kea/hooks/libdhcp_bootp.so
/usr/local/lib/kea/hooks/libdhcp_class_cmds.so
```

(continues on next page)

(continued from previous page)

```

/usr/local/lib/kea/hooks/libdhcp_ddns_tuning.so
/usr/local/lib/kea/hooks/libdhcp_flex_id.so
/usr/local/lib/kea/hooks/libdhcp_flex_option.so
/usr/local/lib/kea/hooks/libdhcp_ha.so
/usr/local/lib/kea/hooks/libdhcp_host_cache.so
/usr/local/lib/kea/hooks/libdhcp_host_cmds.so
/usr/local/lib/kea/hooks/libdhcp_lease_cmds.so
/usr/local/lib/kea/hooks/libdhcp_lease_query.so
/usr/local/lib/kea/hooks/libdhcp_legal_log.so
/usr/local/lib/kea/hooks/libdhcp_limits.so
/usr/local/lib/kea/hooks/libdhcp_mysql.so
/usr/local/lib/kea/hooks/libdhcp_perfmon.so
/usr/local/lib/kea/hooks/libdhcp_pgsql.so
/usr/local/lib/kea/hooks/libdhcp_ping_check.so
/usr/local/lib/kea/hooks/libdhcp_radius.so
/usr/local/lib/kea/hooks/libdhcp_run_script.so
/usr/local/lib/kea/hooks/libdhcp_stat_cmds.so
/usr/local/lib/kea/hooks/libdhcp_subnet_cmds.so

```

with the following subscriber libraries:

```

/usr/local/lib/kea/hooks/libca_rbac.so -> libdhcp_rbac.so
/usr/local/lib/kea/hooks/libdhcp_cb_cmds.so
/usr/local/lib/kea/hooks/libdhcp_rbac.so

```

The exact list returned depends on the packages installed. If the directory was specified via `--prefix`, the hook libraries will be located in `{prefix}directory}/lib/kea/hooks`.

17.3 Configuring Hook Libraries

The hook libraries for a given process are configured using the `hooks-libraries` keyword in the configuration for that process. (Note that the word "hooks" is plural.) The value of the keyword is an array of map structures, with each structure corresponding to a hook library. For example, to set up two hook libraries for the DHCPv4 server, the configuration would be:

```

{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/opt/first_custom_hooks_example.so"
      },
      {
        "library": "/opt/local/second_custom_hooks_example.so",
        "parameters": {
          "mail": "spam@example.com",
          "floor": 13,
          "debug": false,
          "users": [
            "alice",
            "bob",
            "charlie"
          ]
        }
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

        "languages": {
            "french": "bonjour",
            "klíngon": "yl'el"
        }
    }
}

```

Note

Libraries are reloaded even if their lists have not changed, because the parameters specified for the library (or the files those parameters point to) may have changed.

As of Kea 2.7.9, hook libraries may only be loaded from the default installation directory determined during compilation and shown in the config report as "Hooks directory". This value may be overridden at startup by setting the environment variable `KEA_HOOKS_PATH` to the desired path. If a path other than this value is used in a `library` element Kea will emit an error and refuse to load the library. For ease of use `library` elements may simply omit path components, specifying the file name only as shown below:

```

{
    "Dhcp6": {
        "hooks-libraries": [
            {
                "library": "first_custom_hooks_example.so"
            },
            {
                "library": "second_custom_hooks_example.so"
            }
        ]
    }
}

```

This snippet (on Debian 12) is equivalent to:

```

{
    "Dhcp6": {
        "hooks-libraries": [
            {
                "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/first_custom_hooks_
↪example.so"
            },
            {
                "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/second_custom_hooks_
↪example.so"
            }
        ]
    }
}

```

Note

The support for path components in `library` elements is deprecated as of 2.7.8, and will eventually be removed.

Libraries may have additional parameters that are not mandatory, in the sense that there may be libraries that do not require them. However, for any given library there is often a requirement to specify a certain set of parameters. Please consult the documentation for each individual library for details. In the example above, the first library (`/opt/first_custom_hooks_example.so`) has no parameters. The second library (`/opt/local/second_custom_hooks_example.so`) has five parameters: specifying mail (string parameter), floor (integer parameter), debug (boolean parameter), lists (list of strings), and maps (containing strings). Nested parameters can be used if the library supports it. This topic is explained in detail in the [Hooks Developer's Guide](#) section of the [Kea Developer's Guide](#).

Some hooks use user context to set the parameters. See *User Contexts in Hooks*.

Notes:

- The full path to each library should be given.
- As noted above, the order in which the hooks are called may be important; consult the documentation for each library for specifics.
- An empty list has the same effect as omitting the `hooks-libraries` configuration element altogether.

Note

There is one case where this is not true: if Kea is running with a configuration that contains a `hooks-libraries` item, and that item is removed and the configuration reloaded, the removal is ignored and the libraries remain loaded. As a workaround, instead of removing the `hooks-libraries` item, change it to an empty list.

At the moment, only the `kea-dhcp4` and `kea-dhcp6` processes support hook libraries.

17.3.1 Order of Configuration:

It is important to recognize that the order in which hook libraries are configured determines the order in which their callouts will be executed, in cases where more than one hook library implements the same callout. For example, to use the Flexible Identifier (also called Flex ID) hook library to formulate the client IDs in conjunction with the High Availability (HA) hook library for load-balanced HA, it is essential that the Flex ID library be specified first in the server's `hooks-libraries` section. This ensures that the client ID is formulated by the Flex ID library before the HA library uses it for load-balancing. Similarly, it is best to specify the Forensic Logging library last, to ensure that any other installed hooks have already made their contributions to the packet processing before that one is loaded.

17.3.2 User Contexts in Hooks

Hook libraries can have their own configuration parameters, which is convenient if the parameter applies to the whole library. However, sometimes it is useful to extend certain configuration entities with additional configuration data. This is where the concept of user contexts comes in. A system administrator can define an arbitrary set of data and attach it to Kea structures, as long as the data is specified as a JSON map. In particular, it is possible to define fields that are integers, strings, boolean, lists, or maps. It is possible to define nested structures of arbitrary complexity. Kea does not use that data on its own; it simply stores it and makes it available for the hook libraries.

Another use case for user contexts may be storing comments and other information that will be retained by Kea. Regular comments are discarded when the configuration is loaded, but user contexts are retained. This is useful if administrators want their comments to survive `config-set` or `config-get` operations, for example.

If user context is supported in a given context, the parser translates "comment" entries into user context with a "comment" entry.

User context can store configuration for multiple hooks and comments at once.

Some hooks use user context for a configuration that can be easily edited without the need to restart the server.

The DDNS Tuning hook uses user context to configure per-subnet behavior. Here's an example:

```
"subnet4": [{
  "id": 1,
  "subnet": "192.0.2.0/24",
  "pools": [{
    "pool": "192.0.2.10 - 192.0.2.20"
  } ],
  "user-context": {
    "ddns-tuning": {
      "hostname-expr": "'guest-'+int8totext(substring(pkt4.yiaddr, 0,1))+ '-' \
                        +int8totext(substring(pkt4.yiaddr, 1,2))+ '-' \
                        +int8totext(substring(pkt4.yiaddr, 2,3))+ '-' \
                        +int8totext(substring(pkt4.yiaddr, 3,4))"
    },
    "last-modified": "2017-09-04 13:32",
    "phones": [ "x1234", "x2345" ],
    "devices-registered": 42,
    "billing": false
  }
}]
```

The Limits hook uses user-context in classes and subnets to set parameters. For example:

```
{
  "Dhcp6": {
    "client-classes": [
      {
        "name": "gold",
        "user-context": {
          "limits": {
            "address-limit": 2,
            "prefix-limit": 1,
            "rate-limit": "1000 packets per second"
          }
        }
      ]
    },
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/libdhcp_limits.so"
      }
    ],
    "subnet6": [
      {
        "id": 1,
        "pools": [
          {
```

(continues on next page)

(continued from previous page)

```

        "pool": "2001:db8::/64"
    },
    "subnet": "2001:db8::/64",
    "user-context": {
        "limits": {
            "address-limit": 4,
            "prefix-limit": 2,
            "rate-limit": "10 packets per minute"
        }
    }
}
]
}
}

```

17.3.3 Parked-Packet Limit

Kea servers contain a mechanism by which the response to a client packet may be held, pending completion of hook library work. We refer to this as "parking the packet." When work is ready to continue, the server unparks the response and continues processing.

There is a global parameter, `parked-packet-limit`, that may be used to limit the number of responses that may be parked at any given time. This acts as a form of congestion handling and protects the server from being swamped when the volume of client queries is outpacing the server's ability to respond. Once the limit is reached, the server emits a log and drops any new responses until parking spaces are available.

In general, smaller values for the parking lot limit are likely to cause more drops but with shorter response times; larger values are likely to result in fewer drops but with longer response times. Currently, the default value for `parked-packet-limit` is 256.

Warning

Using too small a value may result in an unnecessarily high drop rate, while using too large a value may lead to response times that are simply too long to be useful. A value of 0, while allowed, disables the limit altogether, but this is highly discouraged as it may lead to Kea servers becoming unresponsive to clients. Choosing the best value is very site-specific; we recommend users initially leave it at the default value of 256 and observe how the system behaves over time with varying load conditions.

Here is an example of the global parameter used with `libdhcp_ha.so`. It lowers the number of concurrently parked packets to 128.

```

{
  "Dhcp6": {
    "parked-packet-limit": 128
    "hooks-libraries": [
      {
        "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so"
      },
      {
        "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
        "parameters": {

```

(continues on next page)

(continued from previous page)

```

    "high-availability": [
    {
        "mode": "hot-standby",
        "peers": [
            {
                "name": "server1",
                "role": "primary",
                "url": "http://127.0.0.1:8080/"
            },
            {
                "name": "server2",
                "role": "standby",
                "url": "http://127.0.0.1:8088/"
            }
        ],
        "this-server-name": "server1"
    }
  ]
}

```

17.4 Available Hook Libraries

As described above, the hook functionality provides a way to customize a Kea server without modifying the core code. ISC has chosen to take advantage of this feature to provide functions that may only be useful to a subset of Kea users. To this end, ISC has created some hook libraries, discussed in the following sections.

Note

As of Kea 3.0.0, most of these libraries are available at no cost with the open source base code; a few are only available to organizations that contribute to Kea's development through paid ISC support contracts. Paid support includes professional engineering assistance, advance security notifications, input into ISC's roadmap planning, and many other benefits, while helping keep Kea sustainable in the long term. ISC encourages companies and organizations to consider purchasing a paid support contract; further information can be obtained by completing the form at <https://www.isc.org/contact>.

The following table provides a list of hook libraries currently available from ISC. It is important to pay attention to which libraries may be loaded by which Kea processes. It is a common mistake to configure the *kea-ctrl-agent* process to load libraries that should, in fact, be loaded by the *kea-dhcp4* or *kea-dhcp6* processes. If a library from ISC does not work as expected, please make sure that it has been loaded by the correct process per the table below.

Table 1: List of available hook libraries

Name	Availability	Description
<i>BOOTP</i>	Kea open source	This hook library adds BOOTP support, as defined in RFC 1497. It recognizes received BOOTP requests: they are translated into DHCPREQUEST packets, put into the BOOTP client class, and receive infinite lifetime leases.
<i>Class Commands</i>	Kea open source	This hook library allows configured DHCP client classes to be added, updated, deleted, and fetched without needing to restart the DHCP server.
<i>Configuration Backend Commands</i>	ISC support customers	This hook library implements a collection of commands to manage Kea configuration information in a database. This library may only be used in conjunction with one of the supported Configuration Backend implementations.
<i>DDNS Tuning</i>	Kea open source	This hook library adds custom behaviors related to Dynamic DNS updates on a per-client basis. Its primary feature is to allow the host name used for DNS to be calculated using an expression.
<i>Flexible Identifier</i>	Kea open source	Kea software provides a way to handle host reservations that include addresses, prefixes, options, client classes and other features. The reservation can be based on hardware address, DUID, circuit-id, or client-id in DHCPv4 and on hardware address or DUID in DHCPv6. However, there are sometimes scenarios where the reservation is more complex, e.g. uses other options than mentioned above, uses parts of specific options, or perhaps uses a combination of several options and fields to uniquely identify a client. Those scenarios are addressed by the Flexible Identifier hook application. It allows defining an expression, similar to the one used in client classification, e.g. <code>substring(relay6[0].option[37], 0, 6)</code> . Each incoming packet is evaluated against that expression and its value is then searched in the reservations database.
<i>Flexible Option</i>	Kea open source	This library provides hooks that compute option values instead of static configured values. An expression is evaluated on the query packet. Defined add, supersede, and remove actions are applied on the response packet before it is sent using the evaluation result.
<i>Forensic Logging</i>	Kea open source	This library provides hooks that record a detailed log of lease assignments and renewals in a set of log files. In many legal jurisdictions, companies - especially ISPs - must record information about the addresses they have leased to DHCP clients. This library is designed to help with that requirement. If the information that it records is sufficient, it may be used directly. If a jurisdiction requires a different set of information to be saved, it can be used as a template or example to create custom logging hooks. In Kea 1.9.8, additional parameters were added to give users more flexibility regarding what information should be logged.
<i>GSS-TSIG</i>	Kea open source	This hook library adds support to the Kea D2 server (kea-dhcp-ddns) for using GSS-TSIG to sign DNS updates.
<i>High Availability</i>	Kea open source	The risk of DHCP service unavailability can be minimized by setting up a pair of DHCP servers in a network. Two modes of operation are supported. The first one is called load-balancing, and is sometimes referred to as "active-active." Each server can handle selected groups of clients in this network, or all clients if it detects that its partner has become unavailable. It is also possible to designate one server to serve all DHCP clients, and leave another server as standby. This mode is called "hot standby" and is sometimes referred to as "active-passive." This server activates its DHCP function only when it detects that its partner is not available. Such cooperation between the DHCP servers requires that these servers constantly communicate with each other to send updates about allocated leases, and to periodically test whether their partners are still operational. The hook library also provides an ability to send lease updates to external backup servers, making it much easier to have a replacement that is up to date.
<i>Host Cache</i>	Kea open source	Some database backends, such as RADIUS, may take a long time to respond. Since Kea in general is synchronous, backend performance directly affects DHCP performance. To minimize performance impact, this library provides a way to cache responses from other hosts. This includes negative caching, i.e. the ability to remember that there is no client information in the database.

continues on next page

Table 1 – continued from previous page

Name	Availability	Description
<i>Host Commands</i>	Kea open source	Kea provides a way to store host reservations in a database. In many larger deployments it is useful to be able to manage that information while the server is running. This library provides management commands for adding, querying, and deleting host reservations in a safe way without restarting the server. In particular, it validates the parameters, so an attempt to insert incorrect data, e.g. add a host with conflicting identifier in the same subnet, is rejected. Those commands are exposed via the command channel (JSON over UNIX sockets) and the Control Agent (JSON over RESTful interface).
<i>Lease Commands</i>	Kea open source	This hook library offers a number of commands used to manage leases. Kea can store lease information in various backends: memfile, MySQL, PostgreSQL. This library provides a unified interface to manipulate leases in a unified, safe way. In particular, it allows manipulation of memfile leases while Kea is running, sanity check changes, lease existence checks, and removal of all leases belonging to a specific subnet. It can also catch obscure errors, like the addition of a lease with subnet-id that does not exist in the configuration, or configuration of a lease to use an address that is outside of the subnet to which it is supposed to belong. This library allows easy management of user contexts associated with leases.
<i>Leasequery</i>	Kea open source	This library adds support for DHCPv4 Leasequery (RFC 4388), DHCPv4 Bulk Leasequery (RFC6926); DHCPv6 Leasequery (RFC 5007), and DHCPv6 Bulk Leasequery (RFC5460).
<i>Limits</i>	Kea open source	With this hook library, <i>kea-dhcp4</i> and <i>kea-dhcp6</i> servers can apply a limit to the rate at which packets receive a response. The limit can be applied per-client class or per-subnet.
<i>MySQL Configuration Backend</i>	Kea open source	This hook library is an implementation of the Kea Lease, Host and Configuration Backend for MySQL. It uses a MySQL database as a repository for the Kea leases, host reservations and configuration information. Kea servers use this library to fetch their configurations if Configuration Backend is used.
<i>PerfMon</i>	Kea open source	With this hook library, <i>kea-dhcp4</i> and <i>kea-dhcp6</i> servers can track and report performance data.
<i>Ping Check</i>	Kea open source	With this hook library, the <i>kea-dhcp4</i> server can perform ping checks of candidate lease addresses before offering them to clients.
<i>PostgreSQL Database Backend</i>	Kea open source	This hook library is an implementation of the Kea Lease, Host and Configuration Backend for PostgreSQL. It uses a PostgreSQL database as a repository for the Kea leases, host reservations and configuration information. Kea servers use this library to fetch their configurations if Configuration Backend is used.
<i>RADIUS</i>	Kea open source	The RADIUS hook library allows Kea to interact with RADIUS servers using access and accounting mechanisms. The access mechanism may be used for access control, assigning specific IPv4 or IPv6 addresses reserved by RADIUS, dynamically assigning addresses from designated pools chosen by RADIUS, or rejecting the client's messages altogether. The accounting mechanism allows a RADIUS server to keep track of device activity over time.
<i>RBAC</i>	ISC support customers	This hook library adds support to the Kea Control Agent (kea-ctrl-agent) and Kea servers (kea-dhcp4, kea-dhcp6 and kea-dhcp-ddns) for Role-Based Access Control filtering of commands.
<i>Run Script</i>	Kea open source	This hook library adds support to run external scripts for specific packet-processing hook points. There are several exported environment variables available for the script.
<i>Statistics Commands</i>	Kea open source	This library provides additional commands for retrieving accurate DHCP lease statistics, for Kea DHCP servers that share the same lease database. This setup is common in deployments where DHCP service redundancy is required and a shared lease database is used to avoid lease-data replication between the DHCP servers. This hook library returns lease statistics for each subnet.

continues on next page

Table 1 – continued from previous page

Name	Availability	Description
<i>Subnet Commands</i>	Kea open source	In deployments in which subnet configuration needs to be frequently updated, it is a hard requirement that such updates be performed without the need for a full DHCP server reconfiguration or restart. This hook library allows for incremental changes to the subnet configuration such as adding or removing a subnet. It also allows for listing all available subnets and fetching detailed information about a selected subnet. The commands exposed by this library do not affect other subnets or configuration parameters currently used by the server.
<i>User Check</i>	Kea open source	Reads known users list from a file. Unknown users will be assigned a lease from the last subnet defined in the configuration file, e.g. to redirect them to a captive portal. This demonstrates how an external source of information can be used to influence the Kea allocation engine. This hook is part of the Kea source code and is available in the <code>src/hooks/dhcp/user_chk</code> directory.

ISC hopes to see more hook libraries become available as time progresses, developed both internally and externally. Since this list may evolve dynamically, it is maintained on a wiki page, available at <https://gitlab.isc.org/isc-projects/kea/wikis/Hooks-available>. Developers or others who are aware of any hook libraries not listed there are asked to please send a note to the `kea-users` or `kea-dev` mailing lists for updating. (Information on all of ISC's public mailing lists can be found at <https://www.isc.org/maillinglists/>.)

The libraries developed by ISC are described in detail in the following sections.

17.5 libdhcp_bootp.so: Support for BOOTP Clients

This hook library adds support for BOOTP with vendor-information extensions (RFC 1497). Received BOOTP requests are recognized, translated into DHCPREQUEST packets by adding a `dhcp-message-type` option, and put into the "BOOTP" client class. Members of this class get infinite lifetime leases but the class can also be used to guard a pool of addresses.

The DHCP-specific options, such as `dhcp-message-type`, are removed from the server's responses; responses shorter than the BOOTP minimum size of 300 octets are padded to this size.

Note

`libdhcp_bootp.so` is part of the open source code and is available to every Kea user.

Note

This library can only be loaded by the `kea-dhcp4` process, as there is no BOOTP protocol for IPv6.

This library is loaded similarly to other hook libraries, and it takes no parameters.

```
"Dhcp4": {
  "hooks-libraries": [
    { "library": "/usr/local/lib/libdhcp_bootp.so" },
    ...
  ]
}
```


Note

A host reservation for a BOOTP client should use the hardware address as the identifier (the `client-id` option is a DHCP-specific option).

Incoming BOOTP packets are added to the BOOTP class, allowing administrators to segregate BOOTP clients into separate pools. For example:

```
"Dhcp4": {
  "client-classes": [
    {
      // The DHCP class is the complement of the BOOTP class
      "name": "DHCP",
      "test": "not member('BOOTP')"
    }
  ],
  "subnet4": [
    {
      "subnet": "192.0.2.0/24",
      "pools": [
        {
          // BOOTP clients will be handled here
          "pool": "192.0.2.200 - 192.0.2.254",
          "client-classes": [ "BOOTP" ]
        },
        {
          // Regular DHCP clients will be handled here
          "pool": "192.0.2.1 - 192.0.2.199",
          "client-classes": [ "DHCP" ]
        }
      ],
      ...
    },
    ...
  ],
  ...
}
```

17.5.1 BOOTP Hooks Limitations

Currently the BOOTP library has the following limitation:

- Basic BOOTP, as defined in [RFC 951](#), is not supported. Kea only supports BOOTP with vendor-information extensions.

17.6 libdhcp_cb_cmds.so: Configuration Backend Commands

This hook library is used to manage Kea servers' configurations in a configuration backend database. This library must be used in conjunction with the available CB hook libraries implementing the common APIs to create, read, update, and delete (CRUD) the configuration information in the respective databases. For example: *libdhcp_mysql.so* implements this API for MySQL while *libdhcp_pgsql.so* implements this API for PostgreSQL. To manage the configuration information in a MySQL database, both *libdhcp_mysql.so* and *libdhcp_cb_cmds.so* must be loaded by the server used for configuration management. To manage the configuration information in a PostgreSQL

database, both *libdhcp_pgsql.so* and *libdhcp_cb_cmds.so* must be loaded by the server used for configuration management.

More information on how to configure the Configuration Backend hook library for use with a MySQL or PostgreSQL database can be found in the *Configuration Backend in DHCPv4* and *Configuration Backend in DHCPv6* sections.

Note

libdhcp_cb_cmds.so is available only to ISC customers with a paid support contract. For more information on subscription options, please complete the form at <https://www.isc.org/contact>.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

Note

Please read about *Limitations and Warnings* before using the commands described in this section.

17.6.1 Command Structure

There are 5 types of commands supported by this library:

- **del** - delete the selected object from the database, e.g. *remote-global-parameter4-del*.
- **get** - fetch the selected object from the database, e.g. *remote-subnet4-get-by-id*.
- **get-all** - fetch all objects of the particular type from the database, e.g. *remote-option-def4-get-all*.
- **list** - list all objects of the particular type in the database, e.g. *remote-network4-list*; this class of commands returns brief information about each object compared to the output of **get-all**.
- **set** - creates or replaces an object of the given type in the database, e.g. *remote-option4-global-set*.

All types of commands accept an optional **remote** map which selects the database instance to which the command refers. For example:

```
{
  "command": "remote-subnet4-list",
  "arguments": {
    "remote": {
      "type": "mysql",
      "host": "192.0.2.33",
      "port": 3302
    }
  }
}
```

selects the MySQL database, running on host 192.0.2.33 and port 3302, to fetch the list of subnets from. All parameters in the **remote** argument are optional. The **port** parameter can be only specified in conjunction with the **host**. If no options in the **remote** parameter are to be specified, the parameter should be omitted. In this case, the server will use the first backend listed in the **config-control** map within the configuration of the server receiving the command.

Note

In the current Kea release, it is only possible to configure the Kea server to use a single configuration backend. Strictly speaking, it is possible to point the Kea server to at most one database (either MySQL or PostgreSQL) using the `config-control` parameter. Therefore, the `remote` parameter may be omitted in the commands and `libdhcp_cb_cmds.so` uses the sole backend by default. The example commands below most often show a value of "mysql" for the `type` parameter; it should be assumed that the value is "postgresql" for installations using a PostgreSQL database.

17.6.2 Control Commands for DHCP Servers

This section describes and gives some examples of the control commands implemented by `libdhcp_cb_cmds.so`, to manage the configuration information of the DHCPv4 and DHCPv6 servers. Many of the commands are almost identical between DHCPv4 and DHCPv6; they only differ by the command name. Other commands differ slightly by the structure of the inserted data; for example, the structure of the IPv4 subnet information is different than that of the IPv6 subnet. Nevertheless, they still share the structure of their command arguments and thus it makes sense to describe them together.

In addition, whenever the text in the subsequent sections refers to a DHCP command or DHCP parameter, it refers to both DHCPv4 and DHCPv6 variants. The text specific to the particular server type refers to them as: DHCPv4 command, DHCPv4 parameter, DHCPv6 command, DHCPv6 parameter, etc.

17.6.3 Metadata

The typical response to the `get` or `list` command includes a list of returned objects (e.g. subnets), and each such object contains the metadata map with some database-specific information describing this object. In other words, the metadata contains any information about the fetched object which may be useful for an administrator but which is not part of the object specification from the DHCP server standpoint. In the present Kea release, the metadata is limited to the `server-tag`. It describes the association of the object with a particular server or all servers.

The following is the example response to the `remote-network4-list` command, which includes the metadata:

```
{
  "result": 0,
  "text": "1 IPv4 shared network(s) found.",
  "arguments": {
    "shared-networks": [
      {
        "name": "level3",
        "metadata": {
          "server-tags": [ "all" ]
        }
      }
    ],
    "count": 1
  }
}
```

Client implementations must not assume that the metadata contains only the `server-tags` parameter. In future releases, it is expected that this map will be extended with additional information, e.g. object modification time, log message created during the last modification, etc.

17.6.4 The `remote-server4-del`, `remote-server6-del` Commands

This command is used to delete the information about a selected DHCP server from the configuration database. The server is identified by a unique case insensitive server tag. For example:

```
{
  "command": "remote-server4-del",
  "arguments": {
    "servers": [
      {
        "server-tag": "server1"
      }
    ],
    "remote": {
      "type": "postgresql"
    }
  }
}
```

As a result of this command, all associations of the configuration for the user-defined server called "server1" are removed from the database, including non-shareable configuration information, such as global parameters, option definitions, and global options. Any shareable configuration information, i.e. the configuration elements which may be associated with more than one server, is preserved. In particular, the subnets and shared networks associated with the deleted servers are preserved. If any of the shareable configuration elements was associated only with the deleted server, this object becomes unassigned (orphaned). For example: if a subnet has been created and associated with "server1" using the `remote-subnet4-set` command and "server1" is subsequently deleted, the subnet remains in the database but no servers can use this subnet. The subnet can be updated using the `remote-subnet4-set` command, and can be associated with either another server or with all servers, using the special server tag "all". Such a subnet can be also deleted from the database using the `remote-subnet4-del-by-id` or `remote-subnet4-del-by-prefix` command, if it is no longer needed.

The following is the successful response to the `remote-server4-del` command:

```
{
  "result": 0,
  "text": "1 DHCPv4 server(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Warning

The `remote-server4-del` and `remote-server6-del` commands must be used with care, because an accidental deletion of the server can cause some parts of the existing configurations to be lost permanently from the database. This operation is not reversible. Re-creation of the accidentally deleted server does not revert the lost configuration for that server and such configuration must be re-created manually by the user.

17.6.5 The `remote-server4-get`, `remote-server6-get` Commands

This command is used to fetch the information about the selected DHCP server from the configuration database. For example:

```
{
  "command": "remote-server6-get",
  "arguments": {
    "servers": [
      {
        "server-tag": "server1"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

This command fetches the information about the DHCPv6 server identified by the server tag "server1". The server tag is case-insensitive. A successful response returns basic information about the server, such as the server tag and the user's description of the server:

```
{
  "result": 0,
  "text": "DHCP server server1 found.",
  "arguments": {
    "servers": [
      {
        "server-tag": "server1",
        "description": "A DHCPv6 server located on the first floor."
      }
    ],
    "count": 1
  }
}
```

17.6.6 The remote-server4-get-all, remote-server6-get-all Commands

This command is used to fetch all user-defined DHCPv4 or DHCPv6 servers from the database. The command structure is very simple:

```
{
  "command": "remote-server4-get-all",
  "arguments": {
    "remote": {
      "type": "mysql"
    }
  }
}
```

The response includes basic information about each server, such as its server tag and description:

```
{
  "result": 0,
  "text": "DHCPv4 servers found.",
  "arguments": {
    "servers": [
```

(continues on next page)

(continued from previous page)

```

    {
      "server-tag": "server1",
      "description": "A DHCP server located on the first floor."
    },
    {
      "server-tag": "server2",
      "description": "An old DHCP server to be soon replaced."
    }
  ],
  "count": 2
}

```

17.6.7 The `remote-server4-set`, `remote-server6-set` Commands

This command is used to create or replace an information about a DHCP server in the database. The information about the server must be created when there is a need to differentiate the configurations used by various Kea instances connecting to the same database. Various configuration elements, e.g. global parameters, subnets, etc. may be explicitly associated with the selected servers (using server tags as identifiers), allowing only these servers to use the respective configuration elements. Using the particular server tag to make such associations is only possible when the server information has been stored in the database via the `remote-server4-set` or `remote-server6-set` commands. The following command creates a new (or updates an existing) DHCPv6 server in the database:

```

{
  "command": "remote-server6-set",
  "arguments": {
    "servers": [
      {
        "server-tag": "server1",
        "description": "A DHCP server on the ground floor."
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}

```

The server tag must be unique across all servers in the database. When the server information under the given server tag already exists, it is replaced with the new information. The specified server tag is case-insensitive, and the maximum length of the server tag is 256 characters. The following keywords are reserved and cannot be used as server tags: "all" and "any".

The following is the example response to the above command:

```

{
  "result": 0,
  "text": "DHCPv6 server successfully set.",
  "arguments": {
    "servers": [
      {
        "server-tag": "server1",
        "description": "A DHCP server on the ground floor."
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

17.6.8 The `remote-global-parameter4-del`, `remote-global-parameter6-del` Commands

These commands are used to delete a global DHCP parameter from the configuration database. When the parameter is deleted from the database, the server uses the value specified in the configuration file for this parameter, or a default value if the parameter is not specified in the configuration file.

The following command attempts to delete the DHCPv4 `renew-timer` parameter common for all servers from the database:

```

{
  "command": "remote-global-parameter4-del",
  "arguments": {
    "parameters": [ "renew-timer" ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}

```

If a server-specific parameter is to be deleted, the `server-tags` list must contain the tag of the appropriate server. There must be exactly one server tag specified in this list.

17.6.9 The `remote-global-parameter4-get`, `remote-global-parameter6-get` Commands

These commands are used to fetch a scalar global DHCP parameter from the configuration database.

The following command attempts to fetch the `boot-file-name` parameter for "server1":

```

{
  "command": "remote-global-parameter4-get",
  "arguments": {
    "parameters": [ "boot-file-name" ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}

```

The returned value has one of the four scalar types: string, integer, real, or boolean. Non-scalar global configuration parameters, such as map or list, are not returned by this command.

In the case of the example above, the string value is returned, e.g.:

```

{
  "result": 0,

```

(continues on next page)

(continued from previous page)

```

    "text": "1 DHCPv4 global parameter found.",
    "arguments": {
      "parameters": {
        "boot-file-name": "/dev/null",
        "metadata": {
          "server-tags": [ "all" ]
        }
      },
      "count": 1
    }
  }
}

```

Note that the response above indicates that the returned parameter is associated with "all" servers rather than "server1", used in the command. This indicates that there is no "server1"-specific value in the database and therefore, the value shared by all servers is returned. If there were a "server1"-specific value in the database, that value would be returned instead.

The example response for the integer value is:

```

{
  "result": 0,
  "text": "1 DHCPv4 global parameter found.",
  "arguments": {
    "parameters": {
      "renew-timer": 2000,
      "metadata": {
        "server-tags": [ "server1" ]
      }
    },
    "count": 1
  }
}

```

The real value:

```

{
  "result": 0,
  "text": "1 DHCPv4 global parameter found.",
  "arguments": {
    "parameters": {
      "t1-percent": 0.85,
      "metadata": {
        "server-tags": [ "all" ]
      }
    },
    "count": 1
  }
}

```

Finally, the boolean value:

```

{
  "result": 0,

```

(continues on next page)

(continued from previous page)

```

    "text": "1 DHCPv4 global parameter found.",
    "arguments": {
      "parameters": {
        "match-client-id": true,
        "metadata": {
          "server-tags": [ "server2" ]
        }
      },
      "count": 1
    }
  }
}

```

17.6.10 The remote-global-parameter4-get-all, remote-global-parameter6-get-all Commands

These commands are used to fetch all global DHCP parameters from the database for the specified server. The following example demonstrates how to fetch all global parameters to be used by the server "server1":

```

{
  "command": "remote-global-parameter4-get-all",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}

```

The example response may look as follows:

```

{
  "result": 0,
  "text": "DHCPv4 global parameters found.",
  "arguments": {
    "parameters": [
      {
        "boot-file-name": "/dev/null",
        "metadata": {
          "server-tags": [ "server1" ]
        }
      },
      {
        "match-client-id": true,
        "metadata": {
          "server-tags": [ "all" ]
        }
      }
    ],
    "count": 2
  }
}

```

The example response contains two parameters: one string parameter and one boolean parameter. The metadata re-

turned for each parameter indicates whether this parameter is specific to "server1" or applies to all servers. Since the `match-client-id` value is associated with "all" servers, it indicates that there is no "server1"-specific setting for this parameter. Each parameter always has exactly one server tag associated with it, because global parameters are non-shareable configuration elements.

Note

If the server tag is set to "all" in the command, the response will contain only the global parameters associated with the logical server "all". When the server tag points to the specific server (as in the example above), the returned list combines parameters associated with this server and all servers, but the former take precedence.

17.6.11 The `remote-global-parameter4-set`, `remote-global-parameter6-set` Commands

This command is used to create scalar global DHCP parameters in the database. If any of the parameters already exists, its value is replaced as a result of this command. It is possible to set multiple parameters within a single command, each having one of the four types: string, integer, real, or boolean. For example:

```
{
  "command": "remote-global-parameter4-set",
  "arguments": {
    "parameters": {
      "boot-file-name": "/dev/null",
      "renew-timer": 2000,
      "t1-percent": 0.85,
      "match-client-id": true
    },
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}
```

An error is returned if any of the parameters is not supported by the DHCP server or its type does not match. Care should be taken when multiple parameters are specified in a single command, because it is possible that only some of the parameters will be stored successfully and some will fail. If an error occurs when processing this command, it is recommended to use `remote-global-parameter4-get-all` or `remote-global-parameter6-get-all` to check which of the parameters have been stored/updated successfully and which have failed.

The `server-tags` list is mandatory and must contain a single server tag or the keyword "all". In the example above, all specified parameters are associated with the "server1" server.

17.6.12 The `remote-network4-del`, `remote-network6-del` Commands

These commands are used to delete an IPv4 or IPv6 shared network from the database. The optional parameter `subnets-action` determines whether the subnets belonging to the deleted shared network should also be deleted or preserved. The `subnets-action` parameter defaults to `keep`, which preserves the subnets. If it is set to `delete`, the subnets are deleted along with the shared network.

The following command:

```
{
  "command": "remote-network6-del",
```

(continues on next page)

(continued from previous page)

```

    "arguments": {
      "shared-networks": [
        {
          "name": "level3"
        }
      ],
      "subnets-action": "keep",
      "remote": {
        "type": "mysql"
      }
    }
  }
}

```

deletes the "level3" IPv6 shared network. The subnets are preserved, but they are disassociated from the deleted shared network and become global. This behavior corresponds to the behavior of the [network4-del](#), [network6-del](#) commands with respect to the `subnets-action` parameter.

Note that the `server-tags` parameter cannot be used for this command.

17.6.13 The `remote-network4-get`, `remote-network6-get` Commands

These commands are used to retrieve information about an IPv4 or IPv6 shared network. The optional parameter `subnets-include` denotes whether the subnets belonging to the shared network should also be returned. This parameter defaults to `no`, in which case the subnets are not returned. If this parameter is set to `full`, the subnets are returned together with the shared network.

The following command fetches the "level3" IPv6 shared network along with the full information about the subnets belonging to it:

```

{
  "command": "remote-network6-get",
  "arguments": {
    "shared-networks": [
      {
        "name": "level3"
      }
    ],
    "subnets-include": "full",
    "remote": {
      "type": "mysql"
    }
  }
}

```

Note that the `server-tags` parameter cannot be used for this command.

17.6.14 The `remote-network4-list`, `remote-network6-list` Commands

These commands are used to list all IPv4 or IPv6 shared networks for a server.

The following command retrieves all shared networks to be used by "server1" and "server2":

```

{
  "command": "remote-network4-list",

```

(continues on next page)

(continued from previous page)

```

"arguments": {
  "remote": {
    "type": "mysql"
  },
  "server-tags": [ "server1", "server2" ]
}

```

The `server-tags` parameter is mandatory and contains one or more server tags. It may contain the keyword "all" to fetch the shared networks associated with all servers. When the `server-tags` list contains the `null` value, the returned response contains a list of unassigned shared networks, i.e. the networks which are associated with no servers. For example:

```

{
  "command": "remote-network4-list",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ null ]
  }
}

```

The example response to this command when non-null server tags are specified looks similar to this:

```

{
  "result": 0,
  "text": "3 IPv4 shared network(s) found.",
  "arguments": {
    "shared-networks": [
      {
        "name": "ground floor",
        "metadata": {
          "server-tags": [ "all" ]
        }
      },
      {
        "name": "floor2",
        "metadata": {
          "server-tags": [ "server1" ]
        }
      },
      {
        "name": "floor3",
        "metadata": {
          "server-tags": [ "server2" ]
        }
      }
    ],
    "count": 3
  }
}

```

The returned information about each shared network merely contains the shared network name and the meta-data. To fetch detailed information about the selected shared network, use the `remote-network4-get` or `remote-network6-get` command.

The example response above contains three shared networks. One of the shared networks is associated with all servers, so it is included in the list of shared networks to be used by "server1" and "server2". The remaining two shared networks are returned because one of them is associated with "server1" and another one is associated with "server2".

When listing unassigned shared networks, the response looks similar to this:

```
{
  "result": 0,
  "text": "1 IPv4 shared network(s) found.",
  "arguments": {
    "shared-networks": [
      {
        "name": "fancy",
        "metadata": {
          "server-tags": [ null ]
        }
      }
    ],
    "count": 1
  }
}
```

The null value in the metadata indicates that the returned shared network is unassigned.

17.6.15 The `remote-network4-set`, `remote-network6-set` Commands

These commands create a new or replace an existing IPv4 or IPv6 shared network in the database. The structure of the shared network information is the same as in the Kea configuration file (see *Shared Networks in DHCPv4* and *Shared Networks in DHCPv6* for details), except that specifying subnets along with the shared network information is not allowed. Including the `subnet4` or `subnet6` parameter within the shared network information results in an error.

These commands are intended to be used for managing the shared network-specific information and DHCP options. To associate and disassociate the subnets with the shared networks, the `remote-subnet4-set`, `remote-subnet6-set` commands should be used.

The following command adds the IPv6 shared network "level3" to the database:

```
{
  "command": "remote-network6-set",
  "arguments": {
    "shared-networks": [
      {
        "name": "level3",
        "interface": "eth0",
        "option-data": [ {
          "name": "sntp-servers",
          "data": "2001:db8:1::1"
        } ]
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    "server-tags": [ "all" ]
  }
}

```

This command includes the `interface` parameter, which sets the shared network-level interface name. Any remaining shared-network-level parameters, which are not specified with the command, will be marked as "unspecified" in the database. The DHCP server uses the global values for unspecified parameters or, if the global values are not specified, the default values are used.

The `server-tags` list is mandatory for this command and must include one or more server tags. As a result, the shared network is associated with all listed servers. The shared network may be associated with all servers connecting to the database when the keyword "all" is included.

Note

As with other "set" commands, this command replaces all the information about the given shared network in the database, if the shared network already exists. Therefore, when sending this command, make sure to always include all parameters that must be specified for the updated shared-network instance. Any unspecified parameter will be marked unspecified in the database, even if its value was present prior to sending the command.

17.6.16 The `remote-option-def4-del`, `remote-option-def6-del` Commands

These commands are used to delete a DHCP option definition from the database. The option definition is identified by an option code and option space. For example:

```

{
  "command": "remote-option-def6-del",
  "arguments": {
    "option-defs": [
      {
        "code": 1,
        "space": "isc"
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}

```

deletes the definition of the option associated with "server1", having the code of 1 and belonging to the option space "isc". The default option spaces are "dhcp4" and "dhcp6" for the DHCPv4 and DHCPv6 top-level options, respectively. If there is no such option explicitly associated with "server1", no option is deleted. To delete an option belonging to "all" servers, the keyword "all" must be used as the server tag. The `server-tags` list must contain exactly one tag and cannot include the null value.

17.6.17 The `remote-option-def4-get`, `remote-option-def6-get` Commands

These commands are used to fetch a specified DHCP option definition from the database. The option definition is identified by the option code and option space. The default option spaces are "dhcp4" and "dhcp6" for the DHCPv4 and DHCPv6 top-level options, respectively.

The following command retrieves a DHCPv4 option definition associated with all servers, having the code of 1 and belonging to the option space "isc":

```
{
  "command": "remote-option-def4-get",
  "arguments": {
    "option-defs": [
      {
        "code": 1,
        "space": "isc"
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

The `server-tags` list must include exactly one server tag or the keyword "all", and cannot contain the *null* value.

17.6.18 The `remote-option-def4-get-all`, `remote-option-def6-get-all` Commands

These commands are used to fetch all DHCP option definitions from the database for the given server or all servers. For example:

```
{
  "command": "remote-option-def6-get-all",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

This command attempts to fetch all DHCPv6 option definitions associated with "all" servers. The `server-tags` list is mandatory for this command and must include exactly one server tag or the keyword "all". It cannot include the *null* value.

The following is the example response to this command:

```
{
  "result": 0,
  "text": "1 DHCPv6 option definition(s) found.",
  "arguments": {
    "option-defs": [
      {
```

(continues on next page)

(continued from previous page)

```

        "name": "bar",
        "code": 1012,
        "space": "dhcp6",
        "type": "record",
        "array": true,
        "record-types": "ipv6-address, uint16",
        "encapsulate": "",
        "metadata": {
            "server-tags": [ "all" ]
        }
    },
    "count": 1
}

```

The response contains an option definition associated with all servers, as indicated by the metadata.

17.6.19 The `remote-option-def4-set`, `remote-option-def6-set` Commands

These commands create a new DHCP option definition or replace an existing option definition in the database. The structure of the option definition information is the same as in the Kea configuration file (see *Custom DHCPv4 Options* and *Custom DHCPv6 Options*). The following command creates the DHCPv4 option definition at the top-level "dhcp4" option space and associates it with "server1":

```

{
    "command": "remote-option-def4-set",
    "arguments": {
        "option-defs": [
            {
                "name": "foo",
                "code": 222,
                "type": "uint32",
                "array": false,
                "record-types": "",
                "space": "dhcp4",
                "encapsulate": ""
            }
        ],
        "remote": {
            "type": "mysql"
        },
        "server-tags": [ "server1" ]
    }
}

```

The `server-tags` list must include exactly one server tag or the keyword "all", and cannot contain the `null` value.

17.6.20 The `remote-option4-global-del`, `remote-option6-global-del` Commands

These commands are used to delete a global DHCP option from the database. The option is identified by an option code and option space. For example:

```
{
  "command": "remote-option4-global-del",
  "arguments": {
    "options": [
      {
        "code": 5,
        "space": "dhcp4"
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}
```

The "dhcp4" value represents the top-level option space where the standard DHCPv4 options belong. The `server-tags` parameter is mandatory and must include a single option tag or the keyword "all". If the explicit server tag is specified, this command attempts to delete a global option associated with this server. If there is no such option associated with the given server, no option is deleted. To delete an option associated with all servers, the keyword "all" must be specified.

17.6.21 The `remote-option4-global-get`, `remote-option6-global-get` Commands

These commands are used to fetch a global DHCP option from the database. The option is identified by the code and option space. The top-level option spaces where DHCP standard options belong are called "dhcp4" and "dhcp6" for the DHCPv4 and DHCPv6 servers, respectively.

The following command retrieves the IPv6 "DNS Servers" (code 23) option associated with all servers:

```
{
  "command": "remote-option6-global-get",
  "arguments": {
    "options": [
      {
        "code": 23,
        "space": "dhcp6"
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

The `server-tags` parameter is mandatory and must include exactly one server tag or the keyword "all". It cannot contain the null value.

17.6.22 The `remote-option4-global-get-all`, `remote-option6-global-get-all` Commands

These commands are used to fetch all global DHCP options from the configuration database for the given server or for all servers. The following command fetches all global DHCPv4 options for "server1":

```
{
  "command": "remote-option6-global-get-all",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}
```

The `server-tags` list is mandatory for this command and must contain exactly one server tag or a keyword "all"; it cannot contain the null value.

The following is a example response to this command with a single option being associated with "server1" returned:

```
{
  "result": 0,
  "text": "DHCPv4 options found.",
  "arguments": {
    "options": [
      {
        "name": "domain-name-servers",
        "code": 6,
        "space": "dhcp4",
        "csv-format": false,
        "data": "192.0.2.3",
        "metadata": {
          "server-tags": [ "server1" ]
        }
      }
    ],
    "count": 1
  }
}
```

17.6.23 The `remote-option4-global-set`, `remote-option6-global-set` Commands

These commands create a new global DHCP option or replace an existing option in the database. The structure of the option information is the same as in the Kea configuration file (see *Standard DHCPv4 Options* and *Standard DHCPv6 Options*). For example:

```
{
  "command": "remote-option6-global-set",
  "arguments": {
    "options": [
      {
        "name": "dns-servers",
        "data": "2001:db8:1::1"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "remote": {
    "type": "mysql"
  },
  "server-tags": [ "server1" ]
}

```

The `server-tags` list is mandatory for this command and must include exactly one server tag or the keyword "all"; it cannot include the null value. The command above associates the option with the "server1" server.

Note that specifying an option name instead of the option code only works reliably for standard DHCP options. When specifying a value for a user-defined DHCP option, the option code should be indicated instead of the name. For example:

```

{
  "command": "remote-option6-global-set",
  "arguments": {
    "options": [
      {
        "code": 1,
        "space": "isc",
        "data": "2001:db8:1::1"
      }
    ],
    "server-tags": [ "server1" ]
  }
}

```

17.6.24 The `remote-option4-network-del`, `remote-option6-network-del` Commands

These commands are used to delete a shared-network-specific DHCP option from the database. The option is identified by an option code and option space and these two parameters are passed within the `options` list. Another list, `shared-networks`, contains a map with the name of the shared network from which the option is to be deleted. If the option is not explicitly specified for this shared network, no option is deleted. In particular, the given option may be present for a subnet belonging to the shared network. Such an option instance is not affected by this command as this command merely deletes the shared-network-level option. To delete a subnet-level option, the `remote-option4-subnet-del`, `remote-option6-subnet-del` commands must be used instead.

The following command attempts to delete an option having the option code 5 in the top-level option space from the shared network "fancy".

```

{
  "command": "remote-option4-network-del",
  "arguments": {
    "shared-networks": [
      {
        "name": "fancy"
      }
    ],
    "options": [

```

(continues on next page)

(continued from previous page)

```

    {
      "code": 5,
      "space": "dhcp4"
    }
  ],
  "remote": {
    "type": "mysql"
  }
}

```

The "dhcp4" value represents the top-level option space where the standard DHCPv4 options belong. The `server-tags` parameter cannot be specified for this command.

17.6.25 The `remote-option4-network-set`, `remote-option6-network-set` Commands

These commands create a new shared-network-specific DHCP option or replace an existing option in the database. The structure of the option information is the same as in the Kea configuration file (see *Standard DHCPv4 Options* and *Standard DHCPv6 Options*). The option information is carried in the `options` list. Another list, `shared-networks`, contains a map with the name of the shared network for which the option is to be set. If such an option already exists for the shared network, it is replaced with the new instance.

```

{
  "command": "remote-option6-network-set",
  "arguments": {
    "shared-networks": [
      {
        "name": "fancy"
      }
    ],
    "options": [
      {
        "name": "dns-servers",
        "data": "2001:db8:1::1"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}

```

The `server-tags` parameter cannot be specified for this command.

Specifying an option name instead of the option code only works reliably for standard DHCP options. When specifying a value for a user-defined DHCP option, the option code should be indicated instead of the name.

17.6.26 The `remote-option6-pd-pool-del` Command

This command is used to delete a prefix delegation pool-specific DHCPv6 option from the database. The option is identified by an option code and option space, and these two parameters are passed within the `options` list. Another list, `pd-pools`, contains a map with the prefix-delegation-pool prefix and length identifying the pool. If the option is not explicitly specified for this pool, no option is deleted. In particular, the given option may exist for a subnet containing

the specified pool. Such an option instance is not affected by this command, as this command merely deletes a prefix delegation pool-level option. To delete a subnet-level option, the `remote-option6-subnet-del` command must be used instead.

```
{
  "command": "remote-option6-pd-pool-del",
  "arguments": {
    "pd-pools": [
      {
        "prefix": "3000::",
        "prefix-len": 64
      }
    ],
    "options": [
      {
        "code": 23,
        "space": "dhcp6"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The "dhcp6" value represents the top-level option space where the standard DHCPv6 options belong. The `server-tags` parameter cannot be specified for this command.

17.6.27 The `remote-option6-pd-pool-set` Command

This command creates a new prefix delegation pool-specific DHCPv6 option or replaces an existing option in the database. The structure of the option information is the same as in the Kea configuration file (see *Standard DHCPv4 Options* and *Standard DHCPv6 Options*). The option information is carried in the `options` list. Another list, `pd-pools`, contains a map with the prefix-delegation-pool prefix and the prefix length identifying the pool. If such an option already exists for the prefix delegation pool, it is replaced with the new instance.

For example:

```
{
  "command": "remote-option6-pd-pool-set",
  "arguments": {
    "pd-pools": [
      {
        "prefix": "3001:1::",
        "length": 64
      }
    ],
    "options": [
      {
        "name": "dns-servers",
        "data": "2001:db8:1::1"
      }
    ],
    "remote": {
```

(continues on next page)

(continued from previous page)

```

        "type": "mysql"
    }
}

```

The `server-tags` parameter cannot be specified for this command.

Specifying an option name instead of the option code only works reliably for standard DHCP options. When specifying a value for a user-defined DHCP option, the option code should be indicated instead of the name.

17.6.28 The `remote-option4-pool-del`, `remote-option6-pool-del` Commands

These commands are used to delete an address-pool-specific DHCP option from the database. The option is identified by an option code and option space, and these two parameters are passed within the `options` list. Another list, `pools`, contains a map with the IP address range or prefix identifying the pool. If the option is not explicitly specified for this pool, no option is deleted. In particular, the given option may exist for a subnet containing the specified pool. Such an option instance is not affected by this command, as this command merely deletes a pool-level option. To delete a subnet-level option, the `remote-option4-subnet-del` or `remote-option6-subnet-del` command must be used instead.

The following command attempts to delete an option having the option code 5 in the top-level option space from an IPv4 address pool:

```

{
  "command": "remote-option4-pool-del",
  "arguments": {
    "pools": [
      {
        "pool": "192.0.2.10 - 192.0.2.100"
      }
    ],
    "options": [
      {
        "code": 5,
        "space": "dhcp4"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}

```

The "dhcp4" value represents the top-level option space where the standard DHCPv4 options belong. The `server-tags` parameter cannot be specified for this command.

17.6.29 The `remote-option4-pool-set`, `remote-option6-pool-set` Commands

These commands create a new address-pool-specific DHCP option or replace an existing option in the database. The structure of the option information is the same as in the Kea configuration file (see [Standard DHCPv4 Options](#) and [Standard DHCPv6 Options](#)). The option information is carried in the `options` list. Another list, `pools`, contains a map with the IP address range or prefix identifying the pool. If such an option already exists for the pool, it is replaced with the new instance.

For example:

```
{
  "command": "remote-option4-pool-set",
  "arguments": {
    "pools": [
      {
        "pool": "192.0.2.10 - 192.0.2.100"
      }
    ],
    "options": [
      {
        "name": "domain-name-servers",
        "data": "10.0.0.1"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be specified for this command.

Specifying an option name instead of the option code only works reliably for standard DHCP options. When specifying a value for a user-defined DHCP option, the option code should be indicated instead of the name.

17.6.30 The `remote-option4-subnet-del`, `remote-option6-subnet-del` Commands

These commands are used to delete a subnet-specific DHCP option from the database. The option is identified by an option code and option space, and these two parameters are passed within the `options` list. Another list, `subnets`, contains a map with the identifier of the subnet from which the option is to be deleted. If the option is not explicitly specified for this subnet, no option is deleted.

The following command attempts to delete an option having the option code 5 in the top-level option space from the subnet having an identifier of 123.

```
{
  "command": "remote-option4-subnet-del",
  "arguments": {
    "subnets": [
      {
        "id": 123
      }
    ],
    "options": [
      {
        "code": 5,
        "space": "dhcp4"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The "dhcp4" value represents the top-level option space where the standard DHCPv4 options belong. The `server-tags` parameter cannot be specified for this command.

17.6.31 The `remote-option4-subnet-set`, `remote-option6-subnet-set` Commands

These commands create a new subnet-specific DHCP option or replace an existing option in the database. The structure of the option information is the same as in the Kea configuration file (see *Standard DHCPv4 Options* and *Standard DHCPv6 Options*). The option information is carried in the `options` list. Another list, `subnets`, contains a map with the identifier of the subnet for which the option is to be set. If such an option already exists for the subnet, it is replaced with the new instance.

```
{
  "command": "remote-option6-subnet-set",
  "arguments": {
    "subnets": [
      {
        "id": 123
      }
    ],
    "options": [
      {
        "name": "dns-servers",
        "data": "2001:db8:1::1"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be specified for this command.

Specifying an option name instead of the option code only works reliably for the standard DHCP options. When specifying a value for the user-defined DHCP option, the option code should be indicated instead of the name.

17.6.32 The `remote-subnet4-del-by-id`, `remote-subnet6-del-by-id` Commands

This is the first variant of the commands used to delete an IPv4 or IPv6 subnet from the database. It uses the subnet ID to identify the subnet. For example, to delete the IPv4 subnet with an ID of 5:

```
{
  "command": "remote-subnet4-del-by-id",
  "arguments": {
    "subnets": [
      {
        "id": 5
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```


The `server-tags` parameter cannot be used with this command.

17.6.33 The `remote-subnet4-del-by-prefix`, `remote-subnet6-del-by-prefix` Commands

This is the second variant of the commands used to delete an IPv4 or IPv6 subnet from the database. It uses the subnet prefix to identify the subnet. For example:

```
{
  "command": "remote-subnet6-del-by-prefix",
  "arguments": {
    "subnets": [
      {
        "subnet": "2001:db8:1::/64"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be used with this command.

17.6.34 The `remote-subnet4-get-by-id`, `remote-subnet6-get-by-id` Commands

This is the first variant of the commands used to fetch an IPv4 or IPv6 subnet from the database. It uses a subnet ID to identify the subnet. For example:

```
{
  "command": "remote-subnet4-get-by-id",
  "arguments": {
    "subnets": [
      {
        "id": 5
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be used with this command.

17.6.35 The `remote-subnet4-get-by-prefix`, `remote-subnet6-get-by-prefix` Commands

This is the second variant of the commands used to fetch an IPv4 or IPv6 subnet from the database. It uses a subnet prefix to identify the subnet. For example:

```
{
  "command": "remote-subnet6-get-by-prefix",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "subnets": [
      {
        "subnet": "2001:db8:1::/64"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}

```

The `server-tags` parameter cannot be used with this command.

17.6.36 The `remote-subnet4-list`, `remote-subnet6-list` Commands

These commands are used to list all IPv4 or IPv6 subnets from the database for selected servers or all servers. The following command retrieves all servers to be used by "server1" and "server2":

```

{
  "command": "remote-subnet4-list",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1", "server2" ]
  }
}

```

The `server-tags` parameter is mandatory and contains one or more server tags. It may contain the keyword "all", to fetch the subnets associated with all servers. When the `server-tags` list contains the `null` value, the returned response contains a list of unassigned subnets, i.e. the subnets which are associated with no servers. For example:

```

{
  "command": "remote-subnet4-list",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ null ]
  }
}

```

The example response to this command when non-null server tags are specified looks similar to this:

```

{
  "result": 0,
  "text": "2 IPv4 subnet(s) found.",
  "arguments": {
    "subnets": [
      {
        "id": 1,
        "subnet": "192.0.2.0/24",
        "shared-network-name": null,

```

(continues on next page)

(continued from previous page)

```

        "metadata": {
            "server-tags": [ "server1", "server2" ]
        },
        {
            "id": 2,
            "subnet": "192.0.3.0/24",
            "shared-network-name": null,
            "metadata": {
                "server-tags": [ "all" ]
            }
        }
    ],
    "count": 2
}

```

The returned information about each subnet is limited to the subnet identifier, prefix, and associated shared network name. To retrieve full information about the selected subnet, use the *remote-subnet4-get-by-id* / *remote-subnet6-get-by-id* command or the *remote-subnet4-get-by-prefix* / *remote-subnet6-get-by-prefix* command.

The example response above contains two subnets. One of the subnets is associated with both servers: "server1" and "server2". The second subnet is associated with all servers, so it is also present in the configurations for "server1" and "server2".

When listing unassigned subnets, the response will look similar to this:

```

{
    "result": 0,
    "text": "1 IPv4 subnet(s) found.",
    "arguments": {
        "subnets": [
            {
                "id": 3,
                "subnet": "192.0.4.0/24",
                "shared-network-name": null,
                "metadata": {
                    "server-tags": [ null ]
                }
            }
        ],
        "count": 1
    }
}

```

The null value in the metadata indicates that the returned subnet is unassigned.

17.6.37 The *remote-subnet4-set*, *remote-subnet6-set* Commands

These commands are used to create a new IPv4 or IPv6 subnet or replace an existing subnet in the database. Setting the subnet also associates or disassociates the subnet with a shared network.

The structure of the subnet information is similar to the structure used in the configuration file (see *DHCPv4 Server Configuration* and *DHCPv6 Server Configuration*). The subnet information conveyed in the *remote-subnet4-set*

or `remote-subnet6-set` command must include the additional parameter `shared-network-name`, which denotes whether the subnet belongs to a shared network.

Consider the following example:

```
{
  "command": "remote-subnet4-set",
  "arguments": {
    "subnets": [
      {
        "id": 5,
        "subnet": "192.0.2.0/24",
        "shared-network-name": "level3",
        "pools": [ { "pool": "192.0.2.100-192.0.2.200" } ],
        "option-data": [ {
          "name": "routers",
          "data": "192.0.2.1"
        } ]
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

It creates the subnet and associates it with the "level3" shared network. The "level3" shared network must be created with the `remote-network4-set` command prior to creating the subnet.

If the created subnet must be global - that is, not associated with any shared network - the `shared-network-name` must be explicitly set to `null`:

```
{
  "command": "remote-subnet4-set",
  "arguments": {
    "subnets": [
      {
        "id": 5,
        "subnet": "192.0.2.0/24",
        "shared-network-name": null,
        "pools": [ { "pool": "192.0.2.100-192.0.2.200" } ],
        "option-data": [ {
          "name": "routers",
          "data": "192.0.2.1"
        } ]
      }
    ],
    "server-tags": [ "all" ]
  }
}
```

The subnet created in the previous example is replaced with the new subnet having the same parameters, but it becomes global.

The `shared-network-name` parameter is mandatory for the `remote-subnet4-set` command. The `server-tags`

list is mandatory and must include one or more server tags. As a result, the subnet is associated with all of the listed servers. It may also be associated with all servers connecting to the database when the keyword "all" is used as the server tag.

Note

As with other "set" commands, this command replaces all the information about the particular subnet in the database, if the subnet information is already present. Therefore, when sending this command, make sure to always include all parameters that must be specified for the updated subnet instance. Any unspecified parameter will be marked as unspecified in the database, even if its value was present prior to sending the command.

17.6.38 The `remote-class4-del`, `remote-class6-del` Commands

These commands delete a DHCPv4 or DHCPv6 client class by name. If any client classes in the database depend on the deleted class, an error is returned in response to this command. In this case, to successfully delete the class, the dependent client classes must be deleted first. Use the `remote-class4-get-all` command to fetch all client classes and find the dependent ones.

```
{
  "command": "remote-class4-del",
  "arguments": {
    "client-classes": [
      {
        "name": "foo"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be used for this command because client classes are uniquely identified by name.

17.6.39 The `remote-class4-get`, `remote-class6-get` Commands

These commands retrieve DHCPv4 or DHCPv6 client class information by a client-class name.

```
{
  "command": "remote-class4-get",
  "arguments": {
    "client-classes": [
      {
        "name": "foo"
      }
    ],
    "remote": {
      "type": "mysql"
    }
  }
}
```

The `server-tags` parameter cannot be used for this command because client classes are uniquely identified by name.

A response to the command looks similar to this:

```
{
  "result": 0,
  "text": "DHCPv4 client class 'foo' found.",
  "arguments": {
    "client-classes": [
      {
        "name": "foo",
        "metadata": {
          "server-tags": [ "all" ]
        }
      }
    ],
    "count": 1
  }
}
```

17.6.40 The remote-class4-get-all, remote-class6-get-all Commands

These commands retrieve all DHCPv4 or DHCPv6 client classes for a particular server, multiple explicitly listed servers, and/or all servers. A given server has its own server-specific tag and also has the "all" server tag; these commands retrieve the classes for both an individual server and for "all" servers. For example, the following command retrieves all client classes defined for "server1" as well as the client classes defined for "all" servers:

```
{
  "command": "remote-class4-get-all",
  "arguments": {
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "server1" ]
  }
}
```

The `server-tags` parameter is mandatory and contains one or more server tags. If other server tags are specified, "all" does not need to be included in `server-tags`, as every server automatically also has the "all" server tag. If `server-tags` contains only the keyword "all", only the client classes associated with "all" servers are returned. When the `server-tags` list contains the null value, the returned response contains a list of unassigned client classes, i.e. the networks which are associated with no servers.

A response to the command looks similar to this:

```
{
  "result": 0,
  "text": "2 DHCPv4 client class(es) found.",
  "arguments": {
    "client-classes": [
      {
        "name": "foo",
        "metadata": {
          "server-tags": [ "all" ]
        }
      }
    ],
  },
}
```

(continues on next page)

(continued from previous page)

```

    {
      "name": "bar",
      "test": "member('foo')",
      "metadata": {
        "server-tags": [ "all" ]
      }
    }
  ],
  "count": 2
}

```

17.6.41 The remote-class4-set, remote-class6-set Commands

These commands insert a new or replace an existing DHCPv4 or DHCPv6 client class in the database. The client class information structure is the same as in the Kea configuration file (see *Client Classification in DHCPv4* and *Client Classification in DHCPv6* for details).

```

{
  "command": "remote-class4-set",
  "arguments": {
    "client-classes": [
      {
        "name": "foo",
        "test": "member('KNOWN') or member('bar')",
        "option-def": [
          {
            "name": "configfile",
            "code": 224,
            "type": "string"
          }
        ],
        "option-data": [
          {
            "name": "configfile",
            "data": "1APC"
          }
        ]
      }
    ],
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}

```

Client-class ordering rules described in *Using Expressions in Classification* apply to the classes inserted into the database. They imply that the class *bar* referenced in the test expression must exist in the database when issuing the above command.

By default, a new client class is inserted at the end of the class hierarchy in the database and can reference any class associated with the same server tag or with the special server tag "all". If an existing class is updated, it remains at its

current position within the class hierarchy.

However, the class commands allow the position of the inserted or updated client class to be specified. The optional `follow-class-name` parameter can be included in the command to indicate the name of the existing class after which the managed class should be placed. Suppose there are two DHCPv6 classes in the database: *first-class* and *second-class*. To add a new class, *third-class*, between these two, use a command similar to the following:

```
{
  "command": "remote-class6-set",
  "arguments": {
    "client-classes": [
      {
        "name": "third-class",
        "test": "member('first-class')"
      }
    ],
    "follow-class-name": "first-class",
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

Note that *third-class* can depend on *first-class* because it is placed after *first-class*; *third-class* cannot depend on *second-class* because it is placed before it. However, *second-class* could be updated to depend on *third-class*.

The `follow-class-name` parameter can be explicitly set to `null`, e.g.:

```
{
  "command": "remote-class6-set",
  "arguments": {
    "client-classes": [
      {
        "name": "third-class",
        "test": "member('first-class')"
      }
    ],
    "follow-class-name": null,
    "remote": {
      "type": "mysql"
    },
    "server-tags": [ "all" ]
  }
}
```

It yields the same behavior as if the `follow-class-name` parameter were not included, i.e. the new class is appended at the end of the class hierarchy, and the updated class remains at the current position.

17.7 libdhcp_class_cmds.so: Class Commands

This hook library exposes several control commands for manipulating client classes (part of the Kea DHCP servers' configurations) without the need to restart those servers. Using these commands it is possible to add, update, delete, and list the client classes configured for a given server.

Note

`libdhcp_class_cmds.so` is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the `kea-dhcp4` or `kea-dhcp6` process.

17.7.1 The `class-add` Command

The `class-add` command adds a new client class to the DHCP server configuration. This class is appended at the end of the list of classes used by the server and may depend on any of the already-configured client classes.

The following example demonstrates how to add a new client class to the DHCPv4 server configuration:

```
{
  "command": "class-add",
  "arguments": {
    "client-classes": [
      {
        "name": "ipxe_efi_x64",
        "test": "option[93].hex == 0x0009",
        "next-server": "192.0.2.254",
        "server-hostname": "hal9000",
        "boot-file-name": "/dev/null"
      }
    ]
  }
}
```

Note that the `client-classes` parameter is a JSON list, but it allows only a single client class to be present.

Here is the response to the `class-add` command in our example:

```
{
  "result": 0,
  "text": "Class 'ipxe_efi_x64' added."
}
```

17.7.2 The `class-update` Command

The `class-update` command updates an existing client class in the DHCP server configuration. If the client class with the given name does not exist, the server returns the result code of 3, which means that the server configuration is not modified and the client class does not exist. The `class-add` command must be used instead to create the new client class.

The `class-update` command has the same argument structure as the `class-add` command:

```
{
  "command": "class-update",
  "arguments": {
    "client-classes": [
```

(continues on next page)

(continued from previous page)

```

    {
        "name": "ipxe_efi_x64",
        "test": "option[93].hex == 0x0017",
        "next-server": "0.0.0.0",
        "server-hostname": "xfce",
        "boot-file-name": "/dev/null"
    }
]
}

```

Here is the response for our example:

```

{
    "result": 0,
    "text": "Class 'ipxe_efi_x64' updated."
}

```

Any parameter of the client class can be modified with this command, except `name`. There is currently no way to rename the class, because the class name is used as a key for searching the class to be updated. To achieve a similar effect to renaming the class, an existing class can be removed with the `class-del` command and then added again with a different name using `class-add`. Note, however, that the class with the new name will be added at the end of the list of configured classes.

As with other update commands, this command overwrites all the contents of an entry. If the client class previously had a resource assigned to it, and the `class-update` command is missing the resource, it is deleted from the server configuration. If an incremental update of the class is desired, it can be achieved by issuing a `class-get` to get the current state of the client class, selecting the client class from the response, modifying it to the required outcome, and then issuing the `client-update` command with the resulting client class attached.

17.7.3 The `class-del` Command

The `class-del` command is used to remove a particular class from the server configuration. The class to be removed is identified by name. The class is not removed if there are other classes depending on it; to remove such a class, the dependent classes must be removed first.

The following is a sample command removing the `ipxe_efi_x64` class:

```

{
    "command": "class-del",
    "arguments": {
        "name": "ipxe_efi_x64"
    }
}

```

Here is the response to the `class-del` command in our example, when the specified client class has been found:

```

{
    "result": 0,
    "text": "Class 'ipxe_efi_x64' deleted."
}

```

If the class does not exist, the result of 3 is returned.

17.7.4 The class-list Command

class-list is used to retrieve a list of all client classes. This command includes no arguments:

```
{
  "command": "class-list"
}
```

Here is the response of the server in our example, including the list of client classes:

```
{
  "result": 0,
  "text": "2 classes found",
  "arguments": {
    "client-classes": [
      {
        "name": "ipxe_efi_x64"
      },
      {
        "name": "pxeclient"
      }
    ]
  }
}
```

Note that the returned list does not contain full class definitions, but merely class names. To retrieve full class information, the *class-get* command should be used.

17.7.5 The class-get Command

class-get is used to retrieve detailed information about a specified class. The command structure is very simple:

```
{
  "command": "class-get",
  "arguments": {
    "name": "pxeclient"
  }
}
```

If the class with the specified name does not exist, the status code of 3 is returned. If the specified client class exists, the class details are returned in the following format:

```
{
  "result": 0,
  "text": "Class 'pxeclient' definition returned",
  "arguments": {
    "client-classes": [
      {
        "name": "pxeclient",
        "only-in-additional-list": true,
        "test": "option[vendor-class-identifier].text == 'PXEClient'",
        "option-def": [
          {
            "name": "configfile",
            "code": 209,

```

(continues on next page)

(continued from previous page)

```

        "type": "string"
      }
    ],
    "option-data": [ ],
    "next-server": "0.0.0.0",
    "server-hostname": "xfce",
    "boot-file-name": "/dev/null"
  }
]
}
}

```

Note that the example above is DHCPv4-specific; the last three parameters are only returned by the DHCPv4 server and are never returned by the DHCPv6 server. Also, some of the parameters provided in this example may not be returned if they are not specified for the class. Specifically, `only-in-additional-list`, `test`, and `option-def` are not returned if they are not specified for the class.

Note

As of Kea version 2.7.4, `only-if-required` has been replaced with `only-in-additional-list` and deprecated. In order to allow users time to migrate class commands will still accept it as input but translate it to `only-in-additional-list` on output. Eventually support for the old name will be removed.

17.8 libdhcp_ddns_tuning.so: DDNS Tuning

This hook library adds support for fine-tuning various DNS update aspects. It currently supports procedural host-name generation and the ability to skip performing DDNS updates for select clients.

Note

libdhcp_ddns_tuning.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

The library, which was added in Kea 2.1.5, can be loaded by the *kea-dhcp4* and *kea-dhcp6* daemons by adding it to the `hooks-libraries` element of the server's configuration:

```

{
  "hooks-libraries": [
    {
      "library": "/usr/local/lib/libdhcp_ddns_tuning.so",
      "parameters": {
        ...
      }
    },
    ...
  ],
  ...
}

```

17.8.1 Procedural Host-Name Generation

This hook library provides the ability to generate host names procedurally, based on an expression. The expression can be defined globally in the hook parameters, using *hostname-expr*. If defined globally, it applies to all hosts in all subnets. The expressions can use all tokens defined in *Client Classification*. An example of a global expression is shown below:

```
{
  "hooks-libraries": [
    {
      "library": "/usr/local/lib/libdhcp_ddns_tuning.so",
      "parameters": {
        "hostname-expr": "'host-' + hexstring(pkt4.mac, '-')",
        ...
      }
    },
    ...
  ],
  ...
}
```

The default behavior of *kea-dhcp4* is to prefer the FQDN option (code 81) over the host name option (code 12) when a client sends both. The following example shows the *hostname-expr* one would use to reverse this rule:

```
{
  "hooks-libraries": [
    {
      "library": "/usr/local/lib/libdhcp_ddns_tuning.so",
      "parameters": {
        "hostname-expr" : "ifelse(option[12].exists, option[host-name].text, ↵
↵option[81].text)"
      }
    }
  ]
}
```

It is also possible to define this parameter in a subnet, using the user-context mechanism. If defined at the subnet level, the expression applies to a specific subnet only. If the subnet expression is defined as empty, "", it suppresses (or disables) the use of a global expression for that subnet. An example subnet expression is shown below:

```
{
  "subnet4": [{
    "subnet": "192.0.2.0/24",
    "pools": [{
      "pool": "192.0.2.10 - 192.0.2.20"
    } ],
    "user-context": {
      "ddns-tuning": {
        "hostname-expr": "'guest-' + int8totext(substring(pkt4.yiaddr, 0,1)) + '-' \
+ int8totext(substring(pkt4.yiaddr, 1,2)) + '-' \
+ int8totext(substring(pkt4.yiaddr, 2,3)) + '-' \
+ int8totext(substring(pkt4.yiaddr, 3,4))"
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    }
  },
  ...
}

```

Note

The expression value above uses a backslash, \, to show line continuation. This is for clarity only and is not valid JSON supported by Kea parsing. The actual value must be expressed on a single line.

Note

Privacy should be taken into consideration when generating a host name. The host name is usually inserted into the DNS, which is a public system. Exposing identifiers that can be used to track devices, such as a MAC address, are usually a very bad idea. The global expression example here used a MAC address for simplicity.

17.8.1.1 DHCPv4 Host-Name Generation

With this library installed, the behavior for *kea-dhcp4* when forming host names in response to a client query (e.g. DISCOVER, REQUEST) is as follows:

1. If a host name is supplied via a host reservation, use it with the DDNS behavioral parameters to form the final host name. Go to step 4.
2. If the client supplied an FQDN option (option 81), use the domain name value specified within it, with the DDNS behavioral parameters, to form the final host name. Go to step 4.
3. If the client supplied a host-name option (option 12), use the host name specified within it, with the DDNS behavioral parameters, to form the final host name.
4. If there is a *ddns-tuning* in-scope host-name expression (either global or subnet), calculate the host name using the expression. If the calculated value is not a fully qualified name and there is an in-scope *ddns-qualifying-suffix*, append the suffix.
5. If the value calculated by the hook is not an empty string and is different than the host name formed in steps 1 or 2, the calculated value becomes the final host name.

17.8.1.2 DHCPv6 Host-Name Generation

With this library installed, the behavior for *kea-dhcp6* when forming host names in response to a client query (e.g. SOLICIT, REQUEST, RENEW, REBIND) is as follows:

1. If the client supplied an FQDN option (option 39), use the domain name value specified within it, with the DDNS behavioral parameters, to form the final host name. Go to step 4.
2. If the client did not supply an FQDN but *ddns-replace-client-name* is either *always* or *when-not-present*, then calculate the final form of the host name and use it to create an outbound FQDN. Go to step 4.
3. If there is no outbound FQDN at this point, client-name processing for this packet stops. Without an outbound FQDN there is no way to communicate a host name to the client.
4. If a host name is supplied via a host reservation, use it along with the DDNS behavioral parameters to form the final host name; it supersedes the FQDN value calculated in steps 1 or 2.

5. If there is a `ddns-tuning` in-scope host name expression (either global or subnet), calculate the host name using the expression. If the calculated value is not a fully qualified name and there is an in-scope `ddns-qualifying-suffix`, append the suffix.
6. If the value calculated by the hook is not an empty string and is different than the host name formed in steps 1 or 2, the calculated value becomes the final host name.

17.8.2 Skipping DDNS Updates

libdhcp_ddns_tuning.so also provides the ability to skip DDNS updates on a per-client basis. The library recognizes a special client class, "SKIP_DDNS"; when a client is matched to this class, the Kea servers (*kea-dhcp4* and *kea-dhcp6*) do not send DDNS update requests (NCRs) to *kea-dhcp-ddns*. A common use case would be to skip DDNS updates for fixed-address host reservations. This is done easily by simply assigning the class to the host reservation as shown below:

```
{
  "reservations": [
    {
      "hw-address": "01:02:03:04:05:06",
      "ip-address": "192.0.2.1",
      "client-classes": [ "SKIP_DDNS", "foo", "bar" ]
    }
  ]
}
```

The *libdhcp_ddns_tuning.so* hook library notes the presence of the "SKIP_DDNS" class in the client's class list each time the client requests, renews, or releases its lease, and instructs *kea-dhcp4* to bypass sending DDNS updates. A similar workflow is supported for *kea-dhcp6*:

```
{
  "reservations": [
    {
      "duid": "01:02:03:04:05:06",
      "ip-address": "2001:db8::1",
      "client-classes": [ "SKIP_DDNS", "foo", "bar" ]
    }
  ]
}
```

Although "SKIP_DDNS" is a special class, it can be defined with a test expression. Defining it as shown below would omit DDNS updates for all KNOWN clients:

```
{
  "client-classes": [
    {
      "name": "SKIP_DDNS",
      "test": "member('KNOWN')"
    }
  ]
}
```

Note

The *libdhcp_ddns_tuning.so* hook library must be loaded for the "SKIP_DDNS" class to have an effect.

17.9 libdhcp_flex_id.so: Flexible Identifier for Host Reservations

The Kea software provides a way to handle host reservations that include addresses, prefixes, options, client classes, and other features. The reservation can be based on hardware address, DUID, circuit-id, or client-id in DHCPv4 and on hardware address or DUID in DHCPv6. However, there are sometimes scenarios where the reservation is more complex; it may use options other than those mentioned above, use parts of specific options, or perhaps even use a combination of several options and fields to uniquely identify a client. Those scenarios are addressed by the Flexible Identifiers hook application.

Note

libdhcp_flex_id.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

libdhcp_flex_id.so allows the definition of an expression, using notation initially used only for client classification. (See *Using Expressions in Classification* for a detailed description of the syntax available.) One notable difference is that for client classification, the expression currently has to evaluate to either true or false, while the flexible identifier expression is expected to evaluate to a string that will be used as an identifier. It is a valid case for the expression to evaluate to an empty string (e.g. in cases where a client does not send specific options). This expression is then evaluated for each incoming packet, and this evaluation generates an identifier that is used to identify the client. In particular, there may be host reservations that are tied to specific values of the flexible identifier.

The library can be loaded similarly to other hook libraries. It takes a mandatory parameter *identifier-expression* and some optional boolean parameters like *replace-client-id* and *ignore-iaid*:

```
"Dhcp6": {
  "hooks-libraries": [
    {
      "library": "/path/libdhcp_flex_id.so",
      "parameters": {
        "identifier-expression": "expression",
        "replace-client-id": false,
        "ignore-iaid": false
      }
    },
    ...
  ]
}
```

The flexible identifier library supports both DHCPv4 and DHCPv6.

Let's consider a case of an IPv6 network that has an independent interface for each of its connected customers. Customers are able to plug in whatever device they want, so any type of identifier (e.g. a client-id) is unreliable. Therefore, the operator may decide to use an option inserted by a relay agent to differentiate between clients. In this particular deployment, the operator has verified that the interface-id is unique for each customer-facing interface, so it is suitable for usage as a reservation. However, only the first six bytes of the interface-id are interesting, because the remaining bytes are either randomly changed or not unique between devices. Therefore, the customer decides to use the first six bytes of the interface-id option inserted by the relay agent. After adding *flex-id*, the *host-reservation-identifiers* goal can be achieved by using the following configuration:


```

"Dhcp6": {
  "subnet6": [{
    # subnet definition starts here
    "reservations": [{
      "flex-id": "'port1234'",
      # value of the first 8 bytes of the interface-id
      "ip-addresses": [ "2001:db8::1" ]
    },
    ...
  ],
  ...
},
...
],
# end of subnet definitions
"host-reservation-identifiers": ["duid", "flex-id"],
# add "flex-id" to reservation identifiers
"hooks-libraries": [
  {
    "library": "/path/libdhcp_flex_id.so",
    "parameters": {
      "identifier-expression": "substring(relay6[0].option[18].hex,0,8)"
    }
  },
  ...
],
...
}

```

Note

Care should be taken when adjusting the expression. If the expression changes, then all the `flex-id` values may change, possibly rendering all reservations based on `flex-id` unusable until they are manually updated. It is strongly recommended that administrators start with the expression and a handful of reservations, and then adjust the expression as needed. Once the desired result is obtained with the expression, host reservations can be deployed on a broader scale.

`flex-id` values in host reservations can be specified in two ways. First, they can be expressed as a hex string, e.g. the string "bar" can be represented as 626174. Alternatively, it can be expressed as a quoted value (using double and single quotes), e.g. "'bar'". The former is more convenient for printable characters, while hex string values are more convenient for non-printable characters and do not require the use of the `hexstring` operator.

```

"Dhcp6": {
  "subnet6": [{
    # subnet definition starts here
    "reservations": [{
      "flex-id": "01:02:03:04:05:06",
      # value of the first 8 bytes of the interface-id
      "ip-addresses": [ "2001:db8::1" ]
    },
    ...
  ],
  ...
}

```

(continues on next page)

(continued from previous page)

```

],
...
},
...
],
# end of subnet definitions
"host-reservation-identifiers": ["duid", "flex-id"],
# add "flex-id" to reservation identifiers
"hooks-libraries": [
    {
        "library": "/path/libdhcp_flex_id.so",
        "parameters": {
            "identifier-expression": "vendor[4491].option[1026].hex"
        }
    },
    ...
],
...
}

```

Note

One less common scenario where the examples above may prove useful is for DHCPv6 clients that change their DUIDs between exchanges. Certain PXE clients are known to behave this way.

17.9.1 The replace-client-id Flag

When `replace-client-id` is set to `false` (which is the default setting), `libdhcp_flex_id.so` uses the evaluated flexible identifier solely for identifying host reservations, i.e. searching for reservations within a database. This is the functional equivalent of other identifiers, similar to hardware address or circuit-id. However, this mode of operation implies that if a client device is replaced, it may cause a conflict between an existing lease (allocated to the old device) and the new lease being allocated to the new device. The conflict arises because the same flexible identifier is computed for the replaced device, so the server will try to allocate the same lease. The mismatch between client identifiers sent by the new device and the old device causes the server to refuse this new allocation until the old lease expires. A manifestation of this problem is dependent on the specific expression used as the flexible identifier, and is likely to appear if only options and other parameters are used that identify where the device is connected (e.g. circuit-id), rather than the device identification itself (e.g. MAC address).

`libdhcp_flex_id.so` offers a way to overcome the problem with lease conflicts by dynamically replacing the client identifier (or DUID in DHCPv6) with a value derived from the flexible identifier. The server processes the client's query as if the flexible identifier were sent in the client identifier (or DUID) option. This guarantees that a returning client (for which the same flexible identifier is evaluated) will be assigned the same lease, despite the client identifier and/or MAC address change.

The following is a stub configuration that enables this behavior:

```

"Dhcp4": {
    "hooks-libraries": [
        {
            "library": "/path/libdhcp_flex_id.so",
            "parameters": {
                "identifier-expression": "expression",

```

(continues on next page)

(continued from previous page)

```

        "replace-client-id": true
    },
    ...
]
}

```

In the DHCPv4 case, the value derived from the flexible identifier is formed by prepending one byte with a value of zero to the flexible identifier. In the DHCPv6 case, it is formed by prepending two zero bytes before the flexible identifier.

Note that for this mechanism to take effect, the DHCPv4 server must be configured to respect the client identifier option value during lease allocation, i.e. `match-client-id` must be set to `true`. See [Using Client Identifier and Hardware Address](#) for details. No additional settings are required for DHCPv6.

If the `replace-client-id` option is set to `true`, the value of the `echo-client-id` parameter (which governs whether to send back a client-id option) is ignored.

The *libdhcp_lease_cmds.so: Lease Commands for Easier Lease Management* section describes commands used to retrieve, update, and delete leases using various identifiers, such as `hw-address` and `client-id`. *libdhcp_lease_cmds.so* does not natively support querying for leases by flexible identifier. However, when `replace-client-id` is set to `true`, it makes it possible to query for leases using a value derived from the flexible identifier. In DHCPv4, the query looks similar to this:

```

{
  "command": "lease4-get",
  "arguments": {
    "identifier-type": "client-id",
    "identifier": "00:54:64:45:66",
    "subnet-id": 44
  }
}

```

where the hexadecimal value of "54:64:45:66" is a flexible identifier computed for the client.

In DHCPv6, the corresponding query looks something like this:

```

{
  "command": "lease6-get",
  "arguments": {
    "identifier-type": "duid",
    "identifier": "00:00:54:64:45:66",
    "subnet-id": 10
  }
}

```

17.9.2 The ignore-iaid Flag

When `ignore-iaid` is set to `true` (the default value is `false`), *libdhcp_flex_id.so* causes the Kea DHCPv6 server to ignore the IAID value from incoming IPv6 packets. This parameter is ignored by the Kea DHCPv4 server.

If the packet contains only one IA_NA, the IAID value will be changed to 0 and stored as such in the lease storage. Similarly, if the packet contains only one IA_PD, the IAID value will be changed to 0 and stored as such in the lease storage. The IAID is restored to its initial value in the response back to the client. The change is visible in the identifier expression if the IAID is part of the expression.

Note

To avoid lease conflicts, if the incoming packet contains more than one IA_NA, the IAID value is not changed on any of the IA_NAs. Similarly, if the incoming packet contains more than one IA_PD, the IAID value is not changed on any of the IA_PDs.

Warning

This functionality breaks RFC compliance and should be enabled only if required. When enabled, a warning message is issued at configure time.

17.10 `libdhcp_flex_option.so`: Flexible Option Actions for Option Value Settings

This library allows administrators to define an action to take, for a given option, based upon on the result of an expression. These actions are carried out during the final stages of constructing a query response packet, just before it is sent to the client. The three actions currently supported are add, supersede, and remove.

Note

`libdhcp_flex_option.so` is part of the open source code and is available to every Kea user.

The syntax used for the action expressions is the same syntax used for client classification and the Flexible Identifier hook library; see either *Using Expressions in Classification* or *libdhcp_flex_id.so: Flexible Identifier for Host Reservations* for a detailed description of the syntax.

The add and supersede actions use an expression returning a string, and do nothing if the string is empty. The remove application uses an expression returning true or false, and does nothing on false. When it is necessary to set an option to the empty value this mechanism does not work, but a client class can be used instead.

The add action adds an option only when the option does not already exist and the expression does not evaluate to the empty string. The supersede action is similar, but it overwrites the option value if it already exists. The remove action removes the option from the response packet if it already exists and the expression evaluates to true.

The option to which an action applies may be specified by either its numeric code or its name; either the code or the name must be specified. The option space is DHCPv4 or DHCPv6, depending on the server where the hook library is loaded.

`libdhcp_flex_option.so` can be loaded by either the `kea-dhcp4` or `kea-dhcp6` process. It takes a mandatory options parameter with a list of per-option parameter maps, with code, name, add, supersede, and remove actions. Action entries take a string value representing an expression.

```
{
  "Dhcp4": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/libdhcp_flex_option.so",
        "parameters": {
          "options": [
            {
              "code": 67,
```

(continues on next page)

(continued from previous page)

```

        "add": "ifelse(option[host-name].exists,concat(option[host-
↪name].text, '.boot'), '')"
    }
  ]
}

```

If (and only if) the **query** includes a **host-name** option (code 12), a **boot-file-name** option (code 67) is added to the response with the host name followed by **.boot** for content.

A commonly discussed use case is modifying the DHCPv4 subnet mask option (code 1). The following example demonstrates that capability, as all ingress packets identified by the gateway address 192.168.0.1 are met with a /32 subnet mask in the response:

```

{
  "Dhcp4": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/libdhcp_flex_option.so",
        "parameters": {
          "options": [
            {
              "code": 1,
              "supersede": "ifelse(pkt4.giaddr==192.168.0.1, '255.255.255.
↪255', '')"
            }
          ]
        }
      }
    ]
  }
}

```

The Flexible Option library supports both DHCPv4 and DHCPv6.

The **add** and **supersede** actions take an optional **csv-format** boolean parameter. If not specified or set to **false**, the option data is set using the raw value of the evaluated expression. When it is configured to **true**, this value is parsed using the option definition from the option data specified in the configuration file. This eases option setting for options using complex record formats or fully qualified domain names.

For instance, if the expression evaluation returns "example.com" and the option is defined with the **fqdn** type, the domain name is encoded into DNS binary format.

Since Kea 2.1.4, the **client-class** parameter specifies a class guard, and takes a client-class name. If not empty, the client's packet needs to belong to the specified class for this entry to be used.

It is also possible to have multiple entries for the same option, but each entry must have exactly one action. If the option is not defined in **dhcp4** for DHCPv4 or **dhcp6** for DHCPv6, the location of the option definition can be specified using its name, with the new **space** parameter.

Since Kea 2.1.4, sub-options are supported with a new entry, **sub-options**, which replaces the action in the configuration of the container option, i.e. the option where sub-options are located.

The `sub-options` parameter takes a list of sub-option entries, each containing:

- `code` - specifies the sub-option code; either the code or name must be specified. If both are given, they must match or the configuration is rejected at load time.
- `name` - specifies the sub-option name; either the code or name must be specified. If both are given, they must match or the configuration is rejected at load time.
- `space` - specifies the space where the sub-option can be defined. This parameter is optional, because it can be found in the container option definition. The configuration is rejected if no valid space name is available at load time. Note that vendor spaces are supported for the DHCPv4 `vivso-suboptions` and for the DHCPv6 `vendor-opts`, both pre-defined (e.g. DoCSIS vendor id 4491) or custom.
- `add` - (action) adds a sub-option only if it does not already exist and the expression does not evaluate to the empty string.
- `supersede` - (action) adds or overwrites a sub-option if the expression does not evaluate to the empty string.
- `remove` - (action) removes a sub-option if it already exists and the expression evaluates to true.
- `container-add` - a boolean value that specifies whether the container option should be created, if it does not exist in the `add` and `supersede` action. If not specified, it defaults to `true`.
- `container-remove` - a boolean value that specifies whether the container option should be deleted, if it remains empty after the removal of a sub-option by the `remove` action. If not specified, it defaults to `true`.
- `csv-format` - a boolean value that specifies whether the raw value of the evaluated expression is used (false, default) or parsed using the sub-option definition (true).
- `client-class` - specifies whether the sub-option entry must be skipped when the **query** does not belong to the specified client class. Note that the similar parameter in the container option entry applies to the whole sub-options list.

For instance, this configuration adds a string sub-option in the DHCPv4 `vendor-encapsulated-options` (code 43) option. This option encapsulates the `vendor-encapsulated-options` space.

```
{
  "Dhcp4": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/libdhcp_flex_option.so",
        "parameters": {
          "options": [
            {
              "code": 43,
              "sub-options": [
                {
                  "code": 1,
                  "add": "'foobar'"
                }
              ]
            }
          ]
        }
      }
    ]
  }
}
```

17.11 libddns_gss_tsig.so: Sign DNS Updates With GSS-TSIG

This hook library allows the *kea-dhcp-ddns* server to use GSS-TSIG to sign DNS updates. For a full discussion of GSS-TSIG in Kea, please see *GSS-TSIG*.

Note

libddns_gss_tsig.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

17.12 libdhcp_ha.so: High Availability Outage Resilience for Kea Servers

This hook library can be loaded on a pair of DHCPv4 or DHCPv6 servers, to increase the reliability of the DHCP service in the event of an outage on one server.

Note

libdhcp_ha.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

High Availability (HA) of the DHCP service is provided by running multiple cooperating server instances. If any of these instances becomes unavailable for any reason (DHCP software crash, Control Agent software crash, power outage, hardware failure), a surviving server instance can continue providing reliable service to clients. Many DHCP server implementations include the "DHCP Failover" protocol, whose most significant features are communication between the servers, partner failure detection, and lease synchronization between the servers. However, the DHCPv4 failover standardization process was never completed by the IETF. The DHCPv6 failover standard (RFC 8156) was published, but it is complex and difficult to use, has significant operational constraints, and is different from its v4 counterpart. Although it may be useful to use a "standard" failover protocol, most Kea users are simply interested in a working solution which guarantees high availability of the DHCP service. Therefore, the Kea HA hook library derives major concepts from the DHCP failover protocol but uses its own solutions for communication and configuration. It offers its own state machine, which greatly simplifies its implementation and generally fits better into Kea, and it provides the same features in both DHCPv4 and DHCPv6. This document intentionally uses the term "high availability" rather than "failover" to emphasize that it is not the failover protocol implementation.

The following sections describe the configuration and operation of the Kea HA hook library.

17.12.1 Supported Configurations

The Kea HA hook library supports three configurations, also known as HA modes: *load-balancing*, *hot-standby*, and *passive-backup*. In the *load-balancing* mode, two servers respond to DHCP requests. The *load-balancing* function is implemented as described in [RFC 3074](#), with each server responding to half the received DHCP queries. When one of the servers allocates a lease for a client, it notifies the partner server over the control channel (via the RESTful API), so the partner can save the lease information in its own database. If the communication with the partner is unsuccessful, the DHCP query is dropped and the response is not returned to the DHCP client. If the lease update is successful, the response is returned to the DHCP client by the server which has allocated the lease. By exchanging

lease updates, both servers get a copy of all leases allocated by the entire HA setup, and either server can be switched to handle the entire DHCP traffic if its partner becomes unavailable.

In the **load-balancing** configuration, one of the servers must be designated as **primary** and the other as **secondary**. Functionally, there is no difference between the two during normal operation. However, this distinction is required when the two servers are started at (nearly) the same time and have to synchronize their lease databases. The primary server synchronizes the database first. The secondary server waits for the primary server to complete the lease database synchronization before it starts the synchronization.

In the **hot-standby** configuration, one of the servers is designated as **primary** and the other as **standby**. During normal operation, the primary server is the only one that responds to DHCP requests. The standby server receives lease updates from the primary over the control channel; however, it does not respond to any DHCP queries as long as the primary is running or, more accurately, until the standby considers the primary to be offline. If the standby server detects the failure of the primary, it starts responding to all DHCP queries.

Note

Operators often wonder whether to use **load-balancing** or **hot-standby** mode. The **load-balancing** mode has the benefit of splitting the DHCP load between two instances, reducing the traffic processed by each of them. However, it is not always clear to the operators that using the **load-balancing** mode requires manually splitting the address pools between two Kea instances using client classification, to preclude both servers from allocating the same address to different clients. Such a split is not needed in the **hot-standby** mode. Thus, the benefit of using **hot-standby** over **load-balancing** is that the former has a simpler configuration. Conversely, **load-balancing** has higher performance potential at the cost of more complex configuration. See [Load-Balancing Configuration](#) for details on how to split the pools using client classification.

In the configurations described above, both the primary and secondary/standby are referred to as **active** servers, because they receive lease updates and can automatically react to the partner's failures by responding to the DHCP queries which would normally be handled by the partner. The HA hook library supports another server type/role: **backup**. The use of a backup server is optional, and can be implemented in both **load-balancing** and **hot-standby** setup, in addition to the active servers. There is no limit on the number of backup servers in the HA setup; however, the presence of backup servers may increase the latency of DHCP responses, because not only do active servers send lease updates to each other, but also to the backup servers. The active servers do not expect acknowledgments from the backup servers before responding to the DHCP clients, so the overhead of sending lease updates to the backup servers is minimized.

In the last supported configuration, **passive-backup**, there is only one active server and typically one or more backup servers. A **passive-backup** configuration with no backup servers is also accepted, but it is no different than running a single server with no HA function at all.

The **passive-backup** configuration is used in situations when an administrator wants to take advantage of the backup server(s) as an additional storage for leases without running the full-blown failover setup. In this case, if the primary server fails, the DHCP service is lost; it requires the administrator to manually restart the primary to resume DHCP service. The administrator may also configure one of the backup servers to provide DHCP service to the clients, as these servers should have accurate or nearly accurate information about the allocated leases. The major advantage of the **passive-backup** mode is that it provides some redundancy of the lease information but with better performance of the primary server responding to the DHCP queries. The primary server does not have to wait for acknowledgments to the lease updates from the backup servers before it sends a response to the DHCP client. This reduces the response time compared to the **load-balancing** and **hot-standby** cases, in which the server responding to the DHCP query has to wait for the acknowledgment from the other active server before it can respond to the client.

Note

An interesting use case for a single active server running in the **passive-backup** mode is a notification service, in

which software pretending to be a backup server receives live notifications about allocated and deleted leases from the primary server and can display them on a monitoring screen, trigger alerts, etc.

17.12.2 Clocks on Active Servers

Synchronized clocks are essential for the HA setup to operate reliably. The servers share lease information - via lease updates and during synchronization of the databases - including the time when the lease was allocated and when it expires. Some clock skew between the servers participating in the HA setup usually exists; this is acceptable as long as the clock skew is relatively low, compared to the lease lifetimes. However, if the clock skew becomes too high, the different lease expiration times on different servers may cause the HA system to malfunction. For example, one server may consider a lease to be expired when it is actually still valid. The lease reclamation process may remove a name associated with this lease from the DNS, causing problems when the client later attempts to renew the lease.

Each active server monitors the clock skew by comparing its current time with the time returned by its partner in response to the `ha-heartbeat` command. This gives a good approximation of the clock skew, although it does not take into account the time between the partner sending the response and the receipt of this response by the server which sent the `ha-heartbeat` command. If the clock skew exceeds 30 seconds, a warning log message is issued. The administrator may correct this problem by synchronizing the clocks (e.g. using NTP); the servers should notice the clock skew correction and stop issuing the warning.

If the clock skew is not corrected and exceeds 60 seconds, the HA service on each of the servers is terminated, i.e. the state machine enters the `terminated` state. The servers will continue to respond to DHCP clients (as in the load-balancing or hot-standby mode), but will exchange neither lease updates nor heartbeats and their lease databases will diverge. In this case, the administrator should synchronize the clocks and restart the servers.

Note

It is possible to restart the servers one at a time, in no particular order. The clocks must be in sync before restarting the servers. A restarted server remains in the `waiting` state until the partner is also restarted. If the partner is not restarted within 10 minutes, the restarted server transitions back to the `terminated` state to continue serving the DHCP clients. Restart both servers to resume the normal HA service.

Note

The clock skew is only assessed between two active servers, and only the active servers enter the `terminated` state if the skew is too high. The clock skew between active and backup servers is not assessed, because active servers do not exchange heartbeat messages with backup servers.

17.12.3 HTTPS Support

Since Kea 1.9.7, the High Availability hook library supports HTTPS via TLS, as described in [TLS/HTTPS Support](#).

The HTTPS configuration parameters are:

- `trust-anchor` - specifies the name of a file or directory where the certification authority certificate of a Control Agent can be found.
- `cert-file` - specifies the name of the file containing the end-entity certificate to use.
- `key-file` - specifies the private key of the end-entity certificate to use.

These parameters can be configured at the global and peer levels. When configured at both levels the peer value is used, allowing common values to be shared.

The three parameters must be either all not specified (HTTPS disabled) or all specified (HTTPS enabled). Specification of the empty string is considered not specified; this can be used, for instance, to disable HTTPS for a particular peer when it is enabled at the global level.

As the High Availability hook library is an HTTPS client, there is no `cert-required` parameter in this hook configuration. This parameter can be set in the Control Agent to require and verify a client certificate in client-server communication. It does not affect communication between HA peers at the client side; see below for information on the server side.

Before Kea 2.1.7 using HTTPS in the HA setup required use of the Control Agent on all peers. (See [TLS/HTTPS Support](#) for Control Agent TLS configuration).

Since Kea 2.1.7 the HTTPS server side is supported:

- the peer entry for the server name is used for the TLS setting.
- the new `require-client-certs` parameter specifies whether client certificates are required and verified, i.e. like `cert-required`. It defaults to `true` and is an HA config (vs. peer config) parameter.

Kea 2.1.7 added a new security feature with the `restrict-commands` HA config parameter: when set to `true`, commands which are not used by the hook are rejected. The default is `true` since Kea 3.0.0.

The following is an example of an HA server pair and Control Agent configuration for `hot-standby` with TLS.

Server 1:

```
{
  "Dhcp4": {
    "hooks-libraries": [{
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    }, {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [{
          "this-server-name": "server1",
          "trust-anchor": "/usr/lib/kea/CA.pem",
          "cert-file": "/usr/lib/kea/server1_cert.pem",
          "key-file": "/usr/lib/kea/server1_key.pem",
          "mode": "hot-standby",
          "heartbeat-delay": 10000,
          "max-response-delay": 60000,
          "max-ack-delay": 5000,
          "max-unacked-clients": 5,
          "peers": [{
            "name": "server1",
            "url": "http://192.168.56.33:8005/",
            "role": "primary",
            "auto-failover": true
          }, {
            "name": "server2",
            "url": "http://192.168.56.66:8005/",
            "role": "standby",
            "auto-failover": true
          }
        ]
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    }],
    "subnet4": [{
        "id": 1,
        "subnet": "192.0.3.0/24",
        "pools": [{
            "pool": "192.0.3.100 - 192.0.3.250"
        }]
    }]
}

```

Server 2:

```

{
  "Dhcp4": {
    "hooks-libraries": [{
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    }, {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [{
          "this-server-name": "server2",
          "trust-anchor": "/usr/lib/kea/CA.pem",
          "cert-file": "/usr/lib/kea/server2_cert.pem",
          "key-file": "/usr/lib/kea/server2_key.pem",
          "mode": "hot-standby",
          "heartbeat-delay": 10000,
          "max-response-delay": 60000,
          "max-ack-delay": 5000,
          "max-unacked-clients": 5,
          "peers": [{
            "name": "server1",
            "url": "http://192.168.56.33:8005/",
            "role": "primary",
            "auto-failover": true
          }, {
            "name": "server2",
            "url": "http://192.168.56.66:8005/",
            "role": "standby",
            "auto-failover": true
          }]
        }]
      }
    }],
    "subnet4": [{
      "id": 1,
      "subnet": "192.0.3.0/24",
      "pools": [{
        "pool": "192.0.3.100 - 192.0.3.250"
      }]
    }]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }}
  }
}

```

Control Agent on Server 1:

```

{
  "Control-agent": {
    "http-host": "192.168.56.33",
    "http-port": 8000,
    "control-sockets": {
      "dhcp4": {
        "socket-type": "unix",
        "socket-name": "/var/run/kea/control_socket"
      }
    },
    "trust-anchor": "/var/lib/kea/CA.pem",
    "cert-file": "/var/lib/kea/server1_cert.pem",
    "key-file": "/var/lib/kea/server1_key.pem",
    "cert-required": true
  }
}

```

Control Agent on Server 2:

```

{
  "Control-agent": {
    "http-host": "192.168.56.66",
    "http-port": 8000,
    "control-sockets": {
      "dhcp4": {
        "socket-type": "unix",
        "socket-name": "/var/run/kea/control_socket"
      }
    },
    "trust-anchor": "/var/lib/kea/CA.pem",
    "cert-file": "/var/lib/kea/server2_cert.pem",
    "key-file": "/var/lib/kea/server2_key.pem",
    "cert-required": true
  }
}

```

17.12.4 Server States

A DHCP server operating within an HA setup runs a state machine, and the state of the server can be retrieved by its peers using the [ha-heartbeat](#) command sent over the RESTful API. If the partner server does not respond to the [ha-heartbeat](#) command within the specified amount of time, the communication is considered interrupted and the server may, depending on the configuration, use additional measures (described later in this document) to verify that the partner is still operating. If it finds that the partner is not operating, the server transitions to the **partner-down** state to handle all the DHCP traffic directed to the system.

In this case, the surviving server continues to send the [ha-heartbeat](#) command to detect when the partner wakes up.

At that time, the partner synchronizes the lease database. When it is again ready to operate, the surviving server returns to normal operation, i.e. the `load-balancing` or `hot-standby` state.

The following is the list of all possible server states:

- `backup` - normal operation of the backup server. In this state it receives lease updates from the active server(s).
- `communication-recovery` - an active server running in `load-balancing` mode may transition to this state when it experiences communication issues with a partner server over the control channel. This is an intermediate state between the `load-balancing` and `partner-down` states. In this state the server continues to respond to DHCP queries but does not send lease updates to the partner; lease updates are queued and are sent when normal communication is resumed. If communication does not resume within the time specified, the primary server then transitions to the `partner-down` state. The `communication-recovery` state was introduced to ensure reliable DHCP service when both active servers remain operational but the communication between them is interrupted for a prolonged period of time. Either server can be configured to never enter this state by setting the `delayed-updates-limit` to 0 (please refer to [Load-Balancing Configuration](#), later in this chapter, for details on this parameter). Disabling entry into the `communication-recovery` state causes the server to begin testing for the `partner-down` state as soon as the server is unable to communicate with its partner.

Note

In Kea 1.9.4, with the introduction of `delayed-updates-limit`, the default server's behavior in `load-balancing` mode changed. When a server experiences communication issues with its partner, it now enters the `communication-recovery` state and queues lease updates until communication is resumed. Prior to Kea 1.9.4, a server that could not communicate with its partner in `load-balancing` mode would immediately begin the transition to the `partner-down` state.

- `hot-standby` - normal operation of the active server running in the `hot-standby` mode; both the primary and the standby server are in this state during their normal operation. The primary server responds to DHCP queries and sends lease updates to the standby server and to any backup servers that are present.
- `load-balancing` - normal operation of the active server running in the `load-balancing` mode; both the primary and the secondary server are in this state during their normal operation. Both servers respond to DHCP queries and send lease updates to each other and to any backup servers that are present.
- `in-maintenance` - an active server transitions to this state as a result of being notified by its partner that the administrator requested maintenance of the HA setup. The administrator requests the maintenance by sending the `ha-maintenance-start` command to the server which is supposed to take over the responsibility for responding to the DHCP clients while the other server is taken offline for maintenance. If the server is in the `in-maintenance` state it can be safely shut down. The partner transitions to the `partner-down` state immediately after discovering that the server in maintenance has been shut down.
- `partner-down` - an active server transitions to this state after detecting that its partner (another active server) is offline. The server does not transition to this state if only a backup server is unavailable. In the `partner-down` state the active server responds to all DHCP queries, including those queries which are normally handled by the server that is now unavailable.
- `partner-in-maintenance` - an active server transitions to this state after receiving a `ha-maintenance-start` command from the administrator. The server in this state becomes responsible for responding to all DHCP requests. The server sends a `ha-maintenance-notify` command to the partner, which should enter the `in-maintenance` state. The server remaining in the `partner-in-maintenance` state keeps sending lease updates to the partner until it finds that the partner has stopped responding to those lease updates, heartbeats, or any other commands. In this case, the server in the `partner-in-maintenance` state transitions to the `partner-down` state and keeps responding to the queries, but no longer sends lease updates.
- `passive-backup` - a primary server running in the `passive-backup` HA mode transitions to this state immediately after it boots up. The primary server in this state responds to all DHCP traffic and sends lease updates

to the backup servers it is connected to. By default, the primary server does not wait for acknowledgments from the backup servers and responds to a DHCP query right after sending lease updates to all backup servers. If any of the lease updates fail, a backup server misses the lease update but the DHCP client is still provisioned. This default configuration can be changed by setting the `wait-backup-ack` configuration parameter to `true`, in which case the primary server always waits for the acknowledgements and drops the DHCP query if sending any of the corresponding lease updates fails. This improves lease database consistency between the primary and the secondary. However, if a communication failure between the active server and any of the backups occurs, it effectively causes the failure of the DHCP service from the DHCP clients' perspective.

- **ready** - an active server transitions to this state after synchronizing its lease database with an active partner. This state indicates to the partner (which may be in the `partner-down` state) that it should return to normal operation. If and when it does, the server in the `ready` state also starts normal operation.
- **syncing** - an active server transitions to this state to fetch leases from the active partner and update the local lease database. When in this state, the server issues the `dhcp-disable` command to disable the DHCP service of the partner from which the leases are fetched. The DHCP service is disabled for a maximum time of 60 seconds, after which it is automatically re-enabled, in case the syncing partner was unable to re-enable the service. If the synchronization completes successfully, the synchronizing server issues the `ha-sync-complete-notify` command to notify the partner. In most states, the partner re-enables its DHCP service to continue responding to the DHCP queries. In the `partner-down` state, the partner first ensures that communication between the servers is re-established before enabling the DHCP service. The syncing operation is synchronous; the server waits for an answer from the partner and does nothing else while the lease synchronization takes place. A server that is configured not to synchronize the lease database with its partner, i.e. when the `sync-leases` configuration parameter is set to `false`, will never transition to this state. Instead, it transitions directly from the `waiting` state to the `ready` state.
- **terminated** - an active server transitions to this state when the High Availability hook library is unable to further provide reliable service and a manual intervention of the administrator is required to correct the problem. Various issues with the HA setup may cause the server to transition to this state. While in this state, the server continues responding to DHCP clients based on the HA mode selected (`load-balancing` or `hot-standby`), but lease updates are not exchanged and heartbeats are not sent. Once a server has entered the `terminated` state, it remains in this state until it is restarted. The administrator must correct the issue which caused this situation prior to restarting the server (e.g. synchronize the clocks); otherwise, the server will return to the `terminated` state once it finds that the issue persists.
- **waiting** - each started server instance enters this state. A backup server transitions directly from this state to the `backup` state. An active server sends a heartbeat to its partner to check its state; if the partner appears to be unavailable, the server transitions to the `partner-down` state. If the partner is available, the server transitions to the `syncing` or `ready` state, depending on the setting of the `sync-leases` configuration parameter. If both servers appear to be in the `waiting` state (concurrent startup), the primary server transitions to the next state first. The secondary or standby server remains in the `waiting` state until the primary transitions to the `ready` state.

Note

Currently, restarting the HA service from the `terminated` state requires restarting the DHCP server or reloading its configuration.

Whether the server responds to DHCP queries and which queries it responds to is a matter of the server's state, if no administrative action is performed to configure the server otherwise. The following table provides the default behavior for various states.

The DHCP Service Scopes denote which group of received DHCP queries the server responds to in the given state. The HA configuration must specify a unique name for each server within the HA setup. This document uses the following convention within the provided examples: "server1" for a primary server, "server2" for the secondary or

standby server, and "server3" for the backup server. In real life any names can be used as long as they remain unique. An in-depth explanation of the scopes can be found below.

Table 2: Default behavior of the server in various HA states

State	Server Type	DHCP Service	DHCP Service Scopes
backup	backup server	disabled	none
communication-recovery	primary or secondary (load-balancing mode only)	enabled	"HA_server1" or "HA_server2"
hot-standby	primary or standby (hot-standby mode)	enabled	"HA_server1" if primary, none otherwise
load-balancing	primary or secondary (load-balancing mode)	enabled	"HA_server1" or "HA_server2"
in-maintenance	active server	disabled	none
partner-down	active server	enabled	all scopes
partner-in-maintenance	active server	enabled	all scopes
passive-backup	active server	enabled	all scopes
ready	active server	disabled	none
syncing	active server	disabled	none
terminated	active server	enabled	same as in the load-balancing or hot-standby state
waiting	any server	disabled	none

In the **load-balancing** mode there are two scopes specified for the active servers: "HA_server1" and "HA_server2". The DHCP queries load-balanced to **server1** belong to the "HA_server1" scope and the queries load-balanced to **server2** belong to the "HA_server2" scope. If either server is in the **partner-down** state, the active partner is responsible for serving both scopes.

In the **hot-standby** mode, there is only one scope - "HA_server1" - because only **server1** is responding to DHCP queries. If that server becomes unavailable, **server2** becomes responsible for this scope.

The backup servers do not have their own scopes. In some cases they can be used to respond to queries belonging to the scopes of the active servers. Also, a backup server which is neither in the **partner-down** state nor in normal operation serves no scopes.

The scope names can be used to associate pools, subnets, and networks with certain servers, so that only these servers can allocate addresses or prefixes from those pools, subnets, or networks. This is done via the client classification mechanism (see [Load Balancing With Advanced Classification](#) for more details).

17.12.5 Scope Transition in a Partner-Down Case

When one of the servers finds that its partner is unavailable, it starts serving clients from both its own scope and the scope of the unavailable partner. This is straightforward for new clients, i.e. those sending DHCPDISCOVER (DHCPv4) or Solicit (DHCPv6), because those requests are not sent to any particular server. The available server responds to all such queries when it is in the **partner-down** state.

When a client renews a lease, it sends its DHCPREQUEST (DHCPv4) or Renew (DHCPv6) message directly to the server which has allocated the lease being renewed. If this server is no longer available, the client will get no response. In that case, the client continues to use its lease and attempts to renew until the rebind timer (T2) elapses. The client then enters the rebinding phase, in which it sends a DHCPREQUEST (DHCPv4) or Rebind (DHCPv6) message to any available server. The surviving server receives the rebinding request and typically extends the lifetime of the lease. The client then continues to contact that new server to renew its lease as appropriate.

If and when the other server once again becomes available, both active servers will eventually transition to the load-balancing or hot-standby state, in which they will again be responsible for their own scopes. Some clients belonging to the scope of the restarted server will try to renew their leases via the surviving server, but this server will no longer respond to them; the client will eventually transition back to the correct server via the rebinding mechanism.

17.12.6 Load-Balancing Configuration

The following is the configuration snippet to enable high availability on the primary server within the load-balancing configuration. The same configuration should be applied on the secondary and backup servers, with the only difference that `this-server-name` should be set to "server2" and "server3" on those servers, respectively.

Note

Remember that load-balancing mode requires the address pools and delegated prefix pools to be split between the active servers. During normal operation, the servers use non-overlapping pools to avoid allocating the same lease to different clients by both instances. A server only uses the pool fragments owned by the partner when the partner is not running. See the notes in *Supported Configurations* highlighting differences between the load-balancing and hot-standby modes. The semantics of pool partitioning is explained further in this section. The *Load Balancing With Advanced Classification* section provides advanced pool-partitioning examples.

```
"Dhcp4": {
  "hooks-libraries": [{
    "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
    "parameters": { }
  }, {
    "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
    "parameters": {
      "high-availability": [{
        "this-server-name": "server1",
        "mode": "load-balancing",
        "heartbeat-delay": 10000,
        "max-response-delay": 60000,
        "max-ack-delay": 5000,
        "max-unacked-clients": 5,
        "max-rejected-lease-updates": 10,
        "delayed-updates-limit": 100,
        "peers": [{
          "name": "server1",
          "url": "http://192.168.56.33:8005/",
          "role": "primary",
          "auto-failover": true
        }, {
          "name": "server2",
          "url": "http://192.168.56.66:8005/",
          "role": "secondary",
          "auto-failover": true
        }, {
          "name": "server3",
          "url": "http://192.168.56.99:8005/",
          "role": "backup",
          "basic-auth-user": "foo",
          "basic-auth-password": "1234",
```

(continues on next page)

(continued from previous page)

```

        "auto-failover": false
    }
}
}],
"subnet4": [{
    "id": 1,
    "subnet": "192.0.3.0/24",
    "pools": [{
        "pool": "192.0.3.100 - 192.0.3.150",
        "client-classes": [ "HA_server1" ]
    }, {
        "pool": "192.0.3.200 - 192.0.3.250",
        "client-classes": [ "HA_server2" ]
    }],
    "option-data": [{
        "name": "routers",
        "data": "192.0.3.1"
    }],
    "relay": { "ip-address": "10.1.2.3" }
}]
}

```

Two hook libraries must be loaded to enable HA: *libdhcp_lease_cmds.so* and *libdhcp_ha.so*. The latter implements the HA feature, while the former enables control commands required by HA to fetch and manipulate leases on the remote servers. In the example provided above, it is assumed that Kea libraries are installed in the `/usr/lib` directory. If Kea is not installed in the `/usr` directory, the hook libraries' locations must be updated accordingly.

The HA configuration is specified within the scope of *libdhcp_ha.so*. Note that while the top-level parameter `high-availability` is a list, only a single entry is currently supported.

The following are the global parameters which control the server's behavior with respect to HA:

- `this-server-name` - is a unique identifier of the server within this HA setup. It must match one of the servers specified within the `peers` list.
- `mode` - specifies an HA mode of operation. The currently supported modes are `load-balancing` and `hot-standby`.
- `heartbeat-delay` - specifies a duration in milliseconds between sending the last heartbeat (or other command sent to the partner) and the next heartbeat. Heartbeats are sent periodically to gather the status of the partner and to verify whether the partner is still operating. The default value of this parameter is 10000 ms.
- `max-response-delay` - specifies a duration in milliseconds since the last successful communication with the partner, after which the server assumes that communication with the partner is interrupted. This duration should be greater than the `heartbeat-delay`; typically it should be a multiple of `heartbeat-delay`. When the server detects that communication is interrupted, it may transition to the `partner-down` state (when `max-unacked-clients` is 0) or trigger the failure-detection procedure using the values of the two parameters below. The default value of this parameter is 60000 ms.
- `max-ack-delay` - is one of the parameters controlling partner failure-detection. When communication with the partner is interrupted, the server examines the values of the `"secs"` field (DHCPv4) or `"elapsed time"` option (DHCPv6), which denote how long the DHCP client has been trying to communicate with the DHCP server. This

parameter specifies the maximum time in milliseconds for the client to try to communicate with the DHCP server, after which this server assumes that the client failed to communicate with the DHCP server (is unacknowledged or "unacked"). The default value of this parameter is 10000.

- **max-unacked-clients** - specifies how many "unacked" clients are allowed (see **max-ack-delay**) before this server assumes that the partner is offline and transitions to the **partner-down** state. The special value of 0 is allowed for this parameter, which disables the failure-detection mechanism. In this case, a server that cannot communicate with its partner over the control channel assumes that the partner server is down and transitions to the **partner-down** state immediately. The default value of this parameter is 10.
- **max-rejected-lease-updates** - specifies how many lease updates for distinct clients can fail, due to a conflict between the lease and the partner configuration or state, before the server transitions to the **terminated** state. Conflict can be a sign of a misconfiguration; usually, a small number of conflicted leases are acceptable because they affect only a few devices. However, if the conflicts occur for many devices (e.g., an entire subnet), the HA service becomes unreliable and should be terminated, and the problem must be manually corrected by an administrator. It is up to the administrator to select the highest acceptable value of **max-rejected-lease-updates**. The default value is 10. The special value of 0 configures the server to never terminate the HA service due to lease conflicts. If the value is 1, the server transitions to the **terminated** state when the first conflict occurs. This parameter does not pertain to conflicting lease updates sent to the backup servers.
- **delayed-updates-limit** - specifies the maximum number of lease updates which can be queued while the server is in the **communication-recovery** state. This parameter was introduced in Kea 1.9.4. The special value of 0 configures the server to never transition to the **communication-recovery** state and the server behaves as in earlier Kea versions, i.e. if the server cannot reach its partner, it goes straight into the **partner-down** state. The default value of this parameter is 100.

Note

The **max-rejected-lease-updates** parameter was introduced in Kea 2.3.1. Previously, the server did not differentiate between a lease update failure due to a non-functioning partner and a failure due to a conflict (e.g., configuration issues). As a result, the server could sometimes transition to the **partner-down** state even though the partner was operating normally, but only certain leases had issues. Conflicts should no longer cause such a transition. However, depending on the **max-rejected-lease-updates** setting, too many conflicts can lead to termination of the High Availability service. In that case, both servers continue to respond to DHCP queries but no longer send lease updates.

The values of **max-ack-delay** and **max-unacked-clients** must be selected carefully, taking into account the specifics of the network in which the DHCP servers are operating. The server in question may not respond to some DHCP clients following administrative policy, or the server may drop malformed queries from clients. Therefore, selecting too low a value for the **max-unacked-clients** parameter may result in a transition to the **partner-down** state even though the partner is still operating. On the other hand, selecting too high a value may result in never transitioning to the **partner-down** state if the DHCP traffic in the network is very low (e.g. at night), because the number of distinct clients trying to communicate with the server could be lower than the **max-unacked-clients** setting.

In some cases it may be useful to disable the failure-detection mechanism altogether, if the servers are located very close to each other and network partitioning is unlikely, i.e. failure to respond to heartbeats is only possible when the partner is offline. In such cases, set **max-unacked-clients** to 0.

The **delayed-updates-limit** parameter is used to enable or disable the **communication-recovery** procedure, and controls the server's behavior in the **communication-recovery** state. This parameter can only be used in the load-balancing mode.

If a server in the load-balancing state experiences communication issues with its partner (a heartbeat or lease-update failure), the server transitions to the **communication-recovery** state. In this state, the server keeps responding to DHCP queries but does not send lease updates to the partner. The lease updates are queued until communication is re-established, to ensure that DHCP service remains available even in the event of the communication loss between

the partners. There may appear to be communication loss when either one of the servers has terminated, or when both servers remain available but cannot communicate with each other. In the former case, the surviving server will follow the normal procedure and should eventually transition to the `partner-down` state. In the latter case, both servers should transition to the `communication-recovery` state and should never transition to the `partner-down` state (if `max-unacked-clients` is set to a non-zero value), because all DHCP queries are answered and neither server would see any unacked DHCP queries.

Introduction of the `communication-recovery` procedure was motivated by issues which may appear when two servers remain online but the communication between them remains interrupted for a period of time. In earlier Kea versions, the servers having communication issues used to drop DHCP packets before transitioning to the `partner-down` state. In some cases they both transitioned to the `partner-down` state, which could potentially result in allocations of the same IP addresses or delegated prefixes to different clients by both servers. By entering the intermediate `communication-recovery` state, these problems are avoided.

If a server in the `communication-recovery` state re-establishes communication with its partner, it tries to send the partner all of the outstanding lease updates it has queued. This is done synchronously and may take a considerable amount of time before the server transitions to the `load-balancing` state and resumes normal operation. The maximum number of lease updates which can be queued in the `communication-recovery` state is controlled by `delayed-updates-limit`. If the limit is exceeded, the server stops queuing lease updates and performs a full database synchronization after re-establishing the connection with the partner, instead of sending outstanding lease updates before transitioning to the `load-balancing` state. Even if the limit is exceeded, the server in the `communication-recovery` state remains responsive to DHCP clients.

It may be preferable to set higher values of `delayed-updates-limit` when there is a risk of prolonged communication interruption between the servers and when the lease database is large, to avoid costly lease-database synchronization. On the other hand, if the lease database is small, the time required to send outstanding lease updates may be longer than the lease-database synchronization. In such cases it may be better to use a lower value, e.g. 10. The default value of 100 is a reasonable compromise and should work well in most deployments with moderate traffic.

Note

This parameter is new and values for it that work well in some environments may not work well in others. Feedback from users will help us build a better working set of recommendations.

The `peers` parameter contains a list of servers within this HA setup. This configuration must contain at least one primary and one secondary server. It may also contain an unlimited number of backup servers. In this example, there is one backup server which receives lease updates from the active servers.

Since Kea version 1.9.0, basic HTTP authentication is available to protect the Kea control agent against local attackers.

These are the parameters specified for each of the peers within this list:

- `name` - specifies a unique name for the server.
- `url` - specifies the URL to be used to contact this server over the control channel. Other servers use this URL to send control commands to that server.
- `basic-auth-user` - specifies the user ID for basic HTTP authentication. If not specified or specified as an empty string, no authentication header is added to HTTP transactions. It must not contain the colon (:) character.
- `basic-auth-password` - specifies the password for basic HTTP authentication. This parameter is ignored when the user ID is not specified or is empty. The password is optional; if not specified, an empty password is used.
- `basic-auth-password-file` - is an alternative to `basic-auth-password`: instead of presenting the password in the configuration file it is specified in the file indicated by this parameter.
- `role` - denotes the role of the server in the HA setup. The following roles are supported in the `load-balancing` configuration: `primary`, `secondary`, and `backup`. There must be exactly one primary and one secondary server

in the load-balancing setup.

- **auto-failover** - a boolean value which denotes whether a server detecting a partner's failure should automatically start serving the partner's clients. The default value of this parameter is **true**.

In our example configuration above, both active servers can allocate leases from the subnet "192.0.3.0/24". This subnet contains two address pools: "192.0.3.100 - 192.0.3.150" and "192.0.3.200 - 192.0.3.250", which are associated with HA server scopes using client classification. When **server1** processes a DHCP query, it uses the first pool for lease allocation. Conversely, when **server2** processes a DHCP query it uses the second pool. If either of the servers is in the **partner-down** state, the other can serve leases from both pools; it selects the pool which is appropriate for the received query. In other words, if the query would normally be processed by **server2** but this server is not available, **server1** allocates the lease from the pool of "192.0.3.200 - 192.0.3.250". The Kea control agent in front of **server3** requires basic HTTP authentication, and authorizes the user ID "foo" with the password "1234".

Note

The url schema can be **http** or **https**, but since Kea version 1.9.6 the **https** schema requires a TLS setup. The hostname part must be an IPv4 address or an IPv6 address between square brackets, e.g. **http://[2001:db8::1]:8080/**. Names are not accepted.

17.12.7 Load Balancing With Advanced Classification

In the previous section, we provided an example of a load-balancing configuration with client classification limited to the "HA_server1" and "HA_server2" classes, which are dynamically assigned to the received DHCP queries. In many cases, HA is needed in deployments which already use some other client classification.

Suppose there is a system which classifies devices into two groups: "phones" and "laptops", based on some classification criteria specified in the Kea configuration file. Both types of devices are allocated leases from different address pools. Introducing HA in load-balancing mode results in a further split of each of those pools, as each server allocates leases for some phones and some laptops. This requires each of the existing pools to be split between "HA_server1" and "HA_server2", so we end up with the following classes:

- "phones_server1"
- "laptops_server1"
- "phones_server2"
- "laptops_server2"

The corresponding server configuration, using advanced classification (and the **member** expression), is provided below. For brevity's sake, the HA hook library configuration has been removed from this example.

```
{
  "Dhcp4": {
    "client-classes": [{
      "name": "phones",
      "test": "substring(option[60].hex,0,6) == 'Aastra'"
    }, {
      "name": "laptops",
      "test": "not member('phones')"
    }, {
      "name": "phones_server1",
      "test": "member('phones') and member('HA_server1')"
    }, {
      "name": "phones_server2",
```

(continues on next page)

(continued from previous page)

```

        "test": "member('phones') and member('HA_server2')"
    }, {
        "name": "laptops_server1",
        "test": "member('laptops') and member('HA_server1')"
    }, {
        "name": "laptops_server2",
        "test": "member('laptops') and member('HA_server2')"
    }
  ],

  "hooks-libraries": [{
    "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
    "parameters": { }
  }, {
    "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
    "parameters": {
      "high-availability": [{
      }]
    }
  }
  ],

  "subnet4": [{
    "id": 1,
    "subnet": "192.0.3.0/24",
    "pools": [{
      "pool": "192.0.3.100 - 192.0.3.125",
      "client-classes": [ "phones_server1" ]
    }, {
      "pool": "192.0.3.126 - 192.0.3.150",
      "client-classes": [ "laptops_server1" ]
    }, {
      "pool": "192.0.3.200 - 192.0.3.225",
      "client-classes": [ "phones_server2" ]
    }, {
      "pool": "192.0.3.226 - 192.0.3.250",
      "client-classes": [ "laptops_server2" ]
    }
  ]],

  "option-data": [{
    "name": "routers",
    "data": "192.0.3.1"
  }],

  "relay": { "ip-address": "10.1.2.3" }
}
}

```

The configuration provided above splits the address range into four pools: two pools dedicated to "HA_server1" and two to "HA_server2". Each server can assign leases to both phones and laptops. Both groups of devices are assigned addresses from different pools. The "HA_server1" and "HA_server2" classes are built-in (see *Built-in Client Classes*) and do not need to be declared. They are assigned dynamically by the HA hook library as a result of the load-balancing algorithm. "phones_*" and "laptop_*" evaluate to `true` when the query belongs to a given combination of other classes, e.g. "HA_server1" and "phones". The pool is selected accordingly as a result of such an evaluation.

Consult *Client Classification* for details on how to use the `member` expression and class dependencies.

17.12.8 Hot-Standby Configuration

The following is an example configuration of the primary server in a hot-standby configuration:

```
"Dhcp4": {
  "hooks-libraries": [{
    "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
    "parameters": { }
  }, {
    "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
    "parameters": {
      "high-availability": [{
        "this-server-name": "server1",
        "mode": "hot-standby",
        "heartbeat-delay": 10000,
        "max-response-delay": 60000,
        "max-ack-delay": 5000,
        "max-unacked-clients": 5,
        "max-rejected-lease-updates": 10,
        "peers": [{
          "name": "server1",
          "url": "http://192.168.56.33:8005/",
          "role": "primary",
          "auto-failover": true
        }, {
          "name": "server2",
          "url": "http://192.168.56.66:8005/",
          "role": "standby",
          "auto-failover": true
        }, {
          "name": "server3",
          "url": "http://192.168.56.99:8005/",
          "basic-auth-user": "foo",
          "basic-auth-password": "1234",
          "role": "backup",
          "auto-failover": false
        }
      ]
    }
  }
}],

  "subnet4": [{
    "id": 1,
    "subnet": "192.0.3.0/24",
    "pools": [{
      "pool": "192.0.3.100 - 192.0.3.250",
      "client-classes": [ "HA_server1" ]
    }
  ],

  "option-data": [{
    "name": "routers",
    "data": "192.0.3.1"
```

(continues on next page)

(continued from previous page)

```

    }],
    "relay": { "ip-address": "10.1.2.3" }
  }
}

```

This configuration is very similar to the load-balancing configuration described in *Load-Balancing Configuration*, with a few notable differences.

The mode is now set to `hot-standby`, in which only one server responds to DHCP clients. If the primary server is online, it responds to all DHCP queries. The `standby` server takes over all DHCP traffic only if it discovers that the primary is unavailable.

In this mode, the non-primary active server is called `standby` and that is its role.

Finally, because there is always only one server responding to DHCP queries, there is only one scope - `"HA_server1"` - in use within pool definitions. In fact, the `client-class` parameter could be removed from this configuration without harm, because there can be no conflicts in lease allocations by different servers as they do not allocate leases concurrently. The `client-class` remains in this example mostly for demonstration purposes, to highlight the differences between the `hot-standby` and `load-balancing` modes of operation.

17.12.9 Passive-Backup Configuration

The following is an example configuration file for the primary server in a `passive-backup` configuration:

```

{
  "Dhcp4": {
    "hooks-libraries": [{
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    }, {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [{
          "this-server-name": "server1",
          "mode": "passive-backup",
          "wait-backup-ack": false,
          "peers": [{
            "name": "server1",
            "url": "http://192.168.56.33:8005/",
            "role": "primary"
          }, {
            "name": "server2",
            "url": "http://192.168.56.66:8005/",
            "role": "backup"
          }, {
            "name": "server3",
            "url": "http://192.168.56.99:8005/",
            "basic-auth-user": "foo",
            "basic-auth-password": "1234",
            "role": "backup"
          }
        ]
      }
    }
  }
}

```

(continues on next page)

(continued from previous page)

```

    }],
    "subnet4": [{
        "id": 1,
        "subnet": "192.0.3.0/24",
        "pools": [{
            "pool": "192.0.3.100 - 192.0.3.250"
        }],
        "option-data": [{
            "name": "routers",
            "data": "192.0.3.1"
        }],
        "relay": { "ip-address": "10.1.2.3" }
    }]
}

```

The configurations of three peers are included: one for the primary and two for the backup servers.

Many of the parameters present in the load-balancing and hot-standby configuration examples are not relevant in the passive-backup mode, thus they are not specified here. For example: heartbeat-delay, max-unacked-clients, max-rejected-lease-updates, and others related to the failover mechanism should not be specified in the passive-backup mode.

The wait-backup-ack is a boolean parameter not present in previous examples. It defaults to false and must not be modified in the load-balancing and hot-standby modes. In the passive-backup mode this parameter can be set to true, which causes the primary server to expect acknowledgments to the lease updates from the backup servers prior to responding to the DHCP client. It ensures that the lease has propagated to all servers before the client is given the lease, but it poses a risk of losing a DHCP service if there is a communication problem with one of the backup servers. This setting also increases the latency of the DHCP response, because of the time that the primary spends waiting for the acknowledgements. We recommend that the wait-backup-ack setting be left at its default value (false) if the DHCP service reliability is more important than consistency of the lease information between the primary and the backups, and in all cases when the DHCP service latency should be minimal.

Note

Currently, active servers place lease updates to be sent to peers onto internal queues (one queue per peer/URL). In passive-backup mode, active servers do not wait for lease updates to be acknowledged; thus during times of heavy client traffic it is possible for the number of lease updates queued for transmission to accumulate faster than they can be delivered. As client traffic lessens the queues begin to empty. Since Kea 2.0.0, active servers monitor the size of these queues and emit periodic warnings (see HTTP_CLIENT_QUEUE_SIZE_GROWING in *Kea Messages Manual*) if they perceive a queue as growing too quickly. The warnings cease once the queue size begins to shrink. These messages are intended as a bellwether and seeing them sporadically during times of heavy traffic load does not necessarily indicate a problem. If, however, they occur continually during times of routine traffic load, they likely indicate potential mismatches in server capabilities and/or configuration; this should be investigated, as the size of the queues may eventually impair an active server's ability to respond to clients in a timely manner.

17.12.10 Lease Information Sharing

An HA-enabled server informs its active partner about allocated or renewed leases by sending appropriate control commands, and the partner updates the lease information in its own database. When the server starts up for the first time or recovers after a failure, it synchronizes its lease database with its partner. These two mechanisms guarantee consistency of the lease information between the servers and allow the designation of one of the servers to handle the entire DHCP traffic load if the other server becomes unavailable.

In some cases, though, it is desirable to disable lease updates and/or database synchronization between the active servers, if the exchange of information about the allocated leases is performed using some other mechanism. Kea supports various database types that can be used to store leases, including MySQL and PostgreSQL. Those databases include built-in solutions for data replication which are often used by Kea administrators to provide redundancy.

The HA hook library supports such scenarios by disabling lease updates over the control channel and/or lease-database synchronization, leaving the server to rely on the database replication mechanism. This is controlled by the two boolean parameters `send-lease-updates` and `sync-leases`, whose values default to `true`:

```
"Dhcp4": {
  "hooks-libraries": [
    {
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    },
    {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [ {
          "this-server-name": "server1",
          "mode": "load-balancing",
          "send-lease-updates": false,
          "sync-leases": false,
          "peers": [
            {
              "name": "server1",
              "url": "http://192.168.56.33:8005/",
              "role": "primary"
            },
            {
              "name": "server2",
              "url": "http://192.168.56.66:8005/",
              "role": "secondary"
            }
          ]
        }
      ]
    }
  ],
  ...
}
```

In the most typical use case, both parameters are set to the same value, i.e. both are `false` if database replication is in use, or both are `true` otherwise. Introducing two separate parameters to control lease updates and lease-database synchronization is aimed at possible special use cases; for example, when synchronization is performed by copying a lease file (therefore `sync-leases` is set to `false`), but lease updates should be conducted as usual (`send-lease-updates` is set to `true`). It should be noted that Kea does not natively support such use cases, but users may develop their own scripts and tools around Kea to provide such mechanisms. The HA hook library configuration is designed to maximize

flexibility of administration.

17.12.11 Controlling Lease-Page Size Limit

An HA-enabled server initiates synchronization of the lease database after downtime or upon receiving the *ha-sync* command. The server uses the commands *lease4-get-page* and *lease6-get-page* to fetch leases from its partner server (lease queries). The size of the results page (the maximum number of leases to be returned in a single response to one of these commands) can be controlled via configuration of the HA hook library. Increasing the page size decreases the number of lease queries sent to the partner server, but it causes the partner server to generate larger responses, which lengthens transmission time as well as increases memory and CPU utilization on both servers. Decreasing the page size helps to decrease resource utilization, but requires more lease queries to be issued to fetch the entire lease database.

The default value of the *sync-page-limit* command controlling the page size is 10000. This means that the entire lease database can be fetched with a single command if the size of the database is equal to or less than 10000 lines.

17.12.12 Timeouts

In deployments with a large number of clients connected to the network, lease-database synchronization after a server failure may be a time-consuming operation. The synchronizing server must gather all leases from its partner, which yields a large response over the RESTful interface. The server receives leases using the paging mechanism described in *Controlling Lease-Page Size Limit*. Before the page of leases is fetched, the synchronizing server sends a *dhcp-disable* command to disable the DHCP service on the partner server. If the service is already disabled, this command resets the timeout for the DHCP service being disabled, which by default is set to 60 seconds. If fetching a single page of leases takes longer than the specified time, the partner server assumes that the synchronizing server has died and resumes its DHCP service. The connection of the synchronizing server with its partner is also protected by the timeout. If the synchronization of a single page of leases takes longer than the specified time, the synchronizing server terminates the connection and the synchronization fails. Both timeout values are controlled by a single configuration parameter, *sync-timeout*. The following configuration snippet demonstrates how to modify the timeout for automatic re-enabling of the DHCP service on the partner server and how to increase the timeout for fetching a single page of leases from 60 seconds to 90 seconds:

```
"Dhcp4": {
  "hooks-libraries": [
    {
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    },
    {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [ {
          "this-server-name": "server1",
          "mode": "load-balancing",
          "sync-timeout": 90000,
          "peers": [
            {
              "name": "server1",
              "url": "http://192.168.56.33:8005/",
              "role": "primary"
            },
            {
              "name": "server2",
              "url": "http://192.168.56.66:8005/",
```

(continues on next page)

(continued from previous page)

```

        "role": "secondary"
    }
  ]
}
],
...
}

```

It is important to note that extending this `sync-timeout` value may sometimes be insufficient to prevent issues with timeouts during lease-database synchronization. The control commands travel via the Control Agent, which also monitors incoming (with a synchronizing server) and outgoing (with a DHCP server) connections for timeouts. The DHCP server also monitors the connection from the Control Agent for timeouts. Those timeouts cannot currently be modified via configuration; extending these timeouts is only possible by modifying them in the Kea code and recompiling the server. The relevant constants are located in the Kea source at: `src/lib/config/timeouts.h`.

17.12.13 Pausing the HA State Machine

The high-availability state machine includes many different states described in detail in [Server States](#). The server enters each state when certain conditions are met, most often taking into account the partner server's state. In some states the server performs specific actions, e.g. synchronization of the lease database in the `syncing` state, or responding to DHCP queries according to the configured mode of operation in the `load-balancing` and `hot-standby` states.

By default, transitions between the states are performed automatically and the server administrator has no direct control over when the transitions take place; in most cases, the administrator does not need such control. In some situations, however, the administrator may want to "pause" the HA state machine in a selected state to perform some additional administrative actions before the server transitions to the next state.

Consider a server failure which results in the loss of the entire lease database. Typically, the server rebuilds its lease database when it enters the `syncing` state by querying the partner server for leases, but it is possible that the partner was also experiencing a failure and lacks lease information. In this case, it may be required to reconstruct lease databases on both servers from some external source, e.g. a backup server. If the lease database is to be reconstructed via the RESTful API, the servers should be started in the initial, i.e. `waiting`, state and remain in this state while leases are being added. In particular, the servers should not attempt to synchronize their lease databases nor start serving DHCP clients.

The HA hook library provides configuration parameters and a command to control pausing and resuming the HA state machine. The following configuration causes the HA state machine to pause in the `waiting` state after server startup.

```

"Dhcp4": {
  "hooks-libraries": [
    {
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    },
    {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [ {
          "this-server-name": "server1",
          "mode": "load-balancing",
          "peers": [
            {

```

(continues on next page)

(continued from previous page)

```

        "name": "server1",
        "url": "http://192.168.56.33:8005/",
        "role": "primary"
    },
    {
        "name": "server2",
        "url": "http://192.168.56.66:8005/",
        "role": "secondary"
    }
],
"state-machine": {
    "states": [
        {
            "state": "waiting",
            "pause": "once"
        }
    ]
}
}
}
}
},
...
}

```

The pause parameter value `once` denotes that the state machine should be paused upon the first transition to the waiting state; later transitions to this state will not cause the state machine to pause. Two other supported values of the pause parameter are `always` and `never`. The latter is the default value for each state, which instructs the server never to pause the state machine.

In order to "unpause" the state machine, the `ha-continue` command must be sent to the paused server. This command does not take any arguments. See *Control Commands for High Availability* for details about commands specific to `libdhcp_ha.so`.

It is possible to configure the state machine to pause in more than one state. Consider the following configuration:

```

"Dhcp4": {
    "hooks-libraries": [
        {
            "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
            "parameters": { }
        },
        {
            "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
            "parameters": {
                "high-availability": [ {
                    "this-server-name": "server1",
                    "mode": "load-balancing",
                    "peers": [
                        {
                            "name": "server1",
                            "url": "http://192.168.56.33:8005/",
                            "role": "primary"
                        }
                    ]
                }
            ]
        }
    ]
}

```

(continues on next page)

(continued from previous page)

```

    },
    {
        "name": "server2",
        "url": "http://192.168.56.66:8005/",
        "role": "secondary"
    }
],
"state-machine": {
    "states": [
        {
            "state": "ready",
            "pause": "always"
        },
        {
            "state": "partner-down",
            "pause": "once"
        }
    ]
}
}
}
}
...
}

```

This configuration instructs the server to pause the state machine every time it transitions to the `ready` state and upon the first transition to the `partner-down` state.

Refer to [Server States](#) for a complete list of server states. The state machine can be paused in any of the supported states; however, it is not practical to pause in the `backup` or `terminated` states because the server never transitions out of these states anyway.

Note

In the `syncing` state the server is paused before it makes an attempt to synchronize the lease database with a partner. To pause the state machine after lease-database synchronization, use the `ready` state instead.

Note

The state of the HA state machine depends on the state of the cooperating server. Therefore, pausing the state machine of one server may affect the operation of the partner server. For example: if the primary server is paused in the `waiting` state, the partner server will also remain in the `waiting` state until the state machine of the primary server is resumed and that server transitions to the `ready` state.

17.12.14 Control Agent Configuration

The *The Kea Control Agent* describes in detail the Kea daemon, which provides a RESTful interface to control the Kea servers. The same functionality is used by the High Availability hook library to establish communication between the HA peers. Therefore, the HA library requires that the Control Agent (CA) be started for each DHCP instance within the HA setup. If the Control Agent is not started, the peers cannot communicate with a particular DHCP server (even

if the DHCP server itself is online) and may eventually consider this server to be offline.

The following is an example configuration for the CA running on the same machine as the primary server. This configuration is valid for both the load-balancing and the hot-standby cases presented in previous sections.

```
{
  "Control-agent": {
    "http-host": "192.168.56.33",

    // If enabling HA and multi-threading, the 8000 port is used by the HA
    // hook library http listener. When using HA hook library with
    // multi-threading to function, make sure the port used by dedicated
    // listener is different (e.g. 8001) than the one used by CA. Note
    // the commands should still be sent via CA. The dedicated listener
    // is specifically for HA updates only.
    "http-port": 8000,

    "control-sockets": {
      "dhcp4": {
        "socket-type": "unix",
        "socket-name": "kea-dhcp4-ctrl.sock"
      },
      "dhcp6": {
        "socket-type": "unix",
        "socket-name": "kea-dhcp6-ctrl.sock"
      }
    }
  }
}
```

Since Kea 1.9.0, basic HTTP authentication is supported.

17.12.15 Multi-Threaded Configuration (HA+MT)

It is possible to configure HA to use direct multi-threaded communication between peers. We refer to this mode as HA+MT. With HA+MT enabled, each peer runs its own dedicated, internal HTTP listener (i.e. server) which receives and responds to commands directly, thus eliminating the need for an agent to carry out the HA protocol between peers. In addition, both the listener and client components use multi-threading to support multiple concurrent connections between peers. By eliminating the agent and executing multiple command exchanges in parallel, HA throughput between peers improves considerably over earlier versions of Kea in most situations.

The following parameters have been added to the HA configuration, to support HA+MT operation:

- **enable-multi-threading** - enables or disables multi-threading HA peer communication (HA+MT). Kea core multi-threading must be enabled for HA+MT to operate. When **false**, the server relies on *kea-ctrl-agent* for communication with its peer, and uses single-threaded HTTP client processing. The default is **true**.
- **http-dedicated-listener** - enables or disables the creation of a dedicated, internal HTTP listener through which the server receives HA messages from its peers. The internal listener replaces the role of *kea-ctrl-agent* traffic, allowing peers to send their HA commands directly to each other. The listener listens on the peer's url. When **false**, the server relies on *kea-ctrl-agent*. This parameter has been provided largely for flexibility and testing; running HA+MT without dedicated listeners enabled will substantially limit HA throughput. The default is **true**.
- **http-listener-threads** - indicates the maximum number of threads the dedicated listener should use. A value of **0** instructs the server to use the same number of threads that the Kea core is using for DHCP multi-threading. The default is **0**.

- `http-client-threads` - indicates the maximum number of threads that should be used to send HA messages to its peers. A value of 0 instructs the server to use the same number of threads that the Kea core is using for DHCP multi-threading. The default is 0.

These parameters are grouped together under a map element, `multi-threading`, as illustrated below:

```
"Dhcp4": {
  "hooks-libraries": [
    {
      "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
      "parameters": { }
    },
    {
      "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
      "parameters": {
        "high-availability": [ {
          "this-server-name": "server1",
          "multi-threading": {
            "enable-multi-threading": true,
            "http-dedicated-listener": true,
            "http-listener-threads": 4,
            "http-client-threads": 4
          }
        } ],
        "peers": [
          // This is the configuration of this server instance.
          {
            "name": "server1",
            // This specifies the URL of our server instance.
            // Since the HA+MT uses a direct connection, the
            // DHCPv4 server open its own socket. Note that it
            // must be different than the one used by the CA
            // (typically 8000). In this example, 8005 is used.
            "url": "http://192.0.2.1:8005/",
            // This server is primary. The other one must be
            // secondary.
            "role": "primary"
          },
          // This is the configuration of our HA peer.
          {
            "name": "server2",
            // This specifies the URL of our server instance.
            // Since the HA+MT uses a direct connection, the
            // DHCPv4 server open its own socket. Note that it
            // must be different than the one used by the CA
            // (typically 8000). In this example, 8005 is used.
            "url": "http://192.0.2.2:8005/",
            // The partner is a secondary. This server is a
            // primary as specified in the previous "peers"
            // entry and in "this-server-name" before that.
            "role": "secondary"
          },
          ...
        ],
        ...
      }
    }
  ],
  ...
}
```

(continues on next page)

(continued from previous page)

```

        },
        ...
    ]
    }
    },
    ...
],
...
}

```

In the example above, HA+MT is enabled with four threads for the listener and four threads for the client.

Note

It is essential to configure the ports correctly. One common mistake is to configure CA to listen on port 8000 and also configure dedicated listeners on port 8000. In such a configuration, the communication will still work over CA, but it will be slow and the DHCP server will fail to bind sockets. Administrators should ensure that dedicated listeners use a different port (8001 is a suggested alternative); if ports are misconfigured or the ports dedicated to CA are used, the performance bottlenecks caused by the single-threaded nature of CA and the sequential nature of the UNIX socket that connects CA to DHCP servers will nullify any performance gains offered by HA+MT.

17.12.16 Parked-Packet Limit

Refer to *Parked-Packet Limit* for a basic introduction to packet parking.

The HA hook library uses this mechanism. When an HA server needs to send a lease update to its peers to notify them of the change to the lease, it parks the client response until the peers acknowledge the lease update. At that point, the server unparks the response and sends it to the client. This applies to client queries that cause lease changes, such as DHCPREQUEST for DHCPv4 and Request, Renew, and Rebind for DHCPv6. It does not apply to DHCPDISCOVERs (v4) or Solicits (v6).

17.12.17 Controlled Shutdown and Maintenance of DHCP Servers

Having a pair of servers providing High Availability allows for controlled shutdown and maintenance of those servers without disrupting the DHCP service. For example, an administrator can perform an upgrade of one of the servers while the other one continues to respond to DHCP queries. When the first server is upgraded and back online, the upgrade can be performed for the second server.

A typical problem reported with early versions of the High Availability hook library was that the administrator did not have direct control over the state of the DHCP server. Shutting down one of the servers for maintenance did not necessarily cause the other server to start responding to all DHCP queries, because the failure-detection algorithm described in *Scope Transition in a Partner-Down Case* requires that the partner not respond for a configured period of time and, depending on the configuration, may also require that a number of DHCP requests not be responded to for a specified period of time. The maintenance procedure, however, requires that the administrator be able to instruct one of the servers to instantly start serving all DHCP clients, and the other server to instantly stop serving any DHCP clients, so it can be safely shut down.

The maintenance feature of the High Availability hook library addresses this situation. The `ha-maintenance-start` command allows the administrator to put the pairs of the active servers into a state in which one of them is responding to all DHCP queries and the other one is awaiting shutdown.

Suppose that the HA setup includes two active servers, `server1` and `server2`, and the latter needs to be shut down for maintenance. The administrator can send the `ha-maintenance-start` command to `server1`, as this is the server which is going to handle the DHCP traffic while the other one is offline. `server1` responds with an error if its state

or the partner's state does not allow for a maintenance shutdown: for example, if maintenance is not supported for the backup server or if the server is in the `terminated` state. Also, an error is returned if the `ha-maintenance-start` request was already sent to the other server.

Upon receiving the `ha-maintenance-start` command, `server1` sends the `ha-maintenance-notify` command to `server2` to put it in the `in-maintenance` state. If `server2` confirms, `server1` transitions to the `partner-in-maintenance` state. This is similar to the `partner-down` state, except that in the `partner-in-maintenance` state `server1` continues to send lease updates to `server2` until the administrator shuts down `server2`. `server1` now responds to all DHCP queries.

The administrator can now safely shut down `server2` in the `in-maintenance` state and perform any necessary maintenance actions. While `server2` is offline, `server1` will obviously not be able to communicate with its partner, so it will immediately transition to the `partner-down` state; it will continue to respond to all DHCP queries but will no longer send lease updates to `server2`. Restarting `server2` after the maintenance will trigger normal state negotiation, lease-database synchronization, and, ultimately, a transition to the normal load-balancing or `hot-standby` state. Maintenance can then be performed on `server1`, after sending the `ha-maintenance-start` command to `server2`.

If the `ha-maintenance-start` command was sent to the server and the server has transitioned to the `partner-in-maintenance` state, it is possible to transition both it and its partner back to their previous states to resume the normal operation of the HA pair. This is achieved by sending the `ha-maintenance-cancel` command to the server that is in the `partner-in-maintenance` state. However, if the server has already transitioned to the `partner-down` state as a result of detecting that the partner is offline, canceling the maintenance is no longer possible. In that case, it is necessary to restart the other server and allow it to complete its normal state negotiation process.

If the server has many relationships with different partners, the `ha-maintenance-start` command attempts to transition all of the relationships into the `partner-in-maintenance` state by sending `ha-maintenance-notify` to all partner servers. If this step fails for any server, an error is returned. If that happens, the `ha-maintenance-cancel` command can be used to resume normal operations and fix the issue.

17.12.18 Control Commands for High Availability

Even though the HA hook library is designed to automatically resolve issues with DHCP service interruptions by redirecting the DHCP traffic to a surviving server and synchronizing the lease database as needed, it may be useful for the administrator to have more control over both servers' behavior. In particular, it may be useful to be able to trigger lease-database synchronization on demand, or to manually set the HA scopes that are being served.

The backup server can sometimes be used to handle DHCP traffic if both active servers are down. The backup server does not perform the failover function automatically; thus, in order to use the backup server to respond to DHCP queries, the server administrator must enable this function manually.

The following sections describe commands supported by the HA hook library which are available for the administrator.

17.12.18.1 The `ha-sync` Command

The `ha-sync` command instructs the server to synchronize its local lease database with the selected peer. The server fetches all leases from the peer and updates any locally stored leases which are older than those fetched. It also creates new leases when any of those fetched do not exist in the local database. All leases that are not returned by the peer but are in the local database are preserved. The database synchronization is unidirectional; only the database on the server to which the command has been sent is updated. To synchronize the peer's database, a separate `ha-sync` command must be issued to that peer.

Database synchronization may be triggered for both active and backup server types. The `ha-sync` command has the following structure (in a DHCPv4 example):

```
{
  "command": "ha-sync",
  "service": [ "dhcp4 "],
```

(continues on next page)

(continued from previous page)

```

    "arguments": {
        "server-name": "server2",
        "max-period": 60
    }
}

```

When the server receives this command it first disables the DHCP service of the server from which it will be fetching leases, by sending the [dhcp-disable](#) command to that server. The `max-period` parameter specifies the maximum duration (in seconds) for which the DHCP service should be disabled. If the DHCP service is successfully disabled, the synchronizing server fetches leases from the remote server by issuing one or more [lease4-get-page](#) commands. When the lease-database synchronization is complete, the synchronizing server sends the [dhcp-enable](#) command to the peer to re-enable its DHCP service.

The `max-period` value should be sufficiently long to guarantee that it does not elapse before the synchronization is completed. Otherwise, the DHCP server will automatically enable its DHCP function while the synchronization is still in progress. If the DHCP server subsequently allocates any leases during the synchronization, those new (or updated) leases will not be fetched by the synchronizing server, leading to database inconsistencies.

17.12.18.2 The `ha-scopes` Command

This command allows an administrator to modify the HA scopes being served. It can be useful when recovering from a failure or when performing maintenance. Consult [Load-Balancing Configuration](#) and [Hot-Standby Configuration](#) to learn which scopes are available for the different HA modes of operation.

Warning

The `ha-scopes` command can put Kea servers into conflicting states, which can lead to unexpected behavior. Changing scopes does not automatically change the state of the server or its relationship with a partner.

For example, when we add primary scopes to the standby server it will start responding to DHCP traffic for those scopes. This can lead to a situation where both servers are responding to the same traffic, which can lead to IP address collisions.

The `ha-scopes` command has the following structure; using an example in DHCPv4.:

```

{
    "command": "ha-scopes",
    "service": [ "dhcp4" ],
    "arguments": {
        "scopes": [ "HA_server1", "HA_server2" ],
        "server-name": "server2"
    }
}

```

This command configures the server to handle traffic from both the "HA_server1" and "HA_server2" scopes. To disable all scopes specify an empty list:

```

{
    "command": "ha-scopes",
    "service": [ "dhcp4" ],
    "arguments": {
        "scopes": [ ],
        "server-name": "server2"
    }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

17.12.18.3 The `ha-continue` Command

This command is used to resume the operation of the paused HA state machine, as described in *Pausing the HA State Machine*. It takes no arguments, so the command structure is simply:

```
{
  "command": "ha-continue",
  "service": [ "dhcp4" ],
  "arguments": {
    "scopes": [ ],
    "server-name": "server2"
  }
}
```

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

17.12.18.4 The `ha-heartbeat` Command

The *Server States* section describes how the `ha-heartbeat` command is used by a pair of active HA servers to detect one partner's failure. This command, however, can also be sent by the system administrator to one or both servers to check their HA state. This allows a monitoring system to be deployed on the HA enabled servers to periodically check whether they are operational or whether any manual intervention is required. The `ha-heartbeat` command takes no arguments:

```
{
  "command": "ha-heartbeat",
  "service": [ "dhcp4" ],
  "arguments": {
    "scopes": [ ],
    "server-name": "server2"
  }
}
```

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

Upon successful communication with the server, a response similar to this should be returned:

```
{
  "result": 0,
  "text": "HA peer status returned.",
  "arguments":
    {
      "state": "partner-down",
```

(continues on next page)

(continued from previous page)

```

    "date-time": "Thu, 07 Nov 2019 08:49:37 GMT",
    "scopes": [ "server1" ],
    "unsent-update-count": 123
  }
}

```

The returned `state` value should be one of the values listed in *Server States*. In the example above, the `partner-down` state is returned, which indicates that the server which responded to the command believes that its partner is offline; thus, it is serving all DHCP requests sent to the servers. To ensure that the partner is indeed offline, the administrator should send the `ha-heartbeat` command to the second server. If sending the command fails, e.g. due to an inability to establish a TCP connection to the Control Agent, or if the Control Agent reports issues with communication with the DHCP server, it is very likely that the server is not running.

The `date-time` parameter conveys the server's notion of time.

The `unsent-update-count` value is a cumulative count of all unsent lease updates since the server was booted; its value is set to 0 when the server is started. It is never reset to 0 during the server's operation, even after the partner synchronizes the database. It is incremented by the partner sending the heartbeat response when it cannot send the lease update. For example, suppose the failure is a result of a temporary communication interruption. In that case, the partner receiving the `partner-down` heartbeat response tracks the value changes and can determine, once communication is reestablished, whether there are any new lease updates that it did not receive. If the values on both servers do not match, it is an indication that the partner should synchronize its lease database. A non-zero value itself is not an indication of any present issues with lease updates, but a constantly incrementing value is.

The typical response returned by one server when both are operational is:

```

{
  "result": 0,
  "text": "HA peer status returned.",
  "arguments":
    {
      "state": "load-balancing",
      "date-time": "Thu, 07 Nov 2019 08:49:37 GMT",
      "scopes": [ "server1" ],
      "unsent-update-count": 0
    }
}

```

In most cases, the `ha-heartbeat` command should be sent to both HA-enabled servers to verify the state of the entire HA setup. In particular, if one of the servers indicates that it is in the `load-balancing` state, it means that this server is operating as if its partner is functional. When a partner goes down, it takes some time for the surviving server to realize it. The *Scope Transition in a Partner-Down Case* section describes the algorithm which the surviving server follows before it transitions to the `partner-down` state. If the `ha-heartbeat` command is sent during the time window between the failure of one of the servers and the transition of the surviving server to the `partner-down` state, the response from the surviving server does not reflect the failure. Resending the command detects the failure once the surviving server has entered the `partner-down` state.

17.12.18.5 The `status-get` Command

`status-get` is a general-purpose command supported by several Kea daemons, not only the DHCP servers. However, when sent to a DHCP server with HA enabled, it can be used to get insight into the details of the HA-specific server status. Not only does the response contain the status information of the server receiving this command, but also the information about its partner if it is available.

The following is an example response to the `status-get` command, including the HA status of two load-balancing

servers:

```
{
  "result": 0,
  "text": "",
  "arguments": {
    "pid": 1234,
    "uptime": 3024,
    "reload": 1111,
    "high-availability": [
      {
        "ha-mode": "load-balancing",
        "ha-servers": {
          "local": {
            "role": "primary",
            "scopes": [ "server1" ],
            "server-name": "server1",
            "state": "load-balancing",
            "system-time": "2024-01-01 12:00:00"
          },
          "remote": {
            "age": 10,
            "analyzed-packets": 8,
            "clock-skew": 0,
            "communication-interrupted": true,
            "connecting-clients": 2,
            "in-touch": true,
            "last-scopes": [ "server2" ],
            "last-state": "load-balancing",
            "role": "secondary",
            "server-name": "server2",
            "system-time": "2024-01-01 12:00:00",
            "unacked-clients": 1,
            "unacked-clients-left": 2
          }
        }
      }
    ],
    "multi-threading-enabled": true,
    "thread-pool-size": 4,
    "packet-queue-size": 64,
    "packet-queue-statistics": [ 0.2, 0.1, 0.1 ],
    "sockets": {
      "status": "ready"
    }
  }
}
```

The high-availability argument is a list which currently comprises only one element.

The ha-servers map contains two structures: local and remote. The former contains the status information of the server which received the command, while the latter contains the status information known to the local server about the partner. The role of the partner server is gathered from the local configuration file, and thus should always be available. The remaining status information, such as last-scopes and last-state, is not available until the local server communicates with the remote by successfully sending the *ha-heartbeat* command. If at least one such

communication has taken place, the returned value of the `in-touch` parameter is set to `true`. By examining this value, the command's sender can determine whether the information about the remote server is reliable.

The `last-scopes` and `last-state` parameters contain information about the HA scopes served by the partner and its state. This information is gathered during the `ha-heartbeat` command exchange, so it may not be accurate if a communication problem occurs between the partners and this status information is not refreshed. In such a case, it may be useful to send the `status-get` command to the partner server directly to check its current state. The `age` parameter specifies the age of the information from the partner, in seconds.

The `communication-interrupted` boolean value indicates whether the server receiving the `status-get` command (the local server) has been unable to communicate with the partner longer than the duration specified as `max-response-delay`. In such a situation, the active servers are considered to be in the `communication-interrupted` state. At this point, the local server may start monitoring the DHCP traffic directed to the partner to see if the partner is responding to this traffic. More about the failover procedure can be found in *Load-Balancing Configuration*.

The `system-time` parameter holds the UTC time when skew between local and partner node was last calculated. It is displayed in `%Y-%m-%d %H:%M:%S` format for each active node: local, and remote, respectively. The `clock-skew` parameter is available in the `remote` map and holds the difference in seconds between the two times. Local time is subtracted from the partner's time. A positive value means that the partner is ahead, while a negative value means that the partner is behind. Both `system-time` and `clock-skew` parameters can be null if the clock skew was not calculated yet.

The `connecting-clients`, `unacked-clients`, `unacked-clients-left`, and `analyzed-packets` parameters were introduced along with the `communication-interrupted` parameter and they convey useful information about the state of the DHCP traffic monitoring in the `communication-interrupted` state. Once the server leaves the `communication-interrupted` state, these parameters are all reset to 0.

These parameters have the following meaning in the `communication-interrupted` state:

- `connecting-clients` - this is the number of different clients which have attempted to get a lease from the remote server. These clients are differentiated by their MAC address and client identifier (in DHCPv4) or DUID (in DHCPv6). This number includes "unacked" clients (for which the "secs" field or "elapsed time" value exceeded the `max-response-delay`).
- `unacked-clients` - this is the number of different clients which have been considered "unacked", i.e. the clients which have been trying to get the lease longer than the value of the "secs" field, or for which the "elapsed time" exceeded the `max-response-delay` setting.
- `unacked-clients-left` - this indicates the number of additional clients which have to be considered "unacked" before the server enters the `partner-down` state. This value decreases when the `unacked-clients` value increases. The local server enters the `partner-down` state when this value decreases to 0.
- `analyzed-packets` - this is the total number of packets directed to the partner server and analyzed by the local server since entering the communication interrupted state. It includes retransmissions from the same clients.

Monitoring these values helps to predict when the local server will enter the `partner-down` state or to understand why the server has not yet entered this state.

The `ha-mode` parameter returns the HA mode of operation selected using the `mode` parameter in the configuration file. It can hold one of the following values: `load-balancing`, `hot-standby`, or `passive-backup`.

The `status-get` response has the format described above only in the `load-balancing` and `hot-standby` modes. In the `passive-backup` mode the `remote` map is not included in the response because in this mode there is only one active server (local). The response includes no information about the status of the backup servers.

17.12.18.6 The `ha-maintenance-start` Command

This command initiates the transition of the server's partners into the `in-maintenance` state, and the transition of the server receiving the command into the `partner-in-maintenance` state, in each HA relationship. See the *Controlled Shutdown and Maintenance of DHCP Servers* section for details.

```
{
  "command": "ha-maintenance-start",
  "service": [ "dhcp4" ]
}
```

17.12.18.7 The `ha-maintenance-cancel` Command

This command is used to cancel the maintenance previously initiated using the `ha-maintenance-start` command. The server receiving this command first sends `ha-maintenance-notify`, with the `cancel` flag set to `true`, to its partners. Next, the server reverts from the `partner-in-maintenance` state to its previous state. See the *Controlled Shutdown and Maintenance of DHCP Servers* section for details.

```
{
  "command": "ha-maintenance-cancel",
  "service": [ "dhcp4" ]
}
```

17.12.18.8 The `ha-maintenance-notify` Command

This command is sent by the server receiving the `ha-maintenance-start` or the `ha-maintenance-cancel` command to its partner, to cause the partner to transition to the `in-maintenance` state or to revert from this state to a previous state. See the *Controlled Shutdown and Maintenance of DHCP Servers* section for details.

```
{
  "command": "ha-maintenance-notify",
  "service": [ "dhcp4" ],
  "arguments": {
    "cancel": false,
    "state": "ready",
    "server-name": "server2"
  }
}
```

The `state` argument informs the recipient about the state of the sending server. The recipient can instantly resume the operation of the state machine without sending a heartbeat to check the partner's state.

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

Warning

The `ha-maintenance-notify` command is not meant to be used by system administrators. It is used for internal communication between a pair of HA-enabled DHCP servers. Direct use of this command is not supported and may produce unintended consequences.

17.12.18.9 The `ha-reset` Command

This command causes the server to reset its High Availability state machine by transitioning it to the `waiting` state. A partner in the `communication-recovery` state may send this command to cause the server to synchronize its lease database. Database synchronization is required when the partner has failed to send all lease database updates after re-establishing connection after a temporary connection failure. It is also required when the `delayed-updates-limit` is exceeded, when the server is in the `communication-recovery` state.

A server administrator may send this command to reset a misbehaving state machine.

```
{
  "command": "ha-reset",
  "service": [ "dhcp4" ],
  "arguments": {
    "server-name": "server2"
  }
}
```

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

It elicits the response:

```
{
  "result": 0,
  "text": "HA state machine reset."
}
```

If the server receiving this command is already in the `waiting` state, the command has no effect.

17.12.18.10 The `ha-sync-complete-notify` Command

A server sends this command to its partner to signal that it has completed lease-database synchronization. The partner may enable its DHCP service if it can allocate new leases in its current state. The partner does not enable the DHCP service in the `partner-down` state until it sends a successful `ha-heartbeat` test to its partner server. If the connection is still unavailable, the server in the `partner-down` state enables its own DHCP service to continue responding to clients.

```
{
  "command": "ha-sync-complete-notify",
  "service": [ "dhcp4" ],
  "arguments": {
    "origin-id": 2000,
    "server-name": "server2"
  }
}
```

The optional `server-name` parameter specifies the name of one of the partners in the HA relationship that this command pertains to. This parameter can be omitted if the server receiving this command has only one HA relationship in its configuration.

The `origin-id` parameter is used to select the HA service for which the receiving server should enable the DHCP service when it receives this notification. This is the same origin the sending server used previously to disable the DHCP service before synchronization. The `origin-id` parameter deprecates the `origin` parameter used in some earlier Kea versions.

It elicits the response:


```
{
  "result": 0,
  "text": "Server successfully notified about the synchronization completion."
}
```

Warning

The `ha-sync-complete-notify` command is not meant to be used by system administrators. It is used for internal communication between a pair of HA-enabled DHCP servers. Direct use of this command is not supported and may produce unintended consequences.

17.12.19 Hub and Spoke Configuration

A hub-and-spoke arrangement, with one central server and multiple branch servers, is a common way to ensure resiliency of the DHCP servers. The branch servers are the primary servers in the `hot-standby` mode and respond to the local DHCP traffic in their respective locations. The central server acts as a standby server for each branch server. It maintains independent state machines with the branch servers, called relationships. If one of the branch servers experiences a failure, the central server can take over its DHCP traffic. In this case, one of the central server's relationships is in the `partner-down` state. The remaining relationships may still be in the `hot-standby` state and not actively respond to DHCP traffic. When the branch server becomes active again, it synchronizes the lease database with the central server and the central server becomes fully passive again. In rare cases, if multiple branch servers stop, the central server takes responsibility for all their traffic, and potentially all DHCP traffic in the entire network if every branch server is down. A simple hub-and-spoke arrangement consisting of two branch servers and one central server is shown below.



Each branch server's configuration comprises a set of subnets appropriate for the branch server; different branch servers serve different subnets. The central server's configuration includes all subnets of the branch servers, so that it can respond to DHCP traffic directed to any branch servers that fail. The subnets in the central server must be grouped into relationships as in the snippet below:

```
{
  "Dhcp6": {
    "interfaces-config": {
      "interfaces": [ "enp0s8", "enp0s9" ]
    },
    "hooks-libraries": [
      {
        "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
        "parameters": {}
      },
      {
        "library": "/usr/lib/kea/hooks/libdhcp_ha.so",

```

(continues on next page)

(continued from previous page)

```

"parameters": {
  "high-availability": [
    {
      "this-server-name": "server2",
      "mode": "hot-standby",
      "multi-threading": {
        "enable-multi-threading": true,
        "http-dedicated-listener": true,
        "http-listener-threads": 4,
        "http-client-threads": 4
      },
      "peers": [
        {
          "name": "server1",
          "url": "http://192.168.56.66:8007/",
          "role": "primary",
          "auto-failover": true
        },
        {
          "name": "server2",
          "url": "http://192.168.56.33:8007/",
          "role": "standby",
          "auto-failover": true
        }
      ]
    },
    {
      "this-server-name": "server4",
      "mode": "hot-standby",
      "multi-threading": {
        "enable-multi-threading": true,
        "http-dedicated-listener": true,
        "http-listener-threads": 4,
        "http-client-threads": 4
      },
      "peers": [
        {
          "name": "server3",
          "url": "http://192.168.57.99:8007/",
          "role": "primary",
          "auto-failover": true
        },
        {
          "name": "server4",
          "url": "http://192.168.57.33:8007/",
          "role": "standby",
          "auto-failover": true
        }
      ]
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "subnet6": [
    {
      "id": 1,
      "subnet": "2001:db8:1::/64",
      "pools": [ { "pool": "2001:db8:1::/80" } ],
      "interface": "enp0s8",
      "user-context": {
        "ha-server-name": "server2"
      }
    },
    {
      "id": 2,
      "subnet": "2001:db8:2::/64",
      "pools": [ { "pool": "2001:db8:2::/80" } ],
      "interface": "enp0s9",
      "user-context": {
        "ha-server-name": "server4"
      }
    }
  ]
}

```

The peer names in the relationships must be unique. The user context for each subnet contains the `ha-server-name` parameter, associating a subnet with a relationship. The `ha-server-name` can be any of the peer names in the relationship. For example, if a relationship contains peer names `server1` and `server2`, it does not matter whether the `ha-server-name` is `server1` or `server2`. In both cases, it associates a subnet with that relationship.

It is not required to specify the `ha-server-name` in the branch servers, assuming that the branch servers only contain the subnets they serve. Consider the following configuration for branch server3:

```

{
  "Dhcp6": {
    "interfaces-config": {
      "interfaces": [ "enp0s8" ]
    },
    "hooks-libraries": [
      {
        "library": "/usr/lib/kea/hooks/libdhcp_lease_cmds.so",
        "parameters": {}
      },
      {
        "library": "/usr/lib/kea/hooks/libdhcp_ha.so",
        "parameters": {
          "high-availability": [
            {
              "this-server-name": "server3",
              "mode": "hot-standby",
              "multi-threading": {
                "enable-multi-threading": true,
                "http-dedicated-listener": true,

```

(continues on next page)

(continued from previous page)

```

        "http-listener-threads": 4,
        "http-client-threads": 4
    },
    "peers": [
        {
            "name": "server3",
            "url": "http://192.168.57.99:8007/",
            "role": "primary",
            "auto-failover": true
        },
        {
            "name": "server4",
            "url": "http://192.168.57.33:8007/",
            "role": "standby",
            "auto-failover": true
        }
    ]
}
],
"subnet6": [
    {
        "id": 2,
        "subnet": "2001:db8:2::/64",
        "pools": [
            {
                "pool": "2001:db8:2::/80"
            }
        ],
        "interface": "enp0s8",
        "user-context": {
            "ha-server-name": "server3"
        }
    }
]
}
}

```

Note

Even though it is possible to omit the `ha-server-name` user context parameters in the branch servers, we recommend including them. The servers fetch all leases from the partners during the database synchronization; if the subnets are not explicitly associated with the relationship, the branch server inserts all fetched leases from the central server (including those from other relationships) into its database. Specifying the `ha-server-name` parameter for each configured subnet in the branch server guarantees that only the leases belonging to its relationship are inserted into the branch server's database.

Note

The peer names in the branch servers must match the peer names in the respective central server's relationships, because these names are used for signaling between the HA partners.

Note

The hub-and-spoke configuration is currently incompatible with the `reselect-subnet-address` and `reselect-subnet-pool` parameters of the `libdhcp_radius.so`. The High Availability hook library uses an originally selected subnet for choosing an HA relationship to process a packet. The subnet reselection may interfere with this choice. See the *RADIUS Hook Library Configuration* for details.

17.13 libdhcp_host_cache.so: Host Cache Reservations for Improved Performance

Some database backends, such as RADIUS, are slow and may take a long time to respond. Since Kea in general is synchronous, backend performance directly affects DHCP performance. To minimize the impact and improve performance, the Host Cache library provides a way to cache information from the database locally. This includes negative caching, i.e. the ability to remember that there is no client information in the database.

Note

`libdhcp_host_cache.so` is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the `kea-dhcp4` or `kea-dhcp6` process.

In principle, this hook library can be used with any backend that may introduce performance degradation (MySQL, PostgreSQL or RADIUS). Host Cache must be loaded for the RADIUS accounting mechanism to work.

The Host Cache hook library is very simple. It takes only one optional parameter (`maximum`), which defines the maximum number of hosts to be cached. If not specified, the default value of 0 is used, which means there is no limit. This hook library can be loaded the same way as any other hook library; for example, this configuration could be used:

```
"Dhcp4": {
# Your regular DHCPv4 configuration parameters here.

"hooks-libraries": [
{
  "library": "/usr/local/lib/kea/hooks/libdhcp_host_cache.so",
  "parameters": {
    # Tells Kea to never cache more than 1000 hosts.
    "maximum": 1000
  }
}
]
```

(continues on next page)

(continued from previous page)

```

    }
} ],
...
}

```

Once loaded, the Host Cache hook library provides a number of new commands which can be used either over the control channel (see *Using the Control Channel*) or the RESTful API (see *Overview of the Kea Control Agent*). An example RESTful API client is described in *Overview of the Kea Shell*. The following sections describe the commands available.

17.13.1 The cache-flush Command

This command allows removal of a specified number of cached host entries. It takes one parameter, which defines the number of hosts to be removed. An example usage looks as follows:

```

{
  "command": "cache-flush",
  "arguments": 1000
}

```

This command removes 1000 hosts; to delete *all* cached hosts, use *cache-clear* instead. The hosts are stored in FIFO (first-in, first-out) order, so the oldest entries are always removed.

17.13.2 The cache-clear Command

This command allows removal of all cached host entries. An example usage looks as follows:

```

{
  "command": "cache-clear"
}

```

This command removes all hosts. To delete only a certain number of cached hosts, please use *cache-flush* instead.

17.13.3 The cache-size Command

This command returns the number of host entries. An example usage looks as follows:

```

{
  "command": "cache-size"
}

```

17.13.4 The cache-write Command

In general, the cache content is considered a runtime state and the server can be shut down or restarted as usual; the cache is then repopulated after restart. However, there are some cases when it is useful to store the contents of the cache. One such case is RADIUS, where the cached hosts also retain additional cached RADIUS attributes; there is no easy way to obtain this information again, because renewing clients send their packet to the DHCP server directly. Another use case is when an administrator wants to restart the server and, for performance reasons, wants it to start with a hot (populated) cache.

This command allows writing the contents of the in-memory cache to a file on disk. It takes one parameter, which defines the filename. An example usage looks as follows:

```
{
  "command": "cache-write",
  "arguments": "/usr/local/var/lib/kea/kea-host-cache.json"
}
```

This causes the contents to be stored in the `/usr/local/var/lib/kea/kea-host-cache.json` file. That file can then be loaded with the `cache-load` command or processed by any other tool that is able to understand JSON format.

Note

As of Kea 2.7.9, the cache file may only be written to the data directory determined during compilation: `[kea-install-dir]/var/lib/kea`. This path may be overridden at startup by setting the environment variable `KEA_DHCP_DATA_DIR` to the desired path. For ease of use in specifying a custom file name simply omit the path portion from filename.

17.13.5 The `cache-load` Command

See the previous section for a discussion of use cases where it may be useful to write and load contents of the host cache to disk.

This command allows the contents of a file on disk to be loaded into an in-memory cache. It takes one parameter, which defines the filename. An example usage looks as follows:

```
{
  "command": "cache-load",
  "arguments": "/usr/local/share/kea/kea-host-cache.json"
}
```

This command stores the contents to the `/usr/local/share/kea/kea-host-cache.json` file. That file can then be loaded with the `cache-load` command or processed by any other tool that is able to understand JSON format.

17.13.6 The `cache-get` Command

This command is similar to `cache-write`, but instead of writing the cache contents to disk, it returns the contents to whoever sent the command.

This command allows the contents of a file on disk to be loaded into an in-memory cache. It takes one parameter, which defines the filename. An example usage looks as follows:

```
{
  "command": "cache-get"
}
```

This command returns all the cached hosts; the response may be large.

17.13.7 The `cache-get-by-id` Command

This command is similar to `cache-get`, but instead of returning the whole content it returns only the entries matching the given identifier.

It takes one parameter, which defines the identifier of wanted cached host reservations. An example usage looks as follows:

```
{
  "command": "cache-get-by-id",
  "arguments": {
    "hw-address": "01:02:03:04:05:06"
  }
}
```

This command returns all the cached hosts with the given hardware address.

17.13.8 The cache-insert Command

This command may be used to manually insert a host into the cache; there are very few use cases when this command might be useful. This command expects its arguments to follow the usual syntax for specifying host reservations (see *Host Reservations in DHCPv4* or *Host Reservations in DHCPv6*), with one difference: the `subnet-id` value must be explicitly specified.

An example command to insert an IPv4 host into the host cache looks as follows:

```
{
  "command": "cache-insert",
  "arguments": {
    "hw-address": "01:02:03:04:05:06",
    "subnet-id4": 4,
    "subnet-id6": 0,
    "ip-address": "192.0.2.100",
    "hostname": "somehost.example.org",
    "client-classes4": [ ],
    "client-classes6": [ ],
    "option-data4": [ ],
    "option-data6": [ ],
    "next-server": "192.0.0.2",
    "server-hostname": "server-hostname.example.org",
    "boot-file-name": "bootfile.efi",
    "host-id": 0
  }
}
```

An example command to insert an IPv6 host into the host cache looks as follows:

```
{
  "command": "cache-insert",
  "arguments": {
    "hw-address": "01:02:03:04:05:06",
    "subnet-id4": 0,
    "subnet-id6": 6,
    "ip-addresses": [ "2001:db8::cafe:babe" ],
    "prefixes": [ "2001:db8:dead:beef::/64" ],
    "hostname": "",
    "client-classes4": [ ],
    "client-classes6": [ ],
    "option-data4": [ ],
    "option-data6": [ ],
    "next-server": "0.0.0.0",
    "server-hostname": ""
  }
}
```

(continues on next page)

(continued from previous page)

```

    "boot-file-name": "",
    "host-id": 0
  }
}

```

17.13.9 The `cache-remove` Command

Sometimes it is useful to remove a single entry from the host cache: for example, consider a situation where the device is active, Kea has already provided configuration, and the host entry is in cache. As a result of administrative action (e.g. the customer hasn't paid their bills or has been upgraded to better service), the information in the backend database (e.g. MySQL or RADIUS) is being updated. However, since the cache is in use, Kea does not notice the change as the cached values are used. The `cache-remove` command can solve this problem by removing a cached entry after administrative changes.

The `cache-remove` command works similarly to the `reservation-get` command. It allows querying by two parameters: either `subnet-id4` or `subnet-id6`; or `ip-address` (may be an IPv4 or IPv6 address), `hw-address` (specifies a hardware/MAC address), `duid`, `circuit-id`, `client-id`, or `flex-id`.

An example command to remove an IPv4 host with reserved address 192.0.2.1 from a subnet with a `subnet-id` 123 looks as follows:

```

{
  "command": "cache-remove",
  "arguments": {
    "ip-address": "192.0.2.1",
    "subnet-id": 123
  }
}

```

Another example that removes an IPv6 host identifier by DUID and specific `subnet-id` is:

```

{
  "command": "cache-remove",
  "arguments": {
    "duid": "00:01:ab:cd:f0:a1:c2:d3:e4",
    "subnet-id": 123
  }
}

```

17.14 `libdhcp_host_cmds.so`: Host Commands

Kea can store host reservations in a database; in many larger deployments, it is useful to be able to manage that information while the server is running. The Host Commands library provides management commands for adding, querying, and deleting host reservations in a safe way without restarting the server. In particular, it validates the parameters, so an attempt to insert incorrect data - such as adding a host with a conflicting identifier in the same subnet - is rejected. Those commands are exposed via the command channel (JSON over UNIX sockets) and the Control Agent (JSON over a RESTful interface).

Note

`libdhcp_host_cmds.so` is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

Currently, the following commands are supported:

- *reservation-add*, which adds a new host reservation.
- *reservation-get*, which returns an existing reservation if specified criteria are matched.
- *reservation-get-all*, which returns all reservations in a specified subnet.
- *reservation-get-page*, a variant of *reservation-get-all* that returns reservations by pages, either all or in a specified subnet.
- *reservation-get-by-address*, which returns all reservations with a specified IP address or a delegated prefix (without a prefix length), and optionally a subnet id.
- *reservation-get-by-hostname*, which returns all reservations with a specified hostname and optionally in a subnet.
- *reservation-get-by-id*, which returns all reservations with a specified identifier (since Kea version 1.9.0).
- *reservation-del*, which attempts to delete a reservation matching specified criteria.
- *reservation-update*, which updates (replaces) an existing reservation matching the given identifiers in a subnet.

To use the commands that change reservation information (i.e. *reservation-add*, *reservation-del*, and *reservation-update*) to modify data stored in the host database, the hosts database must be specified and it must not operate in read-only mode (for details, see the *hosts-databases* descriptions in *DHCPv4 Hosts Database Configuration* and *DHCPv6 Hosts Database Configuration*). If the *hosts-databases* are not specified or are running in read-only mode, *libdhcp_host_cmds.so* will load, but any attempts to use *reservation-add*, *reservation-del*, and *reservation-update* to modify data in that database will fail.

These commands can also modify hosts from the JSON configuration independently from the hosts database. The changes provided in the JSON configuration are volatile and can be made permanent by sending the *config-write* command.

For a description of proposed future commands, see the [Control API Requirements](#) document.

All host commands use JSON syntax. They can be issued either using the control channel (see *Management API*) or via the Control Agent (see *The Kea Control Agent*).

The library can be loaded similarly to other hook libraries. It does not take any parameters, and it supports both the DHCPv4 and DHCPv6 servers.

```
"Dhcp6": {
  "hooks-libraries": [
    {
      "library": "/path/libdhcp_host_cmds.so"
    },
    ...
  ]
}
```

17.14.1 The subnet-id Parameter

Before examining the individual commands, it is worth discussing the parameter `subnet-id`. Currently, this parameter is mandatory for all of the commands supplied by this library, with the exception of [reservation-get-by-hostname](#) and [reservation-get-by-address](#), where it is optional. `subnet-id` is also optional in [reservation-get-page](#), and it is forbidden in [reservation-get-by-id](#).

Reservations can be specified globally, and are not necessarily specific to any subnet. When reservations are supplied via the configuration file, the ID of the containing subnet (or lack thereof) is implicit in the configuration structure. However, when managing reservations using host commands, it is necessary to explicitly identify the scope to which the reservation belongs. This is done via the `subnet-id` parameter. For global reservations, use a value of zero (0). For reservations scoped to a specific subnet, use that subnet's ID. See also the note in [The reservation-add Command](#) about the [reservation-add](#) command for an empty (i.e. with no reserved address or prefix) reservation.

On the other hand, when the `subnet-id` is not specified in the command parameters, it is added to each host in responses. If the `subnet-id` has the unused special value, this means the host entry belongs only to the other IP version (i.e. IPv6 in the DHCPv4 server or IPv4 in the DHCPv6 server) and this entry is ignored.

17.14.2 The operation-target Parameter

Most host commands accept the `operation-target` parameter, which selects the host data source. The commands may process data from the server configuration (i.e., memory operation target), a database (database target), or both (all sources). The `operation-target` parameter is optional. By default, the commands that only read the data use all data sources (memory and database); the commands that modify the state (i.e. [reservation-add](#), [reservation-del](#), and [reservation-update](#)) only use the database target.

The `operation-target` parameter accepts the following values:

- `memory` - query or update the runtime server configuration
- `database` - query or update host database(s)
- `all` - query or update both runtime configuration and host database(s)
- `default` - query or update a default host data source - it is command-specific

17.14.3 The reservation-add Command

[reservation-add](#) allows for the insertion of a new host. It takes a set of arguments that vary depending on the nature of the host reservation. Any parameters allowed in the configuration file that pertain to host reservation are permitted here. For details regarding IPv4 reservations, see [Host Reservations in DHCPv4](#); for IPv6 reservations, see [Host Reservations in DHCPv6](#). The `subnet-id` is mandatory. Use a value of zero (0) to add a global reservation, or the ID of the subnet to which the reservation should be added. The command can be as simple as having only the two mandatory entries:

```
{
  "command": "reservation-add",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "hw-address": "1a:1b:1c:1d:1e:1f"
    }
  }
}
```

In that case, however, it does not assign any resources to the host. An IPv4 address can be assigned as in this example:

```
{
  "command": "reservation-add",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "hw-address": "1a:1b:1c:1d:1e:1f",
      "ip-address": "192.0.2.202"
    }
  }
}
```

It can also take many more parameters, such as:

```
{
  "command": "reservation-add",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "client-id": "01:0a:0b:0c:0d:0e:0f",
      "ip-address": "192.0.2.205",
      "next-server": "192.0.2.1",
      "server-hostname": "hal9000",
      "boot-file-name": "/dev/null",
      "option-data": [
        {
          "name": "domain-name-servers",
          "data": "10.1.1.202,10.1.1.203"
        }
      ],
      "client-classes": [ "special_snowflake", "office" ]
    }
  }
}
```

Here is an example of a complex IPv6 reservation:

```
{
  "command": "reservation-add",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "duid": "01:02:03:04:05:06:07:08:09:0A",
      "ip-addresses": [ "2001:db8:1:cafe::1" ],
      "prefixes": [ "2001:db8:2:abcd::/64" ],
      "hostname": "foo.example.com",
      "option-data": [
        {
          "name": "vendor-opts",
          "data": "4491"
        },
        {
          "name": "tftp-servers",
          "space": "vendor-4491",

```

(continues on next page)

(continued from previous page)

```

        "data": "3000:1::234"
    }
}

```

The command accepts the `operation-target` argument. By default, it adds the reservation to the hosts database only.

```

{
  "command": "reservation-add",
  "arguments": {
    "reservation": {},
    "operation-target": "all"
  }
}

```

The command returns a status that indicates either success (result 0) or failure (result 1). A failed command always includes a text parameter that explains the cause of the failure. Here's an example of a successful addition:

```

{
  "result": 0,
  "text": "Host added."
}

```

And here's an example of a failure:

```

{
  "result": 1,
  "text": "Mandatory 'subnet-id' parameter missing."
}

```

As `reservation-add` is expected to store the host, the `hosts-databases` parameter must be specified in the configuration, and databases must not run in read-only mode.

Note

Since kea version 2.7.0 an empty reservation is implicitly global, i.e. if no `subnet-id` entry is present and for DHCPv4 `ip-address` or for DHCPv6 `ip-addresses` and `prefixes` do not specify a reserved address or prefix the command is valid and adds a global (`subnet-id` 0) host reservation.

17.14.4 The reservation-get Command

`reservation-get` can be used to query the host database and retrieve existing reservations. This command supports two types of parameters: (`subnet-id`, `address`) or (`subnet-id`, `identifier-type`, `identifier`). The first type of query is used when the address (either IPv4 or IPv6) is known, but the details of the reservation are not. One common use for this type of query is to find out whether a given address is reserved. The second query uses identifiers. For maximum flexibility, Kea stores the host identifying information as a pair of values: the type and the actual identifier. Currently supported identifiers are `"hw-address"`, `"duid"`, `"circuit-id"`, `"client-id"`, and `"flex-id"`. The `subnet-id` is mandatory. Use a value of zero (0) to fetch a global reservation, or the ID of the subnet to which the reservation belongs.

An example command for getting a host reservation by a (subnet-id, address) pair looks as follows:

```
{
  "command": "reservation-get",
  "arguments": {
    "subnet-id": 1,
    "ip-address": "192.0.2.202"
  }
}
```

An example query by (subnet-id, identifier-type, identifier) looks as follows:

```
{
  "command": "reservation-get",
  "arguments": {
    "subnet-id": 4,
    "identifier-type": "hw-address",
    "identifier": "01:02:03:04:05:06"
  }
}
```

The command accepts the `operation-target` argument. By default, it gets the reservation from both the JSON configuration and the hosts database.

```
{
  "command": "reservation-get",
  "arguments": {
    "subnet-id": 1,
    "ip-address": "192.0.2.202",
    "operation-target": "database"
  }
}
```

Command `reservation-get` typically returns the result 0 when a query was conducted properly. In particular, 0 is returned when the host was not found. If the query was successful, the host parameters are returned. An example of a query that did not find the host looks as follows:

```
{ "result": 0, "text": "Host not found." }
```

Here's an example of a result returned when the host was found successfully:

```
{
  "arguments": {
    "boot-file-name": "bootfile.efi",
    "client-classes": [

    ],
    "hostname": "somehost.example.org",
    "hw-address": "01:02:03:04:05:06",
    "ip-address": "192.0.2.100",
    "next-server": "192.0.0.2",
    "option-data": [

    ],
  },
}
```

(continues on next page)

(continued from previous page)

```

    "server-hostname": "server-hostname.example.org",
    "subnet-id": 4
  },
  "result": 0,
  "text": "Host found."
}

```

An example result returned when the query was malformed might look like this:

```

{ "result": 1, "text": "No 'ip-address' provided and 'identifier-type' is either missing,
↳or not a string." }

```

17.14.5 The reservation-get-all Command

reservation-get-all can be used to query the host database and retrieve all reservations in a specified subnet. This command uses parameters providing the mandatory `subnet-id`. Global host reservations can be retrieved by using a `subnet-id` value of zero (0).

For instance, retrieving host reservations for the subnet 1:

```

{
  "command": "reservation-get-all",
  "arguments": {
    "subnet-id": 1
  }
}

```

returns some IPv4 hosts:

```

{
  "arguments": {
    "hosts": [
      {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "somehost.example.org",
        "hw-address": "01:02:03:04:05:06",
        "ip-address": "192.0.2.100",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org",
        "subnet-id": 1
      },
      {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "otherhost.example.org",
        "hw-address": "01:02:03:04:05:ff",
        "ip-address": "192.0.2.200",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org",
        "subnet-id": 1
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

        },
        ...
    ]
},
"result": 0,
"text": "72 IPv4 host(s) found."
}

```

The response returned by *reservation-get-all* can be very long. The DHCP server does not handle DHCP traffic while preparing a response to *reservation-get-all*, so if there are many reservations in a subnet, this may be disruptive; use with caution. For larger deployments, please consider using *reservation-get-page* instead.

The command accepts the *operation-target* argument. By default, it gets the reservation from both the JSON configuration and the hosts database.

```

{
  "command": "reservation-get-all",
  "arguments": {
    "subnet-id": 1,
    "operation-target": "database"
  }
}

```

For more information, see *The reservation-get-all Command*.

17.14.6 The reservation-get-page Command

reservation-get-page can be used to query the host database and retrieve all reservations in a specified subnet, by pages. This command uses parameters providing the mandatory *subnet-id*. Use a value of zero (0) to fetch global reservations. The second mandatory parameter is the page size limit. The optional *source-index* and *from-host-id* parameters, both of which default to 0, are used to chain page queries. Since Kea version 1.9.0, the *subnet-id* parameter is optional.

The usage of the *from* and *source-index* parameters requires additional explanation. For the first call, those parameters should not be specified (or should be specified as zeros). For any follow-up calls, they should be set to the values returned in previous calls, in a next map holding *from* and *source-index* values. Subsequent calls should be issued until all reservations are returned. The end is reached once the returned list is empty, the count is 0, no next map is present, and result status 3 (empty) is returned.

Note

The *from* and *source-index* parameters reflect the internal state of the search. There is no need to understand what they represent; it is simply a value that should be copied from one response to the next query. However, for those who are curious, the *from* field represents a 64-bit representation of the host identifier used by a host backend. The *source-index* is an internal representation of multiple host backends: 0 is used to represent hosts defined in a configuration file, and 1 represents the first database backend. In some uncommon cases there may be more than one database backend configured, so potentially there may be a 2. In any case, Kea iterates over all backends configured.

For instance, retrieving host reservations for the subnet 1 and requesting the first page can be done by:

```

{
  "command": "reservation-get-page",

```

(continues on next page)

(continued from previous page)

```

    "arguments": {
        "subnet-id": 1,
        "limit": 10
    }
}

```

Since this is the first call, `source-index` and `from` should not be specified. They are set to their zero default values. Some hosts are returned with information to get the next page:

```

{
    "arguments": {
        "count": 72,
        "hosts": [
            {
                "boot-file-name": "bootfile.efi",
                "client-classes": [ ],
                "hostname": "somehost.example.org",
                "hw-address": "01:02:03:04:05:06",
                "ip-address": "192.0.2.100",
                "next-server": "192.0.0.2",
                "option-data": [ ],
                "server-hostname": "server-hostname.example.org"
            },
            {
                "boot-file-name": "bootfile.efi",
                "client-classes": [ ],
                "hostname": "otherhost.example.org",
                "hw-address": "01:02:03:04:05:ff",
                "ip-address": "192.0.2.200",
                "next-server": "192.0.0.2",
                "option-data": [ ],
                "server-hostname": "server-hostname.example.org"
            },
            ...
        ],
        "next": {
            "from": 1234567,
            "source-index": 1
        },
        "result": 0,
        "text": "72 IPv4 host(s) found."
    }
}

```

Note that the `from` and `source-index` fields were specified in the response in the next map. Those two must be copied to the next command, so Kea continues from the place where the last command finished. To get the next page the following command can be sent:

```

{
    "command": "reservation-get-page",
    "arguments": {
        "subnet-id": 1,

```

(continues on next page)

(continued from previous page)

```

    "source-index": 1,
    "from": 1234567,
    "limit": 10
  }
}

```

The response will contain a list of hosts with updated `source-index` and `from` fields. Continue calling the command until the last page is received. Its response will look like this:

```

{
  "arguments": {
    "count": 0,
    "hosts": [ ]
  },
  "result": 3,
  "text": "0 IPv4 host(s) found."
}

```

The command does not accept the `operation-target` argument.

This command is more complex than *reservation-get-all*, but it lets users retrieve larger host reservations lists in smaller chunks. For small deployments with few reservations, it is easier to use *reservation-get-all*.

17.14.7 The *reservation-get-by-address* Command

reservation-get-by-address can be used to query the host database and retrieve all reservations for a given IP address or a delegated prefix (without a prefix length), in a specified subnet or in all subnets. This command uses parameters providing the mandatory `ip-address` and the optional `subnet-id` and `operation-target`.

For instance, retrieving host reservations for the IPv4 address "192.0.200.181" in the subnet 1:

```

{
  "command": "reservation-get-by-address",
  "arguments": {
    "ip-address": "192.0.200.181",
    "subnet-id": 1
  },
  "service": [
    "dhcp4"
  ]
}

```

can return two IPv4 hosts:

```

{
  "arguments": {
    "hosts": [
      {
        "boot-file-name": "",
        "client-classes": [],
        "hostname": "",
        "hw-address": "99:99:99:99:99:01",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",

```

(continues on next page)

(continued from previous page)

```

        "option-data": [],
        "server-hostname": "",
        "subnet-id": 1
    },
    {
        "boot-file-name": "",
        "circuit-id": "1234",
        "client-classes": [],
        "hostname": "",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",
        "option-data": [],
        "server-hostname": "",
        "subnet-id": 1
    }
]
},
"result": 0,
"text": "2 IPv4 host(s) found."
}

```

To search for all reservations in all subnets, simply skip the `subnet-id` parameter:

```

{
  "command": "reservation-get-by-address",
  "arguments": {
    "ip-address": "192.0.200.181"
  },
  "service": [
    "dhcp4"
  ]
}

```

An example response could be:

```

{
  "arguments": {
    "hosts": [
      {
        "boot-file-name": "",
        "client-classes": [],
        "hostname": "",
        "hw-address": "99:99:99:99:99:01",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",
        "option-data": [],
        "server-hostname": "",
        "subnet-id": 1
      },
      {
        "boot-file-name": "",
        "circuit-id": "1234",

```

(continues on next page)

(continued from previous page)

```

        "client-classes": [],
        "hostname": "",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",
        "option-data": [],
        "server-hostname": "",
        "subnet-id": 1
    },
    {
        "boot-file-name": "",
        "client-classes": [],
        "hostname": "",
        "hw-address": "99:99:99:99:99:02",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",
        "option-data": [],
        "server-hostname": "",
        "subnet-id": 0
    },
    {
        "boot-file-name": "",
        "client-classes": [],
        "hostname": "",
        "hw-address": "99:99:99:99:99:03",
        "ip-address": "192.0.200.181",
        "next-server": "0.0.0.0",
        "option-data": [],
        "server-hostname": "",
        "subnet-id": 2
    }
]
},
"result": 0,
"text": "4 IPv4 host(s) found."
}

```

When using the command for retrieving DHCP6 host reservations, one can provide either an IPv6 address or a delegated prefix as the `ip-address` parameter. Please notice that this command does not take prefix length as a parameter in the current implementation. Thus, searching by an IP address `2001:db8:2:cafe::` can return host reservations with delegated prefixes of different lengths matching this search. For example: `2001:db8:2:cafe::/63`, `2001:db8:2:cafe::/64`, etc. Please consider the example below:

```

{
    "command": "reservation-get-by-address",
    "arguments": {
        "ip-address": "2001:db8:2:cafa::"
    },
    "service": [
        "dhcp6"
    ]
}

```

Response:

```
{
  "arguments": {
    "hosts": [
      {
        "client-classes": [],
        "duid": "01:02:03:04:05:06:07:88:98:fa",
        "hostname": "foo.example.com",
        "ip-addresses": [
          "2001:db8:1:cafe::2"
        ],
        "option-data": [],
        "prefixes": [
          "2001:db8:2:abcd::/64",
          "2001:db8:2:cafa::/63"
        ],
        "subnet-id": 8
      },
      {
        "client-classes": [],
        "duid": "01:02:03:04:05:06:07:88:98:fb",
        "hostname": "foo.example.com",
        "ip-addresses": [
          "2001:db8:1:cafe::2"
        ],
        "option-data": [],
        "prefixes": [
          "2001:db8:2:abcd::/64",
          "2001:db8:2:cafa::/64"
        ],
        "subnet-id": 8
      }
    ]
  },
  "result": 0,
  "text": "2 IPv6 host(s) found."
}
```

The command accepts the `operation-target` argument. By default, it gets the reservation from both the JSON configuration and the hosts database.

```
{
  "command": "reservation-get-by-address",
  "arguments": {
    "ip-address": "192.0.200.181",
    "subnet-id": 1,
    "operation-target": "database"
  },
  "service": [
    "dhcp4"
  ]
}
```

Note

This command is useful in specific cases. By default, having more than one host reservation for the same IP address in a given subnet is not allowed, as explained in *Multiple Reservations for the Same IPv4* or *Multiple Reservations for the Same IPv6*. That's why this command comes in handy when the `ip-reservations-unique` boolean parameter is set to `false`. Another use case of this command is having overlapping subnets with the same IP address reservations (but with different identifiers) in different subnets.

17.14.8 The `reservation-get-by-hostname` Command

`reservation-get-by-hostname` can be used to query the host database and retrieve all reservations with a specified hostname or in a specified subnet. This command uses parameters providing the mandatory `hostname` and the optional `subnet-id`. Global host reservations can be retrieved by using a `subnet-id` value of zero (0). Hostname matching is case-insensitive.

For instance, retrieving host reservations for "foobar" in the subnet 1:

```
{
  "command": "reservation-get-by-hostname",
  "arguments": {
    "hostname": "foobar.example.org",
    "subnet-id": 1
  }
}
```

returns some IPv4 hosts:

```
{
  "arguments": {
    "hosts": [
      {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "foobar.example.org",
        "hw-address": "01:02:03:04:05:06",
        "ip-address": "192.0.2.100",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org"
      },
      {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "foobar.example.org",
        "hw-address": "01:02:03:04:05:ff",
        "ip-address": "192.0.2.200",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org"
      },
      ...
    ]
  },
}
```

(continues on next page)

(continued from previous page)

```

    "result": 0,
    "text": "5 IPv4 host(s) found."
}

```

The response returned by *reservation-get-by-hostname* can be long, particularly when responses are not limited to a subnet.

The command accepts the *operation-target* argument. By default, it gets the reservation from both the JSON configuration and the hosts database.

```

{
  "command": "reservation-get-by-hostname",
  "arguments": {
    "hostname": "foobar.example.org",
    "subnet-id": 1,
    "operation-target": "database"
  }
}

```

For more information, see *The reservation-get-by-hostname Command*.

Note

When using MySQL as the host backend, this command relies on the fact that the *hostname* column in the *hosts* table uses a case-insensitive collation, as explained in the *MySQL* section of *Kea Database Administration*.

17.14.9 The *reservation-get-by-id* Command

reservation-get-by-id can be used to query the host database and retrieve all reservations with a specified identifier (*identifier-type* and *identifier* parameters), independently of subnets. The syntax for parameters is the same as for *reservation-get*. The *subnet-id* parameter cannot be used, to avoid confusion. This command is available since Kea version 1.9.0.

For instance, retrieving host reservations for the 01:02:03:04:05:06 MAC address:

```

{
  "command": "reservation-get-by-id",
  "arguments": {
    "identifier-type": "hw-address",
    "identifier": "01:02:03:04:05:06"
  }
}

```

returns some IPv4 hosts:

```

{
  "arguments": {
    "hosts": [
      {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "foo.example.org",
        "hw-address": "01:02:03:04:05:06",

```

(continues on next page)

(continued from previous page)

```

        "ip-address": "192.0.2.100",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org",
        "subnet-id": 123
    },
    {
        "boot-file-name": "bootfile.efi",
        "client-classes": [ ],
        "hostname": "bar.example.org",
        "hw-address": "01:02:03:04:05:06",
        "ip-address": "192.0.2.200",
        "next-server": "192.0.0.2",
        "option-data": [ ],
        "server-hostname": "server-hostname.example.org",
        "subnet-id": 345
    },
    ...
]
},
"result": 0,
"text": "5 IPv4 host(s) found."
}

```

The response returned by *reservation-get-by-id* can be long, particularly when responses are not limited to a subnet.

The command accepts the *operation-target* argument. By default, it gets the reservation from both the JSON configuration and the hosts database.

```

{
  "command": "reservation-get-by-id",
  "arguments": {
    "identifier-type": "hw-address",
    "identifier": "01:02:03:04:05:06",
    "operation-target": "database"
  }
}

```

For more information, see *The reservation-get-by-id Command*.

17.14.10 The reservation-del Command

reservation-del can be used to delete a reservation from the host database and/or JSON configuration. This command supports two types of parameters: (subnet-id, address) or (subnet-id, identifier-type, identifier). The first type of query is used when the address (either IPv4 or IPv6) is known, but the details of the reservation are not. One common use for this type of query is to remove a reservation (e.g. a specific address should no longer be reserved). The second query uses identifiers. For maximum flexibility, Kea stores the host identifying information as a pair of values: the type and the actual identifier. Currently supported identifiers are "hw-address", "duid", "circuit-id", "client-id", and "flex-id". The subnet-id is mandatory. Use a value of zero (0) to delete a global reservation, or the ID of the subnet from which the reservation should be deleted.

An example command for deleting a host reservation by (subnet-id, address) pair looks as follows:


```
{
  "command": "reservation-del",
  "arguments": {
    "subnet-id": 1,
    "ip-address": "192.0.2.202"
  }
}
```

An example deletion by (subnet-id, identifier-type, identifier) looks as follows:

```
{
  "command": "reservation-del",
  "arguments": {
    "subnet-id": 4,
    "identifier-type": "hw-address",
    "identifier": "01:02:03:04:05:06"
  }
}
```

Command *reservation-del* returns a result of 0 when the host deletion was successful, or 1 if it failed. Descriptive text is provided in the event of an error. Here are some examples of possible results:

```
{
  "result": 1,
  "text": "Host not deleted (not found)."
}
```

or

```
{
  "result": 0,
  "text": "Host deleted."
}
```

or

```
{
  "result": 1,
  "text": "Unable to delete a host because there is no hosts-database configured."
}
```

The command accepts the *operation-target* argument. By default, it removes the reservation from the hosts database only.

```
{
  "command": "reservation-del",
  "arguments": {
    "subnet-id": 4,
    "identifier-type": "hw-address",
    "identifier": "01:02:03:04:05:06",
    "operation-target": "memory"
  }
}
```

17.14.11 The reservation-update Command

reservation-update allows for the update of an existing host. It takes the same set of arguments as *reservation-add*, and also requires a host identifier and a subnet ID to identify the host that is being updated. The command can be as simple as having only the two mandatory entries:

```
{
  "command": "reservation-update",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "hw-address": "1a:1b:1c:1d:1e:1f"
    }
  }
}
```

In that case, however, it does not assign any resources to the host. An IPv4 address can be assigned as in this example:

```
{
  "command": "reservation-update",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "hw-address": "1a:1b:1c:1d:1e:1f",
      "ip-address": "192.0.2.202"
    }
  }
}
```

It can also take many more parameters, such as:

```
{
  "command": "reservation-update",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "client-id": "01:0a:0b:0c:0d:0e:0f",
      "ip-address": "192.0.2.205",
      "next-server": "192.0.2.1",
      "server-hostname": "hal9000",
      "boot-file-name": "/dev/null",
      "option-data": [
        {
          "name": "domain-name-servers",
          "data": "10.1.1.202,10.1.1.203"
        }
      ],
      "client-classes": [
        "office",
        "special_snowflake"
      ]
    }
  }
}
```

Here is an example of a complex IPv6 reservation update:

```
{
  "command": "reservation-update",
  "arguments": {
    "reservation": {
      "subnet-id": 1,
      "duid": "01:02:03:04:05:06:07:08:09:0A",
      "ip-addresses": [
        "2001:db8:1:cafe::1"
      ],
      "prefixes": [
        "2001:db8:2:abcd::/64"
      ],
      "hostname": "foo.example.com",
      "option-data": [
        {
          "name": "vendor-opts",
          "data": "4491"
        },
        {
          "name": "tftp-servers",
          "space": "vendor-4491",
          "data": "3000:1::234"
        }
      ]
    }
  }
}
```

The command returns a status that indicates either success (result 0) or failure (result 1), and a text parameter that confirms success or explains the cause of the failure. Here's an example of a successful update:

```
{
  "result": 0,
  "text": "Host updated."
}
```

And here's an example of a failure:

```
{
  "result": 1,
  "text": "Mandatory 'subnet-id' parameter missing."
}
```

The command accepts the `operation-target` argument. By default, it adds the reservation to the hosts database only. As `reservation-update` is expected to store the host, the `hosts-databases` parameter must be specified in the configuration, and databases must not run in read-only mode if the operation target is not the JSON configuration.

As with other update and set commands, this command overwrites all the contents of the entry. If the host previously had a resource assigned to it, and the `reservation-update` command is missing the resource, it is deleted from the database. If an incremental update of the host is desired, this can be achieved by issuing a `reservation-get-by-id` to get the current state of the host, selecting the host from the response, modifying it to the required outcome, and then issuing the `reservation-update` command with the resulting host attached.

17.14.12 General Mentions

Note

The host backends for the host cache and RADIUS hook libraries do not respond to commands that return a collection of host reservations, such as [reservation-get-all](#). Commands returning one host entry or dedicated host cache commands should be used instead.

17.15 libdhcp_lease_cmds.so: Lease Commands for Easier Lease Management

Kea allows users to store lease information in several backends (memfile, MySQL, and PostgreSQL), and the Lease Commands library provides an interface that can manipulate leases in a unified, safe way. In particular, it allows things that were previously impossible: lease manipulation in memfile while Kea is running, sanity check changes, lease existence checks, and removal of all leases belonging to a specific subnet. The hook library can also catch more obscure errors, like an attempt to add a lease with a `subnet-id` that does not exist in the configuration, or configuring a lease to use an address that is outside of the subnet to which it is supposed to belong. The library also provides a non-programmatic way to manage user contexts associated with leases.

As of Kea 2.7.7, the library also provides support for `binding-variables`. Binding variables allow the user to store custom values for each lease. They are discussed here: [Binding Variables](#).

Note

`libdhcp_lease_cmds.so` is part of the open source code and is available to every Kea user.

Note

This library can only be loaded by the `kea-dhcp4` or the `kea-dhcp6` process.

There are many situations where an administrative command may be useful; for example, during migration between servers or different vendors, when a certain network is being retired, or when a device has been disconnected and the system administrator knows that it will not be coming back. The `get` queries may be useful for automating certain management and monitoring tasks, and they can also act as preparatory steps for lease updates and removals.

This library provides the following commands:

- `lease4-add` - adds a new IPv4 lease.
- `lease6-add` - adds a new IPv6 lease.
- `lease6-bulk-apply` - creates, updates, and/or deletes multiple IPv6 leases in a single transaction.
- `lease4-get` - checks whether an IPv4 lease with the specified parameters exists and returns it if it does.
- `lease6-get` - checks whether an IPv6 lease with the specified parameters exists and returns it if it does.
- `lease4-get-all` - returns all IPv4 leases or all IPv4 leases for the specified subnets.
- `lease6-get-all` - returns all IPv6 leases or all IPv6 leases for the specified subnets.
- `lease4-get-page` - returns a set ("page") of leases from the list of all IPv4 leases in the database. By iterating through the pages it is possible to retrieve all the leases.

- `lease6-get-page` - returns a set ("page") of leases from the list of all IPv6 leases in the database. By iterating through the pages it is possible to retrieve all the leases.
- `lease4-get-by-hw-address` - returns all IPv4 leases with the specified hardware address.
- `lease4-get-by-client-id` - returns all IPv4 leases with the specified `client-id`.
- `lease6-get-by-duid` - returns all IPv6 leases with the specified DUID.
- `lease4-get-by-hostname` - returns all IPv4 leases with the specified hostname.
- `lease6-get-by-hostname` - returns all IPv6 leases with the specified hostname.
- `lease4-del` - deletes an IPv4 lease with the specified parameters.
- `lease6-del` - deletes an IPv6 lease with the specified parameters.
- `lease4-update` - updates (replaces) an existing IPv4 lease.
- `lease6-update` - updates (replaces) an existing IPv6 lease.
- `lease4-wipe` - removes all leases from a specific IPv4 subnet or from all subnets.
- `lease6-wipe` - removes all leases from a specific IPv6 subnet or from all subnets.
- `lease4-resend-ddns` - resends a request to update DNS entries for an existing lease.
- `lease6-resend-ddns` - resends a request to update DNS entries for an existing lease.
- `lease4-write` - writes the IPv4 memfile lease database into a file.
- `lease6-write` - writes the IPv6 memfile lease database into a file.

All commands use JSON syntax and can be issued either using the control channel (see [Management API](#)) or Control Agent (see [The Kea Control Agent](#)).

The library can be loaded in the same way as other hook libraries, and it does not take any parameters. It supports both the DHCPv4 and DHCPv6 servers.

```
"Dhcp6": {
  "hooks-libraries": [
    {
      "library": "/path/libdhcp_lease_cmds.so"
    },
    ...
  ]
}
```

17.15.1 The `lease4-add`, `lease6-add` Commands

The `lease4-add` and `lease6-add` commands allow a new lease to be created. Typically Kea creates a lease when it first sees a new device; however, sometimes it may be convenient to create the lease manually. The `lease4-add` command requires at least two parameters: an IPv4 address and an identifier, i.e. hardware (MAC) address. A third parameter, `subnet-id`, is optional. If the `subnet-id` is not specified or the specified value is 0, Kea tries to determine the value by running a subnet-selection procedure. If specified, however, its value must match the existing subnet. The simplest successful call might look as follows:

```
{
  "command": "lease4-add",
  "arguments": {
    "ip-address": "192.0.2.202",
    "hw-address": "1a:1b:1c:1d:1e:1f"
```

(continues on next page)

(continued from previous page)

```
}
}
```

The *lease6-add* command requires three parameters: an IPv6 address, an IAID value (identity association identifier, a value sent by clients), and a DUID. As with *lease4-add*, the *subnet-id* parameter is optional. If the *subnet-id* is not specified or the provided value is 0, Kea tries to determine the value by running a subnet-selection procedure. If specified, however, its value must match the existing subnet. For example:

```
{
  "command": "lease6-add",
  "arguments": {
    "subnet-id": 66,
    "ip-address": "2001:db8::3",
    "duid": "1a:1b:1c:1d:1e:1f:20:21:22:23:24",
    "iaid": 1234
  }
}
```

The *lease6-add* command can also be used to add leases for IPv6 prefixes. In this case there are three additional parameters that must be specified: *subnet-id*, *type* (set to "IA_PD"), and prefix length. The actual prefix is set using the *ip-address* field. Note that Kea cannot guess *subnet-id* values for prefixes; they must be specified explicitly. For example, to configure a lease for prefix 2001:db8:abcd::/48, the following command can be used:

```
{
  "command": "lease6-add",
  "arguments": {
    "subnet-id": 66,
    "type": "IA_PD",
    "ip-address": "2001:db8:abcd:",
    "prefix-len": 48,
    "duid": "1a:1b:1c:1d:1e:1f:20:21:22:23:24",
    "iaid": 1234
  }
}
```

The commands can take several additional optional parameters:

- **valid-lft** - specifies the lifetime of the lease, expressed in seconds. If not specified, the value configured in the subnet related to the specified *subnet-id* is used.
- **expire** - creates a timestamp of the lease expiration time, expressed in UNIX format (seconds since 1 Jan 1970). If not specified, the default value is the current time plus the lease lifetime (the value of *valid-lft*).
- **fqdn-fwd** - specifies whether the lease should be marked as if a forward DNS update were conducted. This only affects the data stored in the lease database, and no DNS update will be performed. If configured, a DNS update to remove the A or AAAA records will be conducted when the lease is removed due to expiration or being released by a client. If not specified, the default value is *false*. The *hostname* parameter must be specified if *fqdn-fwd* is set to *true*.
- **fqdn-rev** - specifies whether the lease should be marked as if reverse DNS update were conducted. This only affects the data stored in the lease database, and no DNS update will be performed.. If configured, a DNS update to remove the PTR record will be conducted when the lease is removed due to expiration or being released by a client. If not specified, the default value is *false*. The *hostname* parameter must be specified if *fqdn-fwd* is set to *true*.

- **hostname** - specifies the hostname to be associated with this lease. Its value must be non-empty if either **fqdn-fwd** or **fqdn-rev** are set to **true**. If not specified, the default value is an empty string.
- **hw-address** - optionally specifies a hardware (MAC) address for an IPv6 lease. It is a mandatory parameter for an IPv4 lease.
- **client-id** - optionally specifies a client identifier for an IPv4 lease.
- **preferred-lft** - optionally specifies a preferred lifetime for IPv6 leases. If not specified, the value configured for the subnet corresponding to the specified **subnet-id** is used. This parameter is not used when adding an IPv4 lease.
- **state** - specifies the state of an added lease, which can be 0 for **default**, 1 for **declined**, 2 for **expired-reclaimed**, and 3 for the **released** state. Any other value causes an error. Using 1 for a "IA_PD" lease type is illegal and will be rejected.
- **user-context** - specifies the user context to be associated with this lease. It must be a JSON map.

Here is an example of a fairly complex lease addition:

```
{
  "command": "lease6-add",
  "arguments": {
    "subnet-id": 66,
    "ip-address": "2001:db8::3",
    "duid": "01:02:03:04:05:06:07:08",
    "iaid": 1234,
    "hw-address": "1a:1b:1c:1d:1e:1f",
    "preferred-lft": 500,
    "valid-lft": 1000,
    "expire": 12345678,
    "fqdn-fwd": true,
    "fqdn-rev": true,
    "state": 0,
    "hostname": "urania.example.org",
    "user-context": { "version": 1 }
  }
}
```

The command returns a status that indicates either success (result 0) or failure (result 1). A failed command always includes a text parameter that explains the cause of failure. For example:

```
{ "result": 0, "text": "Lease added." }
```

Example failure:

```
{ "result": 1, "text": "missing parameter 'ip-address' (<string>:3:19)" }
```

17.15.2 The lease6-bulk-apply Command

The *lease6-bulk-apply* was implemented to address the performance penalty in High-Availability mode when a single DHCPv6 transaction resulted in multiple lease updates sent to the partner, if multiple address and/or prefix leases were allocated. Consider the case when a DHCPv6 client requests the assignment of two IPv6 addresses and two IPv6 prefixes: it may result in the allocation of four leases. In addition, DHCPv6 may assign a different address than the one requested by the client during the renew or rebind stage, and delete the leases previously used by this client. There are six lease changes sent between the HA partners in this case. Sending these updates as individual commands,

e.g. via `lease6-update`, is highly inefficient and produces unnecessary delays in communication, both between the HA partners and in sending the response to the DHCPv6 client.

The `lease6-bulk-apply` command deals with this problem by aggregating all lease changes - both deleted leases and new or updated leases - in a single command. The receiving server iterates over the deleted leases and deletes them from its lease database. Next, it iterates over the new/updated leases and adds them to the database or updates them if they already exist.

Even though High Availability is the major application for this command, it can be freely used in all cases when it is desirable to send multiple lease changes in a single command.

In the following example, we delete two leases and add or update two other leases in the database:

```
{
  "command": "lease6-bulk-apply",
  "arguments": {
    "deleted-leases": [
      {
        "ip-address": "2001:db8:abcd:",
        "type": "IA_PD",
        ...
      },
      {
        "ip-address": "2001:db8:abcd:234",
        "type": "IA_NA",
        ...
      }
    ],
    "leases": [
      {
        "subnet-id": 66,
        "ip-address": "2001:db8:cafe:",
        "type": "IA_PD",
        ...
      },
      {
        "subnet-id": 66,
        "ip-address": "2001:db8:abcd:333",
        "type": "IA_NA",
        ...
      }
    ]
  }
}
```

If any of the leases are malformed, no lease changes are applied to the lease database. If the leases are well-formed but there is a failure to apply any of the lease changes to the database, the command continues to be processed for other leases. All the leases for which the command was unable to apply the changes in the database are listed in the response. For example:

```
{
  "result": 0,
  "text": "Bulk apply of 2 IPv6 leases completed",
  "arguments": {
    "failed-deleted-leases": [
```

(continues on next page)

(continued from previous page)

```

    {
      "ip-address": "2001:db8:abcd:",
      "type": "IA_PD",
      "result": 3,
      "error-message": "no lease found"
    }
  ],
  "failed-leases": [
    {
      "ip-address": "2001:db8:cafe:",
      "type": "IA_PD",
      "result": 1,
      "error-message": "unable to communicate with the lease database"
    }
  ]
}

```

The response above indicates that the hook library was unable to delete the lease for prefix "2001:db8:abcd:" and add or update the lease for prefix "2001:db8:cafe:". However, there are two other lease changes which have been applied as indicated by the text message. The `result` is the status constant that indicates the type of the error experienced for the particular lease. The meanings of the returned codes are the same as the results returned for the commands. In particular, the result of 1 indicates an error while processing the lease, e.g. a communication error with the database. The result of 3 indicates that an attempt to delete the lease was unsuccessful because such a lease doesn't exist (an empty result).

17.15.3 The `lease4-get`, `lease6-get` Commands

`lease4-get` and `lease6-get` can be used to query the lease database and retrieve existing leases. There are two types of parameters the `lease4-get` command supports: (address) or (subnet-id, identifier-type, identifier). There are also two types for `lease6-get`: (address, type) or (subnet-id, identifier-type, identifier, IAID, type). The first type of query is used when the address (either IPv4 or IPv6) is known, but the details of the lease are not; one common use case of this type of query is to find out whether a given address is being used. The second query uses identifiers; currently supported identifiers for leases are: "hw-address" (IPv4 only), "client-id" (IPv4 only), and "duid" (IPv6 only).

An example `lease4-get` command for getting a lease using an IPv4 address is:

```

{
  "command": "lease4-get",
  "arguments": {
    "ip-address": "192.0.2.1"
  }
}

```

An example of the `lease6-get` query is:

```

{
  "command": "lease6-get",
  "arguments": {
    "ip-address": "2001:db8:1234:ab:",
    "type": "IA_PD"
  }
}

```

(continues on next page)

(continued from previous page)

}

An example query by "hw-address" for an IPv4 lease looks as follows:

```
{
  "command": "lease4-get",
  "arguments": {
    "identifier-type": "hw-address",
    "identifier": "08:08:08:08:08:08",
    "subnet-id": 44
  }
}
```

An example query by "client-id" for an IPv4 lease looks as follows:

```
{
  "command": "lease4-get",
  "arguments": {
    "identifier-type": "client-id",
    "identifier": "01:01:02:03:04:05:06",
    "subnet-id": 44
  }
}
```

An example query by (subnet-id, identifier-type, identifier, iaaid, type) for an IPv6 lease is:

```
{
  "command": "lease6-get",
  "arguments": {
    "identifier-type": "duid",
    "identifier": "08:08:08:08:08:08",
    "iaaid": 1234567,
    "type": "IA_NA",
    "subnet-id": 44
  }
}
```

The type is an optional parameter. Supported values are: IA_NA (non-temporary address) and IA_PD (IPv6 prefix). If not specified, IA_NA is assumed.

lease4-get and *lease6-get* return an indication of the result of the operation and lease details, if found. The result has one of the following values: 0 (success), 1 (error), or 3 (empty). An empty result means that a query has been completed properly, but the object (a lease in this case) has not been found. The lease parameters, if found, are returned as arguments. *client-id* is not returned if empty.

An example result returned when the host was found:

```
{
  "arguments": {
    "client-id": "42:42:42:42:42:42:42:42",
    "cltt": 12345678,
    "fqdn-fwd": false,
    "fqdn-rev": true,
    "hostname": "myhost.example.com.",

```

(continues on next page)

(continued from previous page)

```

    "hw-address": "08:08:08:08:08:08",
    "ip-address": "192.0.2.1",
    "state": 0,
    "subnet-id": 44,
    "valid-lft": 3600
  },
  "result": 0,
  "text": "IPv4 lease found."
}
```

Note

The client last transaction time (*cltt* field) is bound to the valid lifetime (*valid-lft*) and to the expire date (not reported here but stored in databases) by the equation $cltt + valid\_lft = expire$

at the exception of the infinite valid lifetime coded by the 0xffffffff (4294967295) special value which makes the expire value to overflow on MySQL and old PostgreSQL backends where timestamps are 32 bit long. So in these lease databases the expire date is the same as the *cltt* i.e. $cltt = expire$ when $valid\_lft = 4294967295$ and the lease backend is MySQL or PostgreSQL.

17.15.4 The lease4-get-all, lease6-get-all Commands

lease4-get-all and *lease6-get-all* are used to retrieve all IPv4 or IPv6 leases, or all leases for the specified set of subnets. All leases are returned when there are no arguments specified with the command, as in the following example:

```

{
  "command": "lease4-get-all"
}
```

If arguments are provided, it is expected that they contain the "subnets" parameter, which is a list of subnet identifiers for which leases should be returned. For example, to retrieve all IPv6 leases belonging to the subnets with identifiers 1, 2, 3, and 4:

```

{
  "command": "lease6-get-all",
  "arguments": {
    "subnets": [ 1, 2, 3, 4 ]
  }
}
```

The returned response contains a detailed list of leases in the following format:

```

{
  "arguments": {
    "leases": [
      {
        "cltt": 12345678,
        "duid": "42:42:42:42:42:42:42:42",
        "fqdn-fwd": false,
        "fqdn-rev": true,
        "hostname": "myhost.example.com.",

```

(continues on next page)

(continued from previous page)

```

        "hw-address": "08:08:08:08:08:08",
        "iaid": 1,
        "ip-address": "2001:db8:2::1",
        "preferred-lft": 500,
        "state": 0,
        "subnet-id": 44,
        "type": "IA_NA",
        "valid-lft": 3600
    },
    {
        "cltt": 12345678,
        "duid": "21:21:21:21:21:21:21:21",
        "fqdn-fwd": false,
        "fqdn-rev": true,
        "hostname": "",
        "iaid": 1,
        "ip-address": "2001:db8:0:0:2::",
        "preferred-lft": 500,
        "prefix-len": 80,
        "state": 0,
        "subnet-id": 44,
        "type": "IA_PD",
        "valid-lft": 3600
    }
]
},
"result": 0,
"text": "2 IPv6 lease(s) found."
}

```

Warning

The *lease4-get-all* and *lease6-get-all* commands may result in very large responses. This may have a negative impact on the DHCP server's responsiveness while the response is generated and transmitted over the control channel, as the server imposes no restriction on the number of leases returned as a result of this command.

17.15.5 The *lease4-get-page*, *lease6-get-page* Commands

The *lease4-get-all* and *lease6-get-all* commands may result in very large responses; generating such a response may consume CPU bandwidth as well as memory. It may even cause the server to become unresponsive. In the case of large lease databases it is usually better to retrieve leases in chunks, using the paging mechanism. *lease4-get-page* and *lease6-get-page* implement a paging mechanism for DHCPv4 and DHCPv6 servers, respectively. The following command retrieves the first 1024 IPv4 leases:

```

{
    "command": "lease4-get-page",
    "arguments": {
        "from": "start",
        "limit": 1024
    }
}

```

The keyword `start` denotes that the first page of leases should be retrieved. Alternatively, an IPv4 zero address can be specified to retrieve the first page:

```
{
  "command": "lease4-get-page",
  "arguments": {
    "from": "0.0.0.0",
    "limit": 1024
  }
}
```

Similarly, the IPv6 zero address can be specified in the `lease6-get-page` command:

```
{
  "command": "lease6-get-page",
  "arguments": {
    "from": "::",
    "limit": 6
  }
}
```

The response has the following structure:

```
{
  "arguments": {
    "leases": [
      {
        "ip-address": "2001:db8:2::1",
        ...
      },
      {
        "ip-address": "2001:db8:2::9",
        ...
      },
      {
        "ip-address": "2001:db8:3::1",
        ...
      },
      {
        "ip-address": "2001:db8:5::3",
        ...
      },
      {
        "ip-address": "2001:db8:4::1",
        ...
      },
      {
        "ip-address": "2001:db8:2::7",
        ...
      },
      ...
    ],
    "count": 6
  },
}
```

(continues on next page)

(continued from previous page)

```

    "result": 0,
    "text": "6 IPv6 lease(s) found."
}

```

Note that the leases' details were excluded from the response above for brevity.

Generally, the returned list is not sorted in any particular order. Some lease database backends may sort leases in ascending order of addresses, but the controlling client must not rely on this behavior.

The `count` parameter contains the number of returned leases on the page.

To fetch the next page, the client must use the last address of the current page as an input to the next *lease4-get-page* or *lease6-get-page* command call. In this example it is:

```

{
  "command": "lease6-get-page",
  "arguments": {
    "from": "2001:db8:2::7",
    "count": 6
  }
}

```

because 2001:db8:2::7 is the last address on the current page.

The client may assume that it has reached the last page when the `count` value is lower than that specified in the command; this includes the case when the `count` is equal to 0, meaning that no leases were found.

17.15.6 The *lease4-get-by-**, *lease6-get-by-** Commands

*lease4-get-by-** and *lease6-get-by-** can be used to query the lease database and retrieve all existing leases matching a given feature (denoted by the `*`). These can include a specified hardware address (IPv4 only), `client-id` (IPv4 only), `duid` (IPv6 only) identifiers, or hostname.

An example *lease4-get-by-hw-address* command for getting IPv4 leases with a given hardware address is:

```

{
  "command": "lease4-get-by-hw-address",
  "arguments": {
    "hw-address": "08:08:08:08:08:08"
  }
}

```

An example of the *lease6-get-by-hostname* is:

```

{
  "command": "lease6-get-by-hostname",
  "arguments": {
    "hostname": "myhost.example.org"
  }
}

```

The `by` key is the only parameter. The returned response contains a detailed list of leases in the same format as *lease4-get-all* or *lease6-get-all*. This list can be empty and is usually not large.

17.15.7 The `lease4-del`, `lease6-del` Commands

`lease4-del` and `lease6-del` can be used to delete a lease from the lease database. There are two types of parameters these commands support, similar to the `lease4-get` and `lease6-get` commands: (address) for both v4 and v6, (subnet-id, identifier-type, identifier) for v4, and (subnet-id, identifier-type, identifier, type, IAID) for v6. The first type of query is used when the address (either IPv4 or IPv6) is known, but the details of the lease are not. One common use case is where an administrator wants a specified address to no longer be used. The second form of the command uses identifiers. For maximum flexibility, this interface uses identifiers as a pair of values: the type and the actual identifier. The currently supported identifiers are "hw-address" (IPv4 only), "client-id" (IPv4 only), and "duid" (IPv6 only).

An example command for deleting an IPv4 lease by address is:

```
{
  "command": "lease4-del",
  "arguments": {
    "ip-address": "192.0.2.202"
  }
}
```

An example IPv4 lease deletion by "hw-address" is:

```
{
  "command": "lease4-del",
  "arguments": {
    "identifier": "08:08:08:08:08:08",
    "identifier-type": "hw-address",
    "subnet-id": 44
  }
}
```

The IPv6 address leases are deleted the same way, but using `lease6-del`. The IPv6 prefix leases are also deleted using `lease6-del`, but with some extra steps. The prefix should be referenced by its address and prefix length should be ignored. For example, to delete 2001:db8:1::/48, the following command can be used. The subnet-id parameter is optional.

```
{
  "command": "lease6-del",
  "arguments": {
    "ip-address": "2001:db8:1::",
    "type": "IA_PD",
    "subnet-id": 1
  }
}
```

Another parameter called `update-ddns`, when `true`, instructs the server to queue a request to `kea-dhcp-ddns` to remove DNS entries after the lease is successfully deleted if:

- DDNS updating is enabled (i.e. "dhcp-ddns":{ "enable-updates": true }).
- The lease's hostname is not empty.
- At least one of the lease's DNS direction flags (`fqdn_fwd` or `fqdn_rev`) is true.

This parameter defaults to `false`. An example of its use is shown below:

```
{
  "command": "lease4-del",
  "arguments": {
    "ip-address": "192.0.2.202",
    "update-ddns": true
  }
}
```

Commands *lease4-del* and *lease6-del* return a result that indicates the outcome of the operation. It has one of the following values: 0 (success), 1 (error), or 3 (empty). The empty result means that a query has been completed properly, but the object (a lease, in this case) has not been found.

17.15.8 The *lease4-update*, *lease6-update* Commands

The *lease4-update* and *lease6-update* commands can be used to update existing leases. Since all lease database backends are indexed by IP addresses, it is not possible to update an address, but all other fields may be altered. If an address needs to be changed, please use *lease4-del* / *lease6-del* followed by *lease4-add* / *lease6-add*.

The *subnet-id* parameter is optional. If not specified, or if the specified value is 0, Kea tries to determine its value by running a subnet-selection procedure. If specified, however, its value must match the existing subnet.

The optional boolean parameter "force-create" specifies whether the lease should be created if it does not exist in the database. It defaults to *false*, which indicates that the lease is not created if it does not exist. In such a case, an error is returned when trying to update a non-existing lease. If the "force-create" parameter is set to *true* and the updated lease does not exist, the new lease is created as a result of receiving the *lease4-update* / *lease6-update* command.

An example of a command to update an IPv4 lease is:

```
{
  "command": "lease4-update",
  "arguments": {
    "ip-address": "192.0.2.1",
    "hostname": "newhostname.example.org",
    "hw-address": "1a:1b:1c:1d:1e:1f",
    "subnet-id": 44,
    "force-create": true
  }
}
```

An example of a command to update an IPv6 lease is:

```
{
  "command": "lease6-update",
  "arguments": {
    "ip-address": "2001:db8::1",
    "duid": "88:88:88:88:88:88:88:88",
    "iaid": 7654321,
    "hostname": "newhostname.example.org",
    "subnet-id": 66,
    "force-create": false
  }
}
```

As with other update commands, this command overwrites all the contents of the entry. If the lease previously had a resource assigned to it, and the *lease4-update* / *lease6-update* command is missing the resource, it is deleted from

the lease database. If an incremental update of the lease is desired, then this can be achieved by issuing a *lease4-get* / *lease6-get* command to get the current state of the lease, selecting the lease from the response, modifying it to the required outcome, and then issuing the *lease4-update* / *lease6-update* command with the resulting lease attached.

17.15.9 The *lease4-wipe*, *lease6-wipe* Commands

lease4-wipe and *lease6-wipe* are designed to remove all leases associated with a given subnet. This administrative task is expected to be used when an existing subnet is being retired. The leases are not properly expired; no DNS updates are carried out, no log messages are created, and hooks are not called for the leases being removed.

An example of *lease4-wipe* is:

```
{
  "command": "lease4-wipe",
  "arguments": {
    "subnet-id": 44
  }
}
```

An example of *lease6-wipe* is:

```
{
  "command": "lease6-wipe",
  "arguments": {
    "subnet-id": 66
  }
}
```

The commands return a text description of the number of leases removed, plus the status code 0 (success) if any leases were removed or 3 (empty) if there were no leases. Status code 1 (error) may be returned if the parameters are incorrect or some other exception is encountered.

subnet-id 0 has a special meaning; it tells Kea to delete leases from all configured subnets. Also, the *subnet-id* parameter may be omitted. If not specified, leases from all subnets are wiped.

Note: currently only memfile lease storage supports this command.

17.15.10 The *lease4-resend-ddns*, *lease6-resend-ddns* Commands

lease4-resend-ddns and *lease6-resend-ddns* can be used to generate a request to *kea-dhcp-ddns* to update the DNS entries for an existing lease. The desired lease is selected by a single parameter, "ip-address". For an update request to be generated, DDNS updating must be enabled and DNS entries must have already been made (or attempted) for the lease. In other words, all of the following must be true:

- DDNS updating must be enabled (i.e. "dhcp-ddns":{ "enable-updates": true}).
- The lease's hostname must not be empty.
- At least one of the lease's DNS direction flags (*fqdn_fwd* or *fqdn_rev*) must be true.

An example *lease4-resend-ddns* command for getting a lease using an IPv4 address is:

```
{
  "command": "lease4-resend-ddns",
  "arguments": {
    "ip-address": "192.0.2.1"
  }
}
```

An example of the *lease6-resend-ddns* query is:

```
{
  "command": "lease6-resend-ddns",
  "arguments": {
    "ip-address": "2001:db8:1::1"
  }
}
```

Commands *lease4-resend-ddns* and *lease6-resend-ddns* return an indication of the result of the operation. It has one of the following values: 0 (success), 1 (error), or 3 (empty). An empty result means that a query has been completed properly, but the object (a lease in this case) has not been found.

A successful result does not mean that DNS has been successfully updated; it indicates that a request to update DNS has been successfully created and queued for transmission to *kea-dhcp-ddns*.

Here's an example of a result returned when the lease was found:

```
{
  "result": 0,
  "text": "NCR generated for: 2001:db8:1::1, hostname: example.com."
}
```

17.15.11 The *lease4-write*, *lease6-write* Commands

lease4-write and *lease6-write* can be used for recovery in emergency situations where the memfile lease file is damaged, e.g. removed by accident or truncated by a full file system, but the in-memory database is still valid. These commands are supported only by the memfile database backend and write the lease database into a CSV file. They take the path of the file as the *filename* argument. If the specified output file is the same as the configured memfile one, the backend closes and reopens the file in an attempt to synchronize both the files and the in-memory images of the lease database. The extension *.bak* and the server PID number are added to the previous filename: for example, *.bak14326*.

Note

As of Kea 2.7.9, the lease file may only be written to the data directory determined during compilation: "[*kea-install-dir*]/var/lib/kea". This path may be overridden at startup by setting the environment variable *KEA_DHCP_DATA_DIR* to the desired path. If a path other than this value is used in *filename*, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use in specifying a custom file name simply omit the path portion from *filename*.

Note

These commands do not replace the LFC mechanism; they should be used only in exceptional circumstances, such as when recovering after running out of disk space.

17.15.12 Binding Variables

Binding variables allow users to store custom values with each lease. The values are calculated using expressions and stored in the lease's *user-context*. The feature is similar to ISC DHCP's *set* statement.

They are configured as parameters for the *lease-cmds* hook library. Each variable has the following parameters:

- **name**
The name of the variable that will appear in the `user-context`. The name must be unique within the list of binding variables..
- **expression**
The expression used to calculate the variable's value (see :ref: *classification-using-expressions*).
- **source**
The packet to use as input for the expression. It is either `query` (the packet sent by the client) or `response` (the packet the server is sending in response).

The following example would store a user-defined option sent by the client along with the domain-name sent by the server for each DHCPv4 lease:

```
{
  "hooks-libraries": [{
    "library": "/path/libdhcp_lease_cmds.so",
    "parameters": {
      "binding-variables": [{
        "name": "opt-222",
        "expression": "hexstring(option[222].hex, ':')",
        "source": "query"
      }, {
        "name": "domain-name",
        "expression": "option[15].text",
        "source": "response"
      }]
    }
  }]
}
```

The values are stored as name-value pairs in the ISC map in the lease's `user-context` contents would look similar to the following:

```
{
  "ISC": {
    "binding-variables": [{
      "domain-name": "example.org",
      "opt-222": "01:02:03:04"
    }]
  }
}
```

When stored in database back ends, the `user-context` will not contain line breaks as shown above and would appear as follows:

```
{ "ISC": { "binding-variables": { "domain-name": "example.org", "opt-222": "01:02:03:04"
↪ } } }
```

When using `memfile` lease storage, the `user-context` is output with commas replaced by the escape sequence `"&#x2c"`. This is necessary to prevent interference with the lease file parsing. The example output above would appear as follows in a `memfile` lease file:

```
{ "ISC": { "binding-variables": { "domain-name": "example.org"&#x2c "opt-222":
↪ "01:02:03:04" } } }
```

Binding variable values are evaluated whenever a lease is assigned or renewed. For *kea-dhcp4*, they are also added if *offer-lifetime* is greater than zero. The lease is only updated in the back end if the values have changed.

Note

When used in conjunction with the Ping Check hook library, the *lease-cmds* hook library should be listed after Ping Check in the server's *hooks-libraries* section of the configuration.

When used in conjunction with the HA hook library, the *lease-cmds* hook library must be listed before HA in the server's *hooks-libraries* section of the configuration. This ensures that the binding variables are evaluated before HA sends lease updates to its peer(s).

17.16 *libdhcp_lease_query.so*: Leasequery Support

This library provides support for DHCPv4 Leasequery as described in [RFC 4388](#); and for DHCPv6 Leasequery as described in ([RFC 5007](#)).

Note

libdhcp_lease_query.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

Kea version 2.3.4 added support for DHCPv6 Bulk Leasequery ([RFC 5460](#)); Kea version 2.3.5 added support for DHCPv4 Bulk Leasequery ([RFC 6926](#)) using the memfile lease backend.

17.16.1 DHCPv4 Leasequery

DHCPv4 simple Leasequery provides a requester the ability to query for active lease information for either a single IP address or a single client. RFC 4388 calls for three such queries:

- Query by IP address

The IP address of interest is contained within the *ciaddr* field of the query.

- Query by hardware address

The hardware address of interest is contained with the *chaddr* field of the query.

- Query by client identifier

The client identifier of interest is sent in the *dhcp-client-identifier* option (61) of the query.

The inbound DHCPLEASEQUERY packet must supply only one of the three values above. Queries which supply more than one of these values are dropped.

In addition, the query must contain the IP address of the requester in *giaddr*. This value is used not only as the destination for the query response but also to validate the requester against a known list of IP addresses which are permitted to query. This list of valid requester addresses is specified as part of the Leasequery hook library's configuration (see the section on configuration below).

In response to a valid query, the server returns one of three message types:

- DHCPLEASEUNKNOWN

Returned when the IP address of interest is not one the server knows about (query by IP address); or there are no active leases for the client of interest (query by hardware address or client ID).

- DHCPLEASEUNASSIGNED

Returned when the IP address is one the server knows of but for which there are no active leases (applies only to query by IP address).

- DHCPLEASEACTIVE

Returned when there is at least one active lease found matching the criteria.

For both DHCPLEASEUNKNOWN and DHCPLEASEUNASSIGNED responses, the only information sent back to the requester in response is the query parameter itself (i.e. one of: IP address, hardware address, or client identifier).

For DHCPLEASEACTIVE the server provides the following information for the newest active lease that matches the criteria, in the response:

- ciaddr - set to the lease's IP address
- chaddr - set to the lease's hardware address

In addition, one or more of the following options are included:

Table 3: DHCPLEASEACTIVE options

Option	Code	Content
dhcp-client-identifier	61	copied from the lease (if appropriate)
client-last-transaction-time	91	the amount of time that has elapsed since the lease's client-last-transaction-time (CLTT). This value is also used by the server to adjust lifetime and timer values.
dhcp-lease-time	51	lease's lifetime reduced by CLTT
dhcp-renewal-time	58	as controlled by kea-dhcp4 configuration and then reduced by CLTT
dhcp-rebind-time	59	as dictated by kea-dhcp4 configuration and then reduced by CLTT
dhcp-agent-options	82	if stored on the lease. (See <i>Storing Extended Lease Information</i>)
associated-ip	92	a list of all other IP addresses for which the client has active leases. (Does not apply to query by IP address)

The `dhcp-server-identifier` option (54) is returned in all responses in keeping with RFC 2131, section 4.3.1.

RFC 4388 allows requesters to ask for specific options via the `dhcp-parameter-request-list` (PRL, option 55). This is not currently supported in Kea.

17.16.2 DHCPv4 Leasequery Configuration

Configuring the Leasequery hook library for use is straightforward. It supports a single parameter, `requesters`, which is a list of IP addresses from which DHCPLEASEQUERY packets are accepted. In other words, it is a list of known requesters. The following code shows an example configuration with two requester addresses:

```
{
  "hooks-libraries": [
    {
      "library": "lib/kea/hooks/libdhcp_lease_query.so",
      "parameters": {
        "requesters": [ "192.0.1.1", "10.0.0.2" ]
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
],
...
}
```

Note

For security purposes, there is no way to specify wildcards. Each requester address must be explicitly listed.

17.16.3 DHCPv6 Leasequery

DHCPv6 simple Leasequery gives a requester the ability to query for active lease information for either a single IP address or a single client DUID. The query type and parameters are conveyed in an `lq-query` option (44) attached to a `DHCPV6_LEASEQUERY` message:

- `query-type`

This is either `query-by-address` (1) or `query-by-clientid` (2)

- `link-address`

The global link address, when not empty, instructs the query to be limited to leases within that "link." Kea uses this value to select only leases that belong to subnets whose prefix matches this value. Active leases for prefix delegations for a matched subnet are included in the query reply, even if the delegated prefix itself falls outside the subnet prefix.

- `query-options`

A single `iaaddr` option (12) must be supplied when querying by address. When querying by client ID, a single `clientid` option (1) must be supplied. RFC 5007 also calls for an optional, `oro` option (6), to request specific options be returned for matched leases. This is not currently implemented.

Note

RFC 5007, Section 3.3 states that querying by IP address should return either a lease (e.g. binding) for the address itself or a lease for a delegated prefix that contains the address.

`DHCPV6_LEASEQUERY` queries are only honored if the source address of the query matches an entry in a list of known IP addresses which are permitted to query. This list of valid requester addresses is specified as part of the Leasequery hook library's configuration (see the section on configuration below). Queries received from unknown requesters are logged and dropped.

In response to a valid query, the server carries out the requisite activities and returns a `DHCPV6_LEASEQUERY_REPLY`. All replies contain at least a `status-code` option (13) that indicates the outcome of the query as detailed in the following table:

Table 4: `DHCPV6_LEASEQUERY_REPLY` status option values per query outcome

Query Outcome	Status Label	Status Code	Status Text
Invalid query type field	<code>STATUS_UnknownQueryType</code>	7	"unknown query-ty
Query by IP address that does not contain an address option	<code>STATUS_Malformed</code>	10	"missing D6O_IAA

continue

Table 4 – continued from previous page

Query Outcome	Status Label	Status Code	Status Text
Query by IP address for an address that does fall within any configured pools	STATUS_NotConfigured	9	"address not in a co
Query by IP address which found only an inactive lease (e.g. expired, declined, reclaimed-expired)	STATUS_Success	0	"inactive lease exists"
Query by IP address that found no leases (active or otherwise)	STATUS_Success	0	"no active lease"
Query by IP address that found an active lease for the address	STATUS_Success	0	"active lease found"
Query by Client ID that does not contain a client ID option	STATUS_Malformed	10	"missing D6O_CL
Query by Client ID with a link address that does not match any configured subnets	STATUS_NotConfigured	9	"not a configured li
Query by client ID which found no matching leases	STATUS_Success	0	"no active leases"
Query by client ID which found one or more active leases	STATUS_Success	0	"active lease(s) fou

For those scenarios where the query was either invalid or for which no matching active leases were found, the DHCPV6_LEASEQUERY_REPLY only contains the status-code option (12) per the above table.

When a query finds active leases in more than one subnet and the query's link-address is empty, then, in addition to the status-code, the DHCPV6_LEASEQUERY_REPLY contains an lq-client-link option (48). The lq-client-link contains a list of IPv6 addresses, one for each subnet in which a lease was found (see RFC 5007, Section 4.1.2.5). If, however, the query's link-address is not empty, the list of queries is pruned to contain only leases that belong to that subnet.

When the query results in one or more active leases which all belong to a single subnet, in addition to the status-code, the DHCPV6_LEASEQUERY_REPLY contains a client-data option (45) (see RFC 5007, Section 4.1.2.2). The client-data option encapsulates the following options:

Table 5: OPTION_CLIENT_DATA returned when active lease(s) are found

Option	Code	Content
clientid	1	copied from the lease (if one exists)
clt-time	46	amount of time that has elapsed since the lease's client-last-transaction-time (CLTT). This value will also be used by the server to adjust lifetime and timer values.
iaaddr	5	One option per matched address. Fields in each option: - lease address - valid lifetime reduced by CLTT - preferred lifetime reduced by CLTT
iaprefix	26	One option per matched prefix. Fields in each option: - prefix - prefix length - valid lifetime reduced by CLTT - preferred lifetime reduced by CLTT

If the lease with the most recent client-last-transaction-time (CLTT) value has relay information in its user context (see *Storing Extended Lease Information*), then an OPTION_LQ_RELAY_DATA option is added to the reply (see RFC 5007, Section 4.1.2.4).

The relay information on the lease is a list with an entry for each relay layer the client packet (e.g. DHCPV6_REQUEST) traversed, with the first entry in the list being the outermost layer (closest to the server). The peer-address field of the lq-rely-option is set to the peer address of this relay. The list of relays is then used to construct a DHCPV6_RELAY_FORW message equivalent to that which contained the client packet, minus the client packet. This message is stored in the DHCP-relay-message field of the lq-relay-data option.

17.16.4 DHCPv6 Leasequery Configuration

Configuring the Leasequery hook library for use is straightforward. It supports a single parameter, `requesters`, which is a list of IP addresses from which DHCPV6_LEASEQUERY packets are accepted. In other words, it is a list of known requesters. The following code shows an example configuration with two requester addresses:

```
{
  "hooks-libraries": [
    {
      "library": "lib/kea/hooks/libdhcp_lease_query.so",
      "parameters": {
        "requesters": [ "2001:db8:1::1", "2001:db8:2::1" ],
        "prefix-lengths": [ 72 ]
      }
    }
  ],
  ...
}
```

Note

For security purposes, there is no way to specify wildcards. Each requester address must be explicitly listed.

When a query by IP address does not match an existing address lease, a search for a matching delegated prefix is conducted. This is carried out by iterating over a list of prefix lengths in descending order, extracting a prefix of that length from the query address, and searching for a delegation matching the resulting prefix. This continues for each length in the list until a match is found or the list is exhausted.

By default, the list of prefix lengths to use in the search is determined dynamically after (re)configuration events. The resulting list contains unique values of `delegated-len`, gleaned from the currently configured set of PD pools.

There is an optional parameter, `prefix-lengths`, which provides the ability to explicitly configure the list rather than having it be determined dynamically. This provides tighter control over which prefix lengths are searched. In the above example, the prefix-length search is restricted to a single pass, using a length of 72, regardless of whether there are pools using other values for `delegated-len`. Specifying an empty list, as shown below:

```
:
    "prefix-lengths": [ ]
:
```

disables the search for delegated prefixes for query by IP address.

17.16.5 DHCPv4 Bulk Leasequery

DHCPv4 Bulk Leasequery gives a requester the ability to query for active lease information over a TCP connection. This allows the server to return all leases matching a given query.

Two of the query types identified by RFC 4388 - Query by MAC address and Query by Client-identifier - are Bulk Leasequery types specified by RFC 6926. That RFC also defines these new Bulk Leasequery types:

- Query by Relay Identifier

The query carries an RAI (dhcp-agent-options (82)) option with a relay-id (12) sub-option.

- Query by Remote ID

The query carries an RAI (dhcp-agent-options (82) option) with a remote-id (2) sub-option.

- Query for All Configured IP Addresses

This query type is selected when no other query type is specified.

RFC 6926 also defines new options for Bulk Leasequery:

- status-code (151)

This reply option carries a status code such as `MalformedQuery` or `NotAllowed`, with an optional text message.

- base-time (152)

This reply option carries the absolute current time that the response was created. All other time-based reply options are related to this value.

- start-time-of-state (153)

This reply option carries the time of the lease's transition into its current state.

- query-start-time (154)

This query option specifies a start-query time; replies will only contain leases that are older than this value.

- query-end-time (155)

This query option specifies an end-query time; replies will only contain leases that are newer than this value.

- dhcp-state (156)

This reply option carries the lease state.

- data-source (157)

This reply option carries the source of the data as a remote flag.

RFC 6926 reuses and extends the Virtual Subnet Selection option (221) defined in RFC 6607.

Note

Kea does not yet support querying for all configured IP addresses, so the `dhcp-state` option cannot be used, as only active leases can be returned in replies. Kea does not keep the start time of the lease's state, nor the local/remote information, so it cannot emit the corresponding `start-time-of-state` and `data-source` options. Kea does not support VPNs so the presence of option 221 in the query is considered a (`NotAllowed`) error.

17.16.6 DHCPv6 Bulk Leasequery

DHCPv6 Bulk Leasequery gives a requester the ability to query for active lease information over a TCP connection. This allows the server to return all active leases matching a query.

New query types are available: `query-by-relay-id` (3), `query-by-link-address` (4), and `query-by-remote-id` (5).

A new status code, `STATUS_QueryTerminated` (11), has been defined but it is not yet used by the hook library.

Note

Kea attempts to map link address parameters to the prefixes of configured subnets. If a given address falls outside all configured subnet prefixes, the query fails with a status code of `STATUS_NotConfigured`. If the link address parameter for `query-by-relay-id` or `query-by-remote-id` is not `::` (i.e. not empty), only delegated prefixes

that lie within matching subnet prefixes are returned. Currently, `query-by-address` does not support finding delegated prefixes by specifying an address that lies within the prefix.

17.16.7 Bulk Leasequery Configuration

Bulk Leasequery configuration is specified via a new map parameter, `advanced`, with these possible entries:

- `bulk-query-enabled`
When `true`, Kea accepts connections from IP addresses in the requesters list and processes received bulk leasequeries. The default is `false`.
- `active-query-enabled`
This is an anticipated parameter: if set, it must be `false`.
- `extended-info-tables-enabled`
When `true`, the lease backend manages DHCPv6 lease extended info (relay info) in tables to support the new DHCPv6 Bulk Leasequery by-relay-id and by-remote-id types. The default is to use the same value as `bulk-query-enabled`.
- `lease-query-ip`
This is the IP address upon which to listen for connections. The address must be of the same family as the server, e.g. IPv6 for the DHCPv6 server.
- `lease-query-port`
This is the port upon which to listen. The default is 67 for IPv4 and 547 for IPv6, i.e. the same value as for the UDP DHCP service, but for TCP.
- `max-bulk-query-threads`
This indicates the maximum number of threads that bulk leasequery processing should use. A value of 0 instructs the server to use the same number of threads that the Kea core is using for DHCP multi-threading. The default is 0.
- `max-requester-connections`
This is the maximum number of concurrent requester connections. The default is 10; the value must be greater than 0.
- `max-concurrent-queries`
This is the maximum number of concurrent queries per connection. The value 0 allows Kea to determine the number, and is the default.
- `max-requester-idle-time`
This is the amount of time that may elapse after receiving data from a requester before its connection is closed as idle, in seconds. The default is 300.
- `max-leases-per-fetch`
This is the maximum number of leases to return in a single fetch. The default is 100.

Once TLS is supported, we expect to implement common TLS parameters.

For instance, for DHCPv4:

```
{
  "hooks-libraries": [
    {
      "library": "lib/kea/hooks/libdhcp_lease_query.so",
      "parameters": {
        "requesters": [ "192.0.2.1", "192.0.2.2" ],
        "advanced" : {
          "bulk-query-enabled": true,
          "active-query-enabled": false,

          "lease-query-ip": "127.0.0.1",
          "lease-query-tcp-port": 67,

          "max-bulk-query-threads": 0,
          "max-requester-connections": 10,
          "max-concurrent-queries": 4,
          "max-requester-idle-time": 300,
          "max-leases-per-fetch": 100
        }
      }
    }
  ],
  ...
}
```

or for DHCPv6:

```
{
  "hooks-libraries": [
    {
      "library": "lib/kea/hooks/libdhcp_lease_query.so",
      "parameters": {
        "requesters": [ "2001:db8:1::1", "2001:db8:2::1" ],
        "advanced" : {
          "bulk-query-enabled": true,
          "active-query-enabled": false,

          "extended-info-tables-enabled": true,

          "lease-query-ip": "::1",
          "lease-query-tcp-port": 547,

          "max-bulk-query-threads": 0,
          "max-requester-connections": 10,
          "max-concurrent-queries": 4,
          "max-requester-idle-time": 300,
          "max-leases-per-fetch": 100
        }
      }
    }
  ],
  ...
}
```

17.16.8 Updating Existing Leases in SQL Lease Backends

Bulk Leasequery requires additions to the lease data that is stored. With SQL lease backends, leases created prior to the server being configured for bulk lease query do not contain the new data required. In order to populate this data, it is necessary to run these API commands:

Note

Existing leases must have been created by Kea with `store-extended-info` enabled, in order for the new data from extended info to be extracted and stored.

For DHCPv4 lease data, the command is:

```
{
  "command": "extended-info4-upgrade"
}
```

For DHCPv6 lease data, the command is:

For extended info used by relay ID and by remote ID, the command is:

```
{
  "command": "extended-info6-upgrade"
}
```

In all cases, the response indicates whether it succeeded or failed and includes either the count of leases updated or the nature of the failure:

```
{
  "result": 0,
  "text": "Upgraded 1000 leases"
}
```

This `extended-info6-upgrade` command must be called when:

- the database schema was upgraded from 2.4.1 or older version. `extended-info4-upgrade` can be used when upgrading from 2.3.8 or older version.
- Bulk Leasequery was not enabled; tables are maintained only when v6 BLQ is enabled.
- data in tables does not seem to be consistent; tables are not maintained in an atomic way, so consistency is not guaranteed. For instance, when a database is shared between several servers, races can happen between updates.

The operation of the extended info command is governed by the `extended-info-checks` parameter, under the `sanity-checks` element. Please see *Sanity Checks in DHCPv4* or *Sanity Checks in DHCPv6*.

For large numbers of leases, this command may take some time to complete.

Note

Both the `extended-info4-upgrade` and `extended-info6-upgrade` commands are intended for one-time use after upgrading to a version of Kea that fully supports Bulk Leasequery in database backends. These commands are not meant for regular use and will be removed in future Kea versions.

17.17 libdhcp_legal_log.so: Forensic Logging

The Forensic Logging hook library provides hooks that record a detailed log of assignments, renewals, releases, and other lease events into a set of log files.

Note

libdhcp_legal_log.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

In many legal jurisdictions, companies - especially ISPs - must record information about the addresses they have leased to DHCP clients. This library is designed to help with that requirement. If the information that it records is sufficient, it may be used directly.

If a jurisdiction requires that different information be saved, users may use the custom formatting capability to extract information from the inbound request packet, or from the outbound response packet. Administrators are advised to use this feature with caution, as it may affect server performance. The custom format cannot be used for control channel commands.

Alternatively, this library may be used as a template or an example for the user's own custom logging hook. The logging is done as a set of hooks to allow it to be customized to any particular need; modifying a hook library is easier and safer than updating the core code. In addition, by using the hooks features, users who do not need to log this information can leave it out and avoid any performance penalties.

17.17.1 Log File Naming

The names of the log files follow a set pattern.

If using day, month, or year as the time unit, the file name follows the format:

```
path/base-name.CCYYMMDD.txt
```

where CC represents the century, YY represents the year, MM represents the month, and DD represents the day.

If using second as the time unit the file name follows the format:

```
path/base-name.TXXXXXXXXXXXXXXXXXXXXX.txt
```

where XXXXXXXXXXXXXXXXXXXX represents the time in seconds since the beginning of the UNIX epoch.

When using second as the time unit, the file is rotated when the count number of seconds pass. In contrast, when using day, month, or year as the time unit, the file is rotated whenever the count of day, month, or year starts, as applicable.

The "path" and "base-name" are supplied in the configuration as described below; see *Configuring the Forensic Logging Hooks*.

Note

When running Kea servers for both DHCPv4 and DHCPv6, the log names must be distinct. See the examples in *Configuring the Forensic Logging Hooks*.

17.17.2 Configuring the Forensic Logging Hooks

To use this functionality, the hook library must be included in the configuration of the desired DHCP server modules. *libdhcp_legal_log.so* can save logs to a text file or to a database (created using *kea-admin*; see *First-Time Creation of the MySQL Database* and *First-Time Creation of the PostgreSQL Database*). The library is installed alongside the Kea libraries in `[kea-install-dir]/var/log/kea`, where *kea-install-dir* is determined by the `--prefix` meson setup option which defaults to `/usr/local`. Assuming the default value, *kea-dhcp4* can be configured to load *libdhcp_legal_log.so* like this:

```
{
  "Dhcp4": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "path": "/var/log/kea",
          "base-name": "kea-forensic4"
        }
      }
    ]
  }
}
```

For *kea-dhcp6*, the configuration is:

```
{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "path": "/var/log/kea",
          "base-name": "kea-forensic6"
        }
      }
    ]
  }
}
```

The hook library parameters for the text file configuration are:

- `path` - the directory in which the forensic file(s) will be written. The default value is `[kea-install-dir]/var/log/kea`. The directory must exist.

Note

As of Kea 2.7.9, legal log files may only be written to the output directory determined during compilation as: `"[kea-install-dir]/var/log/kea"`. This path may be overridden at startup by setting the environment variable `KEA_LEGAL_LOG_DIR` to the desired path. If a path other than this value is used in `path`, Kea will emit an

error and refuse to start or if already running, log an unrecoverable error. For ease of use simply omit the `path` parameter.

- `base-name` - an arbitrary value which is used in conjunction with the current system date to form the current forensic file name. It defaults to `kea-legal`.
- `time-unit` - configures the time unit used to rotate the log file. Valid values are `second`, `day`, `month`, or `year`. It defaults to `day`.
- `count` - configures the number of time units that need to pass until the log file is rotated. It can be any positive number, or 0, which disables log rotation. It defaults to 1.

If log rotation is disabled, a new file is created when the library is loaded; the new file name is different from any previous file name.

Additional actions can be performed just before closing the old file and after opening the new file. These actions must point to an external executable or script and are configured with the following settings:

- `prerotate` - an external executable or script called with the name of the file that will be closed. Kea does not wait for the process to finish.
- `postrotate` - an external executable or script called with the name of the file that was opened. Kea does not wait for the process to finish.

These executables must be stored in the "[`kea-install-dir`]/share/kea/scripts/" directory which can be overridden at startup by setting the environment variable `KEA_HOOK_SCRIPTS_PATH` to a different path.

Custom formatting can be enabled for logging information that can be extracted either from the client's request packet or from the server's response packet. Use with caution as this might affect server performance. The custom format cannot be used for control channel commands. Two parameters can be used towards this goal, either together or separately:

- `request-parser-format` - an evaluated parsed expression used to extract and log data from the incoming packet.
- `response-parser-format` - an evaluated parsed expression used to extract and log data from the server response packet.

See *Using Expressions in Classification* for a list of expressions. If either `request-parser-format` or `response-parser-format` is configured, the default logging format is not used. If both of them are configured, the resulting log message is constructed by concatenating the data extracted from the request and the data extracted from the response.

The custom formatting permits logging on multiple lines using the hexstring `0x0a` (ASCII code for new line). In the log file, each line is prepended with the log timestamp. For the database backend, the data is stored (including the newline character) in the same entry.

Examples:

```
{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "path": "/var/log/kea",
          "base-name": "kea-forensic6",
          "request-parser-format": "'first line' + 0x0a + 'second line'",
          "response-parser-format": "'also second line' + 0x0a + 'third line'"
        }
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

Some data might be available in the request or only in the response; the data in the request packet might differ from that in the response packet.

The lease-client context can only be printed using the default format, as this information is not directly stored in the request packet or in the response packet.

The `timestamp-format` parameter can be used to change the timestamp logged at the beginning of each line. Permissible formatting is the one supported by `strftime` plus the `'%Q'` extra format which adds the microseconds subunits. The default is: `"%Y-%m-%d %H:%M:%S %Z"`. This parameter has no effect for the database backends, where the timestamp is defined at the schema level.

Examples:

```

{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "path": "/var/log/kea",
          "base-name": "kea-forensic6",
          "timestamp-format": "%H%M%S.%F%Q"
        }
      }
    ]
  }
}

```

Additional parameters for the database connection can be specified, e.g:

```

{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "name": "database-name",
          "password": "1234",
          "type": "mysql",
          "user": "user-name"
        }
      }
    ]
  }
}

```

For more specific information about database-related parameters, please refer to [Lease Database Configuration](#) and [Lease Database Configuration](#).

If it is desired to restrict forensic logging to certain subnets, the `"legal-logging"` boolean parameter can be specified

within a user context of these subnets. For example:

```
{
  "Dhcp4": {
    "subnet4": [
      {
        "id": 1,
        "subnet": "192.0.2.0/24",
        "pools": [
          {
            "pool": "192.0.2.1 - 192.0.2.200"
          }
        ],
        "user-context": {
          "legal-logging": false
        }
      }
    ]
  }
}
```

This configuration disables legal logging for the subnet "192.0.2.0/24". If the "legal-logging" parameter is not specified, it defaults to `true`, which enables legal logging for the subnet.

The following example demonstrates how to selectively disable legal logging for an IPv6 subnet:

```
{
  "Dhcp6": {
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8:1::/64",
        "pools": [
          {
            "pool": "2001:db8:1::1-2001:db8:1::ffff"
          }
        ],
        "user-context": {
          "legal-logging": false
        }
      }
    ]
  }
}
```

See *User Contexts in IPv4* and *User Contexts in IPv6* to learn more about user contexts in Kea configuration.

17.17.3 DHCPv4 Log Entries

For DHCPv4, the library creates entries based on DHCPREQUEST, DHCPDECLINE, and DHCPRELEASE messages, et al., and their responses. The resulting packets and leases are taken into account, intercepted through the following hook points:

- `pkt4_receive`
- `leases4_committed`

- pkt4_send
- lease4_release
- lease4_decline

An entry is a single string with no embedded end-of-line markers and a prepended timestamp, and has the following sections:

```
timestamp address duration device-id {client-info} {relay-info} {user-context}
```

Where:

- **timestamp** - the date and time the log entry was written, in "%Y-%m-%d %H:%M:%S %Z" strftime format ("%Z" is the time zone name).
- **address** - the leased IPv4 address given out, and whether it was assigned, renewed, or released.
- **duration** - the lease lifetime expressed in days (if present), hours, minutes, and seconds. A lease lifetime of 0xFFFFFFFF will be denoted with the text "infinite duration." This information is not given when the lease is released.
- **device-id** - the client's hardware address shown as a numerical type and hex-digit string.
- **client-info** - the DHCP client id option (61) if present, shown as a hex string. When its content is printable it is displayed.
- **relay-info** - for relayed packets, the giaddr and the RAI circuit-id, remote-id, and subscriber-id options (option 82 sub options: 1, 2 and 6), if present. The circuit-id and remote-id are presented as hex strings. When their content is printable it is displayed.
- **user-context** - the optional user context associated with the lease.

For instance (line breaks are added here for readability; they are not present in the log file):

```
2018-01-06 01:02:03 CET Address: 192.2.1.100 has been renewed for 1 hrs 52 min 15 secs_
↳to a device with hardware address:
hwtype=1 08:00:2b:02:3f:4e, client-id: 17:34:e2:ff:09:92:54 connected via relay at_
↳address: 192.2.16.33,
identified by circuit-id: 68:6f:77:64:79 (howdy) and remote-id: 87:f6:79:77:ef
```

or for a release:

```
2018-01-06 01:02:03 CET Address: 192.2.1.100 has been released from a device with_
↳hardware address:
hwtype=1 08:00:2b:02:3f:4e, client-id: 17:34:e2:ff:09:92:54 connected via relay at_
↳address: 192.2.16.33,
identified by circuit-id: 68:6f:77:64:79 (howdy) and remote-id: 87:f6:79:77:ef
```

In addition to logging lease activity driven by DHCPv4 client traffic, the hook library also logs entries for the following lease management control channel commands: *lease4-add*, *lease4-update*, and *lease4-del*. These cannot have custom formatting. Each entry is a single string with no embedded end-of-line markers, and it will typically have the following form:

lease4-add:

```
*timestamp* Administrator added a lease of address: *address* to a device with hardware_
↳address: *device-id*
```

Depending on the arguments of the add command, it may also include the client-id and duration.

Example:

```
2018-01-06 01:02:03 CET Administrator added a lease of address: 192.0.2.202 to a device.
↳with hardware address:
1a:1b:1c:1d:1e:1f for 1 days 0 hrs 0 mins 0 secs
```

lease4-update:

```
*timestamp* Administrator updated information on the lease of address: *address* to a
↳device with hardware address: *device-id*
```

Depending on the arguments of the update command, it may also include the client-id and lease duration.

Example:

```
2018-01-06 01:02:03 CET Administrator updated information on the lease of address: 192.0.
↳2.202 to a device
with hardware address: 1a:1b:1c:1d:1e:1f, client-id: 1234567890
```

lease4-del: deletes have two forms, one by address and one by identifier and identifier type:

```
*timestamp* Administrator deleted the lease for address: *address*
```

or

```
*timestamp* Administrator deleted a lease for a device identified by: *identifier-type*
↳of *identifier*
```

Currently only a type of @b hw-address (hardware address) is supported.

Examples:

```
2018-01-06 01:02:03 CET Administrator deleted the lease for address: 192.0.2.202

2018-01-06 01:02:12 CET Administrator deleted a lease for a device identified by: hw-
↳address of 1a:1b:1c:1d:1e:1f
```

If the High Availability hook library is enabled, the partner periodically sends lease commands in a similar format; the only difference is that the issuer of the command is "HA partner" instead of "Administrator."

```
*timestamp* HA partner added ...
```

or

```
*timestamp* HA partner updated ...
```

or

```
*timestamp* HA partner deleted ...
```

The request-parser-format and response-parser-format options can be used to extract and log data from the incoming packet and server response packet, respectively. The configured value is an evaluated parsed expression returning a string. A list of tokens is described in the server classification process. Use with caution as this might affect server performance. If either of them is configured, the default logging format is not used. If both of them are configured, the resulting log message is constructed by concatenating the logged data extracted from the request and the logged data extracted from the response.

The custom formatting permits logging on multiple lines using the hexstring 0x0a (ASCII code for new line). In the case of the log file, each line is prepended with the log timestamp. For the database backend, the data is stored (including the newline character) in the same entry.

Examples:

```
{
  "Dhcp4": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "name": "database-name",
          "password": "1234",
          "type": "mysql",
          "user": "user-name",
          "request-parser-format": "'log entry' + 0x0a + 'same log entry'",
          "response-parser-format": "'also same log entry' + 0x0a + 'again same log entry'"
        }
      }
    ]
  }
}
```

Some data might be available in the request or in the response only, and some data might differ in the incoming packet from the one in the response packet.

Examples:

```
{
  "request-parser-format": "ifelse(pkt4.msgtype == 4 or pkt4.msgtype == 7, 'Address: ' +
  + ifelse(option[50].exists, addrtotext(option[50].hex), addrtotext(pkt4.ciaddr)) + '
  has been released from a device with hardware address: hwtype=' +
  substring(hexstring(pkt4.htype, ''), 7, 1) + ' ' + hexstring(pkt4.mac, ':') +
  ifelse(option[61].exists, ', client-id: ' + hexstring(option[61].hex, ':'), '') +
  ifelse(pkt4.giaddr == 0.0.0.0, '', ' connected via relay at address: ' +
  addrtotext(pkt4.giaddr) + ifelse(option[82].option[1].exists, ', circuit-id: ' +
  hexstring(option[82].option[1].hex, ':'), '') + ifelse(option[82].option[2].exists, ',
  remote-id: ' + hexstring(option[82].option[2].hex, ':'), '') + ifelse(option[82].
  option[6].exists, ', subscriber-id: ' + hexstring(option[82].option[6].hex, ':'), '')),
  '')",
  "response-parser-format": "ifelse(pkt4.msgtype == 5, 'Address: ' + addrtotext(pkt4.
  yiaddr) + ' has been assigned for ' + uint32totext(option[51].hex) + ' seconds to a
  device with hardware address: hwtype=' + substring(hexstring(pkt4.htype, ''), 7, 1) +
  ' ' + hexstring(pkt4.mac, ':') + ifelse(option[61].exists, ', client-id: ' +
  hexstring(option[61].hex, ':'), '') + ifelse(pkt4.giaddr == 0.0.0.0, '', ' connected
  via relay at address: ' + addrtotext(pkt4.giaddr) + ifelse(option[82].option[1].exists,
  ', circuit-id: ' + hexstring(option[82].option[1].hex, ':'), '') + ifelse(option[82].
  option[2].exists, ', remote-id: ' + hexstring(option[82].option[2].hex, ':'), '') +
  ifelse(option[82].option[6].exists, ', subscriber-id: ' + hexstring(option[82].
  option[6].hex, ':'), ''), '')"
}
```

Details:

This will log the following data on request and renew:

```
Address: 192.2.1.100 has been assigned for 6735 seconds to a device with hardware_
↳ address: hwtype=1 08:00:2b:02:3f:4e, client-id: 17:34:e2:ff:09:92:54 connected via_
↳ relay at address: 192.2.16.33, circuit-id: 68:6f:77:64:79, remote-id: 87:f6:79:77:ef,_
↳ subscriber-id: 1a:2b:3c:4d:5e:6f
```

This will log the following data on release and decline:

```
Address: 192.2.1.100 has been released from a device with hardware address: hwtype=1_
↳ 08:00:2b:02:3f:4e, client-id: 17:34:e2:ff:09:92:54 connected via relay at address: 192._
↳ 2.16.33, circuit-id: 68:6f:77:64:79, remote-id: 87:f6:79:77:ef, subscriber-id:_
↳ 1a:2b:3c:4d:5e:6f
```

A similar result can be obtained by configuring only request-parser-format.

Examples:

```
{
  "request-parser-format": "ifelse(pkt4.msgtype == 3, 'Address: ' + ifelse(option[50].
↳ exists, addrtotext(option[50].hex), addrtotext(pkt4.ciaddr)) + ' has been assigned' +_
↳ ifelse(option[51].exists, ' for ' + uint32totext(option[51].hex) + ' seconds', '') + '_
↳ to a device with hardware address: hwtype=' + substring(hexstring(pkt4.htype, ''), 7,_
↳ 1) + ' ' + hexstring(pkt4.mac, ':') + ifelse(option[61].exists, ', client-id: ' +_
↳ hexstring(option[61].hex, ':'), '') + ifelse(pkt4.giaddr == 0.0.0.0, '', ' connected_
↳ via relay at address: ' + addrtotext(pkt4.giaddr) + ifelse(option[82].option[1].exists,
↳ ', circuit-id: ' + hexstring(option[82].option[1].hex, ':'), '') + ifelse(option[82].
↳ option[2].exists, ', remote-id: ' + hexstring(option[82].option[2].hex, ':'), '') +_
↳ ifelse(option[82].option[6].exists, ', subscriber-id: ' + hexstring(option[82].
↳ option[6].hex, ':'), '')), ifelse(pkt4.msgtype == 4 or pkt4.msgtype == 7, 'Address: '_
↳ + ifelse(option[50].exists, addrtotext(option[50].hex), addrtotext(pkt4.ciaddr)) + '_
↳ has been released from a device with hardware address: hwtype=' +_
↳ substring(hexstring(pkt4.htype, ''), 7, 1) + ' ' + hexstring(pkt4.mac, ':') +_
↳ ifelse(option[61].exists, ', client-id: ' + hexstring(option[61].hex, ':'), '') +_
↳ ifelse(pkt4.giaddr == 0.0.0.0, '', ' connected via relay at address: ' +_
↳ addrtotext(pkt4.giaddr) + ifelse(option[82].option[1].exists, ', circuit-id: ' +_
↳ hexstring(option[82].option[1].hex, ':'), '') + ifelse(option[82].option[2].exists, ',_
↳ remote-id: ' + hexstring(option[82].option[2].hex, ':'), '') + ifelse(option[82].
↳ option[6].exists, ', subscriber-id: ' + hexstring(option[82].option[6].hex, ':'), '')),
↳ '')'"
}
```

Details:

17.17.4 DHCPv6 Log Entries

For DHCPv6, the library creates entries based on REQUEST, RENEW, RELEASE, and DECLINE messages, et al. and their responses. The resulting packets and leases are taken into account, intercepted through the following hook points:

- pkt6_receive
- leases6_committed
- pkt6_send
- lease6_release

- lease6_decline

An entry is a single string with no embedded end-of-line markers and a prepended timestamp, and has the following sections:

```
timestamp address duration device-id {relay-info}* {user-context}
```

Where:

- **timestamp** - the date and time the log entry was written, in "%Y-%m-%d %H:%M:%S %Z" strftime format ("%Z" is the time zone name).
- **address** - the leased IPv6 address or prefix given out, and whether it was assigned, renewed, or released.
- **duration** - the lease lifetime expressed in days (if present), hours, minutes, and seconds. A lease lifetime of 0xFFFFFFFF will be denoted with the text "infinite duration." This information is not given when the lease is released.
- **device-id** - the client's DUID and hardware address (if present).
- **relay-info** - for relayed packets the content of relay agent messages, and the **remote-id** (code 37), **subscriber-id** (code 38), and **interface-id** (code 18) options, if present. Note that the **interface-id** option, if present, identifies the whole interface on which the relay agent received the message. This typically translates to a single link in the network, but it depends on the specific network topology. Nevertheless, this is useful information to better pinpoint the location of the device, so it is recorded, if present.
- **user-context** - the optional user context associated with the lease.

For instance (line breaks are added here for readability; they are not present in the log file):

```
2018-01-06 01:02:03 PST Address:2001:db8:1:: has been assigned for 0 hrs 11 mins 53 secs
to a device with DUID: 17:34:e2:ff:09:92:54 and hardware address: hwtype=1
↪08:00:2b:02:3f:4e
(from Raw Socket) connected via relay at address: fe80::abcd for client on link address:
↪3001::1,
hop count: 1, identified by remote-id: 01:02:03:04:0a:0b:0c:0d:0e:0f and subscriber-id:
↪1a:2b:3c:4d:5e:6f
```

or for a release:

```
2018-01-06 01:02:03 PST Address:2001:db8:1:: has been released
from a device with DUID: 17:34:e2:ff:09:92:54 and hardware address: hwtype=1
↪08:00:2b:02:3f:4e
(from Raw Socket) connected via relay at address: fe80::abcd for client on link address:
↪3001::1,
hop count: 1, identified by remote-id: 01:02:03:04:0a:0b:0c:0d:0e:0f and subscriber-id:
↪1a:2b:3c:4d:5e:6f
```

In addition to logging lease activity driven by DHCPv6 client traffic, the hook library also logs entries for the following lease management control channel commands: *lease6-add*, *lease6-update*, and *lease6-del*. Each entry is a single string with no embedded end-of-line markers, and it will typically have the following form:

lease6-add:

```
*timestamp* Administrator added a lease of address: *address* to a device with DUID:
↪*DUID*
```

Depending on the arguments of the add command, it may also include the hardware address and duration.

Example:

```
2018-01-06 01:02:03 PST Administrator added a lease of address: 2001:db8::3 to a device ↵
↵with DUID:
1a:1b:1c:1d:1e:1f:20:21:22:23:24 for 1 days 0 hrs 0 mins 0 secs
```

lease6-update:

```
*timestamp* Administrator updated information on the lease of address: *address* to a ↵
↵device with DUID: *DUID*
```

Depending on the arguments of the update command, it may also include the hardware address and lease duration.

Example:

```
2018-01-06 01:02:03 PST Administrator updated information on the lease of address: ↵
↵2001:db8::3 to a device with
DUID: 1a:1b:1c:1d:1e:1f:20:21:22:23:24, hardware address: 1a:1b:1c:1d:1e:1f
```

lease6-del: deletes have two forms, one by address and one by identifier and identifier type:

```
*timestamp* Administrator deleted the lease for address: *address*
```

or

```
*timestamp* Administrator deleted a lease for a device identified by: *identifier-type* ↵
↵of *identifier*
```

Currently only a type of DUID is supported.

Examples:

```
2018-01-06 01:02:03 PST Administrator deleted the lease for address: 2001:db8::3

2018-01-06 01:02:11 PST Administrator deleted a lease for a device identified by: duid ↵
↵of 1a:1b:1c:1d:1e:1f:20:21:22:23:24
```

If the High Availability hook library is enabled, the partner periodically sends lease commands in a similar format; the only difference is that the issuer of the command is "HA partner" instead of "Administrator."

```
*timestamp* HA partner added ...
```

or

```
*timestamp* HA partner updated ...
```

or

```
*timestamp* HA partner deleted ...
```

The `request-parser-format` and `response-parser-format` options can be used to extract and log data from the incoming packet and server response packet, respectively. The configured value is an evaluated parsed expression returning a string. A list of tokens is described in the server classification process. Use with caution as this might affect server performance. If either of them is configured, the default logging format is not used. If both of them are configured, the resulting log message is constructed by concatenating the logged data extracted from the request and the logged data extracted from the response.

The custom formatting permits logging on multiple lines using the hexstring 0x0a (ASCII code for new line). In the case of the log file, each line is prepended with the log timestamp. For the database backend, the data is stored (including the newline character) in the same entry.

Examples:

```
{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_legal_log.so",
        "parameters": {
          "name": "database-name",
          "password": "1234",
          "type": "mysql",
          "user": "user-name",
          "request-parser-format": "'log entry' + 0x0a + 'same log entry'",
          "response-parser-format": "'also same log entry' + 0x0a + 'again same log entry'"
        }
      }
    ]
  }
}
```

Some data might be available in the request or in the response only, and some data might differ in the incoming packet from the one in the response packet.

Notes:

In the case of IPv6, the packets can contain multiple IA_NA (3) or IA_PD (25) options, each containing multiple options, including OPTION_IAADDR (5) or OPTION_IAPREFIX (25) suboptions. To be able to print the current lease associated with the log entry, the forensic log hook library internally isolates the corresponding IA_NA or IA_PD option and respective suboption matching the current lease. The hook library will iterate over all new allocated addresses and all deleted addresses, making each address available for logging as the current lease for the respective logged entry.

They are accessible using the following parser expressions:

Current lease associated with OPTION_IAADDR:

```
addrtotext(substring(option[3].option[5].hex, 0, 16))
```

Current lease associated with OPTION_IAPREFIX:

```
addrtotext(substring(option[25].option[26].hex, 9, 16))
```

All other parameters of the options are available at their respective offsets in the option. Please read RFC8415 for more details.

Examples:

```
{
  "request-parser-format": "ifelse(pkt6.msgtype == 8 or pkt6.msgtype == 9,
↳ ifelse(option[3].option[5].exists, 'Address: ' + addrtotext(substring(option[3].
↳ option[5].hex, 0, 16)) + ' has been released from a device with DUID: ' +
↳ hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '', '', ' connected via
↳ relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for client on link address: '
↳
```

(continues on next page)

(continued from previous page)

```

→+ addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists, ', remote-id: ' +
→+ hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].option[38].exists,
→', subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '') + ifelse(relay6[0].
→option[18].exists, ', connected at location interface-id: ' + hexstring(relay6[0].
→option[18].hex, ':'), ''), '') + ifelse(option[25].option[26].exists, 'Prefix: ' +
→addrtotext(substring(option[25].option[26].hex, 9, 16)) + '/' +
→uint8totext(substring(option[25].option[26].hex, 8, 1)) + ' has been released from a
→device with DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '',
→'', ' connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for
→client on link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].
→option[37].exists, ', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') +
→ifelse(relay6[0].option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].
→option[38].hex, ':'), '') + ifelse(relay6[0].option[18].exists, ', connected at
→location interface-id: ' + hexstring(relay6[0].option[18].hex, ':'), ''), ''), ''),
  "response-parser-format": "ifelse(pkt6.msgtype == 7, ifelse(option[3].option[5].
→exists and not (substring(option[3].option[5].hex, 20, 4) == 0), 'Address: ' +
→addrtotext(substring(option[3].option[5].hex, 0, 16)) + ' has been assigned for ' +
→uint32totext(substring(option[3].option[5].hex, 20, 4)) + ' seconds to a device with
→DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '', '', '
→connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for client on
→link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists,
→', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].
→option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '')
→+ ifelse(relay6[0].option[18].exists, ', connected at location interface-id: ' +
→hexstring(relay6[0].option[18].hex, ':'), ''), '') + ifelse(option[25].option[26].
→exists and not (substring(option[25].option[26].hex, 4, 4) == 0), 'Prefix: ' +
→addrtotext(substring(option[25].option[26].hex, 9, 16)) + '/' +
→uint8totext(substring(option[25].option[26].hex, 8, 1)) + ' has been assigned for ' +
→uint32totext(substring(option[25].option[26].hex, 4, 4)) + ' seconds to a device with
→DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '', '', '
→connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for client on
→link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists,
→', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].
→option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '')
→+ ifelse(relay6[0].option[18].exists, ', connected at location interface-id: ' +
→hexstring(relay6[0].option[18].hex, ':'), ''), ''), ''), '')"
}

```

Details:

This will log the following data on request, renew, and rebind for NA:

```

Address: 2001:db8:1:: has been assigned for 713 seconds to a device with DUID:
→17:34:e2:ff:09:92:54 connected via relay at address: fe80::abcd for client on link
→address: 3001::1, remote-id: 01:02:03:04:0a:0b:0c:0d:0e:0f, subscriber-id:
→1a:2b:3c:4d:5e:6f, connected at location interface-id: 72:65:6c:61:79:31:3a:65:74:68:30

```

This will log the following data on request, renew and rebind for PD:

```

Prefix: 2001:db8:1::/64 has been assigned for 713 seconds to a device with DUID:
→17:34:e2:ff:09:92:54 connected via relay at address: fe80::abcd for client on link
→address: 3001::1, remote-id: 01:02:03:04:0a:0b:0c:0d:0e:0f, subscriber-id:
→1a:2b:3c:4d:5e:6f, connected at location interface-id: 72:65:6c:61:79:31:3a:65:74:68:30

```

This will log the following data on release and decline for NA:

```
Address: 2001:db8:1:: has been released from a device with DUID: 17:34:e2:ff:09:92:54
→connected via relay at address: fe80::abcd for client on link address: 3001::1, remote-
→id: 01:02:03:04:0a:0b:0c:0d:0e:0f, subscriber-id: 1a:2b:3c:4d:5e:6f, connected at
→location interface-id: 72:65:6c:61:79:31:3a:65:74:68:30
```

This will log the following data on release and decline for PD:

```
Prefix: 2001:db8:1::/64 has been released from a device with DUID: 17:34:e2:ff:09:92:54
→connected via relay at address: fe80::abcd for client on link address: 3001::1, remote-
→id: 01:02:03:04:0a:0b:0c:0d:0e:0f, subscriber-id: 1a:2b:3c:4d:5e:6f, connected at
→location interface-id: 72:65:6c:61:79:31:3a:65:74:68:30
```

A similar result can be obtained by configuring only request-parser-format.

Examples:

```
{
  "request-parser-format": "ifelse(pkt6.msgtype == 3 or pkt6.msgtype == 5 or pkt6.
→msgtype == 6, ifelse(option[3].option[5].exists, 'Address: ' +
→addrtotext(substring(option[3].option[5].hex, 0, 16)) + ' has been assigned for ' +
→uint32totext(substring(option[3].option[5].hex, 20, 4)) + ' seconds to a device with
→DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '', '', '
→connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for client on
→link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists,
→', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].
→option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '')
→+ ifelse(relay6[0].option[18].exists, ', connected at location interface-id: ' +
→hexstring(relay6[0].option[18].hex, ':'), '')), '') + ifelse(option[25].option[26].
→exists, 'Prefix: ' + addrtotext(substring(option[25].option[26].hex, 9, 16)) + '/' +
→uint8totext(substring(option[25].option[26].hex, 8, 1)) + ' has been assigned for ' +
→uint32totext(substring(option[25].option[26].hex, 4, 4)) + ' seconds to a device with
→DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '', '', '
→connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for client on
→link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists,
→', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].
→option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '')
→+ ifelse(relay6[0].option[18].exists, ', connected at location interface-id: ' +
→hexstring(relay6[0].option[18].hex, ':'), '')), ''), ifelse(pkt6.msgtype == 8 or pkt6.
→msgtype == 9, ifelse(option[3].option[5].exists, 'Address: ' +
→addrtotext(substring(option[3].option[5].hex, 0, 16)) + ' has been released from a
→device with DUID: ' + hexstring(option[1].hex, ':') + ifelse(relay6[0].peeraddr == '',
→'', ' connected via relay at address: ' + addrtotext(relay6[0].peeraddr) + ' for
→client on link address: ' + addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].
→option[37].exists, ', remote-id: ' + hexstring(relay6[0].option[37].hex, ':'), '') +
→ifelse(relay6[0].option[38].exists, ', subscriber-id: ' + hexstring(relay6[0].
→option[38].hex, ':'), '') + ifelse(relay6[0].option[18].exists, ', connected at
→location interface-id: ' + hexstring(relay6[0].option[18].hex, ':'), '')), '') +
→ifelse(option[25].option[26].exists, 'Prefix: ' + addrtotext(substring(option[25].
→option[26].hex, 9, 16)) + '/' + uint8totext(substring(option[25].option[26].hex, 8,
→1)) + ' has been released from a device with DUID: ' + hexstring(option[1].hex, ':') +
→ifelse(relay6[0].peeraddr == '', '', ' connected via relay at address: ' +
→addrtotext(relay6[0].peeraddr) + ' for client on link address: ' +
```

(continues on next page)

(continued from previous page)

```

→addrtotext(relay6[0].linkaddr) + ifelse(relay6[0].option[37].exists, ', remote-id: ' +
→hexstring(relay6[0].option[37].hex, ':'), '') + ifelse(relay6[0].option[38].exists, ',
→subscriber-id: ' + hexstring(relay6[0].option[38].hex, ':'), '') + ifelse(relay6[0].
→option[18].exists, ', connected at location interface-id: ' + hexstring(relay6[0].
→option[18].hex, ':'), ''), ''), ''), '')"
}

```

Details:

17.17.5 Database Backend

Log entries can be inserted into a database when Kea is configured with database backend support. Kea uses a table named `logs`, that includes a timestamp generated by the database software, and a text log with the same format as files without the timestamp.

Please refer to [MySQL](#) for information on using a MySQL database; or to [PostgreSQL](#) for PostgreSQL database information. The `logs` table is part of the Kea database schemas.

Configuration parameters are extended by standard lease database parameters as defined in [Lease Database Configuration](#). The `type` parameter should be `mysql`, `postgresql`, `logfile` or `syslog`; when it is absent or set to `logfile`, files are used.

No specific tools are provided to operate the database, but standard tools may be used, for example, to dump the logs table from a MySQL database:

```

$ mysql --user keatest --password 1234 -e "select * from logs;"
+-----+-----+-----+-----+
→-----+-----+-----+-----+
→-----+-----+-----+-----+
| timestamp          | address      | log                                     |
→-----+-----+-----+-----+
→          | id |
→-----+-----+-----+-----+
→-----+-----+-----+-----+
| 2022-03-30 17:38:41 | 192.168.50.1 | Address: 192.168.50.1 has been assigned for 0 hrs,
→10 mins 0 secs to a device with hardware address: hwtype=1 ff:01:02:03:ff:04, client-
→id: 00:01:02:03:04:05:06 | 31 |
| 2022-03-30 17:38:43 | 192.168.50.1 | Address: 192.168.50.1 has been assigned for 0 hrs,
→10 mins 0 secs to a device with hardware address: hwtype=1 ff:01:02:03:ff:04, client-
→id: 00:01:02:03:04:05:06 | 32 |
| 2022-03-30 17:38:45 | 192.168.50.1 | Address: 192.168.50.1 has been assigned for 0 hrs,
→10 mins 0 secs to a device with hardware address: hwtype=1 ff:01:02:03:ff:04, client-
→id: 00:01:02:03:04:05:06 | 33 |
+-----+-----+-----+-----+
→-----+-----+-----+-----+
→-----+-----+-----+-----+

```

Like all the other database-centric features, forensic logging supports database connection recovery, which can be enabled by setting the `on-fail` parameter. If not specified, the `on-fail` parameter in forensic logging defaults to `serve-retry-continue`. This is different than for [libdhcp_lease_cmds.so](#), [libdhcp_host_cmds.so](#), and [libdhcp_cb_cmds.so](#), where `on-fail` defaults to `stop-retry-exit`. In this case, the server continues serving clients and does not shut down even if the recovery mechanism fails. If `on-fail` is set to `serve-retry-exit`, the server will shut down if the connection to the database backend is not restored according to the `max-reconnect-tries` and `reconnect-wait-time` parameters, but it continues serving clients while this mechanism is activated.

During server startup, the inability to connect to any of the configured backends is considered fatal only if `retry-on-startup` is set to `false` (the default). A fatal error is logged and the server exits, based on the idea that the configuration should be valid at startup. Exiting to the operating system allows nanny scripts to detect the problem. If `retry-on-startup` is set to `true`, the server starts reconnection attempts even at server startup or on reconfigure events, and honors the action specified in the `on-fail` parameter.

17.17.6 Syslog Backend

Log entries can be inserted into syslog by setting the type to `syslog`. When syslog type is configured, the `pattern` parameter specifies the details that are used for logging. For more details see [Logging](#). If not configured, it defaults to:

```
"%-5p [%c.%t] %m\n"
```

The `facility` parameter specifies the syslog facility and it defaults to `local0`.

17.18 libdhcp_limits.so: Limits to Manage Lease Allocation and Packet Processing

This hook library enables two types of limits:

1. Lease limiting: allow a maximum of `n` leases assigned at any one time.
2. Rate limiting: allow a maximum of `n` packets per `time_unit` to receive a response.

Note

libdhcp_limits.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

17.18.1 Configuration

The following examples are for *kea-dhcp6*, but they apply equally to *kea-dhcp4*. The wildcards "`<limit-type>`" and "`<limit-value>`" need to be replaced with the respective keys and values for each limit type described in the sections following this one.

The library can be loaded by both *kea-dhcp4* and *kea-dhcp6* servers by adding its path in the `"hooks-libraries"` element of the server's configuration.

```
{
  "Dhcp6": {
    "hooks-libraries": [
      {
        "library": "/usr/local/lib/libdhcp_limits.so"
      }
    ]
  }
}
```

This alone does not limit anything. The desired limits are added to the user context in the configuration portion of the element that identifies the clients to be limited: a client class or a subnet. Upon reconfiguration, if Kea picked up on the configured limits, it logs one line for each configured limit. The log message contains `LIMITS_CONFIGURED` in its identifier.

This is how a lease limit is defined for a client class:

```
{
  "Dhcp6": {
    "client-classes": [
      {
        "name": "cable-modem-1",
        "test": "option[123].hex == 0x000C4B1E",
        "user-context": {
          "limits": {
            "<limit>": "<limit-value>"
          }
        }
      }
    ]
  }
}
```

This is how a lease limit is defined for a global subnet:

```
{
  "Dhcp6": {
    "subnet6": [
      {
        "id": 1,
        "subnet": "2001:db8::/64",
        "user-context": {
          "limits": {
            "<limit>": "<limit-value>"
          }
        }
      }
    ]
  }
}
```

This is how a lease limit is defined for a subnet inside a shared network:

```
{
  "Dhcp6": {
    "shared-networks": [
      {
        "subnet6": [
          {
            "id": 1,
            "subnet": "2001:db8::/64",
            "user-context": {
              "limits": {
                "<limit>": "<limit-value>"
              }
            }
          }
        ]
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

Note

The Limits hook library uses the class name to identify a client class and the subnet ID to identify a subnet. Changing a test expression in a client class or the network range of a subnet while leaving the name or ID unchanged does not reset the lease count for the respective client class or subnet. To reset the lease count, change the client class name or the subnet ID.

17.18.2 Lease Limiting

It is possible to limit the number of leases that a group of clients can get from a Kea DHCP server or from a set of collaborating Kea DHCP servers.

The value of a lease limit can be specified as an unsigned integer in 32 bits, i.e. between 0 and 4,294,967,295. Each lease type can be limited individually. IPv4 leases and IPv6 IA_NA leases are limited through the "address-limit" configuration entry. IPv6 IA_PD leases are limited through the "prefix-limit" configuration entry. Here are some examples:

- "address-limit": 4
- "prefix-limit": 2

For lease limiting, client classes and the associated lease counts - which are checked against the configured limits - are updated for each lease in the following hook callouts:

- lease4_select
- lease4_renew
- lease6_select
- lease6_renew
- lease6_rebind

As a result, classes for which "only-in-additional-list" is "true" cannot be lease-limited. Please refer to [the classification steps](#) for more information on which client classes can be used to limit the number of leases.

Note

Under load, a Kea DHCP server may allocate more leases than the limit strictly allows. This only has a chance of happening during high traffic surges, coming from clients belonging to the same class or the same subnet, depending on what is limited. Users may be interested in following the development of [atomic lease limits](#) in ISC's GitLab instance.

17.18.3 Rate Limiting

It is possible to limit the frequency or rate at which inbound packets receive a response.

The value of a rate limit can be specified in the format "<p> packets per <time-unit>". <p> is any number that can be represented by an unsigned integer in 32 bits, i.e. between 0 and 4,294,967,295. <time-unit> can be any of second, minute, hour, day, week, month, or year. A month is considered to be 30 days for simplicity; similarly, a year is 365 days for limiting purposes. This syntax covers a wide range of rates, from one lease per year to four billion leases per second. This value is assigned to the "rate-limit" configuration entry. Here are some examples:

- "rate-limit": 1 packet per second
- "rate-limit": 4 packets per minute
- "rate-limit": 16 packets per hour

The configured value of 0 packets is a convenient way of disabling packet processing for certain clients entirely. As such, it means its literal value and is not a special value for disabling limiting altogether, as might be imagined. Disabling limiting entirely is achieved by removing the "rate-limit" leaf configuration entry, the "limits" map or user context around it, or the hook library configuration. The same applies to the value of 0 in lease limiting. However, that use case is best achieved with rate limiting; it puts less computational strain on Kea, since the action of dropping the request or sending a NAK is decided earlier.

In terms of rate limiting, client classes are evaluated at the `pkt4_receive` and the `pkt6_receive` callout, respectively, so that rate limits are checked as early as possible in the packet-processing cycle. Thus, only those classes which are assigned to the packet solely via an independent test expression can be used. Classes that depend on host reservations or the special BOOTP or KNOWN classes, and classes that are marked with "only-in-additional-list": true, cannot be rate limited. See *the classification steps* for more details on which client classes can be used to limit the packet rate.

Rate limits based on subnet are enforced only on the initially selected subnet for a given packet. If the selected subnet is subsequently changed, as may be the case for subnets in a shared network or when reselection is enabled in libraries such as the RADIUS hook, rate limits on the newly selected subnet are ignored. In other words, packets are gated only by the rate limit on the original subnet.

Note

It may seem logical to think that assigning a rate limit of *n* packets per time unit results in *n* DORA or *n* SARR exchanges. However, by default, all inbound packets are counted - meaning that a full message exchange accounts for two packets. To achieve the effect of counting an exchange only once, use client-class rate-limiting with a test expression that binds `pkt4.msgtype` to DHCPDISCOVER messages or `pkt6.msgtype` to SOLICIT messages.

17.19 libdhcp_mysql.so: Database Backend for MySQL

This hook library implements MySQL database storage for:

- Lease Backend
- Host Backend
- Configuration Backend

17.19.1 MySQL Configuration Backend

This hook library works in conjunction with *libdhcp_cb_cmds.so* to implement the API to create, read, update, and delete (CRUD) the configuration in a MySQL database. Please see *libdhcp_cb_cmds.so: Configuration Backend Commands* for more details.

Note

libdhcp_mysql.so is part of the open source code and is available to every Kea user. However, this hook library only works with *libdhcp_cb_cmds.so*, which allows easy configuration management with the use of API commands. *libdhcp_cb_cmds.so* is available only to ISC customers with a paid support contract; for more information on subscription options, please complete the form at <https://www.isc.org/contact>.

17.20 libdhcp_perfmon.so: Performance Monitoring

The Performance Monitoring (perfmon) hook library gives the Kea DHCPv4 and DHCPv6 server the ability to track and report performance-related data.

Note

libdhcp_perfmon.so is part of the open source code and is available to every Kea user.

17.20.1 Overview

The perfmon library, added in Kea 2.5.6, can be loaded by the *kea-dhcp4* or *kea-dhcp6* daemon by adding it to the `hooks-libraries` element of the server's configuration:

```
{
  "hooks-libraries": [
    {
      "library": "/usr/local/lib/libdhcp_perfmon.so",
      "parameters": {
        ...
      }
    },
    ...
  ],
  ...
}
```

The library tracks the lifecycle of client query packets as they are processed by Kea, from when the query was received by the kernel to when the response is ready to be sent. This tracking is driven by a packet event stack, which contains a list of event/timestamp pairs for significant milestones that occur during the processing of a client query. The list of possible events is shown below:

1. `socket_received` - kernel placed the client query into the socket buffer
2. `buffer_read` - server read the client query from the socket buffer
3. `mt_queued` - server placed the client query into its thread-pool work queue (multi-threaded mode only)
4. `process_started` - server has begun processing the query
5. `process_completed` - server has constructed the response and is ready to send it

This list is likely to expand over time; it will also be possible for hook developers to add their own events. This will be detailed in future releases.

17.20.2 Passive Event Logging

When the perfmon hook library is loaded, it logs the packet event stack contents for each client query that generates a response packet. The log entry contains client query identifiers, followed by the list of event/timestamp pairs in the order that they occurred.

For *kea-dhcp4*, the log is identified by the label `PERFMON_DHCP4_PKT_EVENTS`, and emitted at logger debug level 50 or higher. For a DHCPDISCOVER it is emitted once the DHCPOFFER is ready to send, and looks similar to the following (see the second entry):


```

2024-03-20 10:52:20.069 INFO [kea-dhcp4.leases/50033.140261252249344] DHCP4_LEASE_OFFER_
↳ [hwtype=1 08:00:27:25:d3:f4], cid=[no info], tid=0xc288f9: lease 178.16.2.0 will be_
↳ offered
2024-03-20 10:52:20.070 DEBUG [kea-dhcp4.perfmon-hooks/50033.140261252249344] PERFMON_
↳ DHCP4_PKT_EVENTS query: [hwtype=1 08:00:27:25:d3:f4], cid=[no info], tid=0xc288f9_
↳ events=[2024-Mar-20 14:52:20.067563 : socket_received, 2024-Mar-20 14:52:20.067810 :_
↳ buffer_read, 2024-Mar-20 14:52:20.067897 : mt_queued, 2024-Mar-20 14:52:20.067952 :_
↳ process_started, 2024-Mar-20 14:52:20.069614 : process_completed]

```

For *kea-dhcp6*, the log is identified by the label `PERFMON_DHCP6_PKT_EVENTS`, and emitted at logger debug level 50 or higher. For a `DHCPV6_SOLICIT` it is emitted once the `DHCPV6_ADVERTISE` is ready to send, and looks similar to the following (see the second entry):

```

2024-03-20 10:22:14.030 INFO [kea-dhcp6.leases/47195.139913679886272] DHCP6_LEASE_
↳ ADVERT duid=[00:03:00:01:08:00:27:25:d3:f4], [no hwaddr info], tid=0xb54806: lease for_
↳ address 3002:: and iaaid=11189196 will be advertised
2024-03-20 10:22:14.031 DEBUG [kea-dhcp6.perfmon-hooks/47195.139913679886272] PERFMON_
↳ DHCP6_PKT_EVENTS query: duid=[00:03:00:01:08:00:27:25:d3:f4], [no hwaddr info],_
↳ tid=0xb54806 events=[2024-Mar-20 14:22:14.028729 : socket_received, 2024-Mar-20_
↳ 14:22:14.028924 : buffer_read, 2024-Mar-20 14:22:14.029005 : process_started, 2024-Mar-
↳ 20 14:22:14.030566 : process_completed]

```

17.20.3 Duration Monitoring

When monitoring is enabled, stack event data is aggregated over a specified interval. These aggregates are referred to as "monitored durations" or simply "durations" for ease. Durations are uniquely identified by a "duration key" which consists of the following values:

- query message type - message type of the client query (e.g. `DHCPDISCOVER`, `DHCPV6_REQUEST`)
- response message type - message type of the server response (e.g. `DHCPOFFER`, `DHCPV6_REPLY`)
- start event label - event that defines the beginning of the task (e.g. `socket_received`, `process_started`)
- stop event label - event that defines the end of the task (e.g. `buffer_read`, `process_completed`)
- subnet id - subnet selected during message processing (or 0 for global durations)

Once the server has finished constructing a response to a query, the query's event stack is processed into a series of updates to monitored durations. If, upon updating, a duration's sample interval is found to have been completed, it is sent to reporting and a new sample interval is begun. The interval width is dictated by configuration parameter `interval-width-secs`.

The event stack for the multi-threaded mode `DHCPDISCOVER/DHCPOFFER` cycle shown above contains the following events:

Event Timestamp	Event Label
2024-Mar-20 14:52:20.067563	socket_received
2024-Mar-20 14:52:20.067810	buffer_read
2024-Mar-20 14:52:20.067897	mt_queued
2024-Mar-20 14:52:20.067952	process_started
2024-Mar-20 14:52:20.069614	process_completed

Assuming the selected subnet's ID was 100, the duration updates created by perfmon from these events are shown below:

Duration Keys for SubnetID 100	Update in microseconds
DHCPDISCOVER.DHCPOFFER.socket_received-buffer_read.100	247
DHCPDISCOVER.DHCPOFFER.buffer_read-mt_queue.100	87
DHCPDISCOVER.DHCPOFFER.mt_queued-process_started.100	55
DHCPDISCOVER.DHCPOFFER.process_started-process_completed.100	1662
DHCPDISCOVER.DHCPOFFER.composite-total_response.100	2051

Notice that in addition to the adjacent event updates, there is an additional duration update for the total duration of the entire stack whose key contains the event-pair `composite-total_response`. This tracks the total time to respond from when the query is received until the response is ready to send. Finally, there are also global duration updates for each of the above:

Global Duration Keys	Update in microseconds
DHCPDISCOVER.DHCPOFFER.socket_received-buffer_read.0	247
DHCPDISCOVER.DHCPOFFER.buffer_read-mt_queue.0	87
DHCPDISCOVER.DHCPOFFER.mt_queued-process_started.0	55
DHCPDISCOVER.DHCPOFFER.process_started-process_completed.0	1662
DHCPDISCOVER.DHCPOFFER.composite-total_response.0	2051

17.20.4 Statistics Reporting

When enabled, `perfmon` reports a duration's data each time the duration completes a sampling interval (see `stats-mgr-reporting` under *Configuration*). Each statistic employs the following naming convention:

```
{subnet-id[x]}.perfmon.<query type>-<response type>.<start event>-<end event>.<value-name>
```

There is both a global and a subnet-specific value for each. Currently, the only value reported for a given duration key is `mean-usecs`; this statistic is the mean time between the duration's event pair over the most recently completed interval. In other words, if during a given interval there were seven occurrences (i.e. updates) totaling 3500us, the `mean-usecs` reported would be 500us. Continuing with the example above, the statistics reported are named as follows for the subnet-level values:

```
subnet[100].perfmon.DHCPDISCOVER.DHCPOFFER.socket_received-buffer_read.mean-usecs
subnet[100].perfmon.DHCPDISCOVER.DHCPOFFER.buffer_read-mt_queue.mean-usecs
subnet[100].perfmon.DHCPDISCOVER.DHCPOFFER.mt_queued-process_started.mean-usecs
subnet[100].perfmon.DHCPDISCOVER.DHCPOFFER.process_started-process_completed.mean-usecs
subnet[100].perfmon.DHCPDISCOVER.DHCPOFFER.composite-total_response.mean-usecs
```

and as shown for global values:

```
perfmon.DHCPDISCOVER.DHCPOFFER.socket_received-buffer_read.mean-usecs
perfmon.DHCPDISCOVER.DHCPOFFER.buffer_read-mt_queue.mean-usecs
perfmon.DHCPDISCOVER.DHCPOFFER.mt_queued-process_started.mean-usecs
perfmon.DHCPDISCOVER.DHCPOFFER.process_started-process_completed.mean-usecs
perfmon.DHCPDISCOVER.DHCPOFFER.composite-total_response.mean-usecs
```

The results are reported to `StatsMgr`, an internal Kea component that reports data as statistics that can be retrieved using statistics commands. They can be fetched using the commands `statistic-get-all` or `statistic-get`.

17.20.5 Alarms

Alarms may be defined to watch specific durations. Each alarm defines a high-water mark, `high-water-ms`, and a low-water mark, `low-water-ms`. If the reported mean value for the duration exceeds the high-water mark a WARN level alarm log is emitted, at which point the alarm is considered "triggered." Once triggered, the WARN level log is repeated at the alarm report interval specified by `alarm-report-secs`, as long as the reported mean for the duration remains above the low-water mark. Once the mean falls below the low-water mark the alarm is cleared and an INFO level log is emitted.

The alarm-triggered WARN log looks similar to the following:

```
2024-03-20 10:22:14.030 WARN [kea-dhcp6.leases/47195.139913679886272] PERFMON_ALARM_
↳TRIGGERED Alarm for DHCPDISCOVER.DHCPOFFER.composite-total_response.0 has been_
↳triggered since 2024-03-20 10:18:20.070000, reported mean duration 00:00:00.700000_
↳exceeds high-water-ms: 500
```

The alarm-cleared INFO log looks like this:

```
2024-03-20 10:30:14.030 INFO [kea-dhcp6.leases/47195.139913679886272] PERFMON_ALARM_
↳CLEARED Alarm for DHCPDISCOVER.DHCPOFFER.composite-total_response.0 has been cleared,_
↳reported mean duration 00:00:00.010000 is now below low-water-ms: 25
```

17.20.6 API Commands

17.20.6.1 The `perfmon-control` Command

This command can be used to enable or disable active monitoring and statistics reporting at runtime without altering or reloading configuration.

```
{
  "command": "perfmon-control",
  "arguments": {
    "enable-monitoring": true,
    "stats-mgr-reporting": false
  }
}
```

Regardless of the arguments (if any) supplied, the resulting values of both flags are always returned:

```
{
  "result": 0,
  "text": "perfmon-control success",
  "arguments": {
    "enable-monitoring": true,
    "stats-mgr-reporting": false
  }
}
```

17.20.6.2 The `perfmon-get-all-durations` Command

This command fetches all monitored duration data currently held in memory by the Perfmon hook library.

```
{
  "command": "perfmon-get-all-durations",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "result-set-format": true
  }
}

```

A result of 0 is returned if command succeeds along with monitored duration data, while a result of 1 is returned if command is invalid or command processing encounters an error.

The format of the monitored duration data returned is determined by the optional argument, `result-set-format`. When false, (the default), the list of durations will be returned as a list of individual elements as shown below:

```

{
  "arguments": {
    "durations": [{
      "duration-key": {
        "query-type": "DHCPDISCOVER",
        "response-type": "DHCPOFFER",
        "start-event": "buffer_read",
        "stop-event": "mt_queued",
        "subnet-id": 0
      },
      "max-duration-usecs": 118,
      "min-duration-usecs": 31,
      "occurrences": 501,
      "start-time": "2024-06-12 17:52:06.814884",
      "total-duration-usecs": 23951,
      "mean-duration-usecs": 47
    },
    ...
  ],
  "result-set-format": false,
  "interval-width-secs": 5,
  "timestamp": "2024-06-12 17:52:22.397233"
},
"result": 0,
"text": "perfmon-get-all-durations: n found"
}

```

When `result-set-format` is true, the list of durations will be returned in a format similar to an SQL result set as follows:

```

{
  "arguments": {
    "durations-result-set": {
      "columns": [
        "query-type",
        "response-type",
        "start-event",
        "end-event",
        "subnet-id",
        "interval-start",
        "occurrence",
        "min-duration-usecs",
        "max-duration-usecs",

```

(continues on next page)

(continued from previous page)

```

        "total-duration-usecs",
        "mean-duration-usecs"
    ],
    "rows": [ [
        "DHCPDISCOVER",
        "DHCPOFFER",
        "buffer_read",
        "mt_queued",
        0,
        "2024-06-12 17:52:06.814884",
        501,
        31,
        118,
        23951,
        47
    ],
        ...
    ]
},
"result-set-format": true,
"interval-width-secs": 5,
"timestamp": "2024-06-12 17:52:22.397233"
},
"result": 0,
"text": "perfmon-get-all-durations: n found"
}

```

The data values for each duration will be from the duration's last completed data interval. If a duration has no such interval, `interval-start` will be reported as "`<none>`" and the remaining values will be zero.

17.20.7 Configuration

Use of the performance monitoring hook library is configured via several parameters:

- **enable-monitoring**
This parameter enables event data aggregation for reporting, statistics, and alarms. It defaults to `false`.
- **interval-width-secs**
This specifies the amount of time, in seconds, that individual task durations are accumulated into an aggregate before they are reported. The default is 60 seconds.
- **stats-mgr-reporting**
This enables the reporting of aggregates to StatsMgr. It defaults to `true`.
- **alarm-report-secs**
This specifies the amount of time, in seconds, between logging instances for an alarm once it has been triggered. It defaults to 300 seconds.
- **alarms**
This is an optional list of alarms that monitor specific duration aggregates. Each alarm is defined by these four parameters:
 - **duration-key**
Identifies the monitored duration for the following:
 - * `query-type` - message type of the client query (e.g. `DHCPDISCOVER`, `DHCPV6_REQUEST`)

- * response-type - message type of the server response (e.g. DHCPOFFER, DHCPV6_REPLY)
- * start-event - event that defines the beginning of the task (e.g. socket_received, process_started)
- * stop-event - event that defines the end of the task
- * subnet-id - subnet selected during message processing (or 0 for global durations)

– **enable-alarm**

Enables or disables this alarm; defaults to `true`.

– **high-water-ms**

Specifies the value, in milliseconds, over which the duration must occur to trigger this alarm; must be greater than zero.

– **low-water-ms**

Specifies the value, in milliseconds, under which the duration must fall to clear this alarm; must be greater than zero but less than `high-water-ms`.

Note

Passive event logging is always enabled, even without specifying the "parameters" section.

A sample configuration is shown below:

```
{
  "hooks-libraries": [
    {
      "library": "lib/kea/hooks/libdhcp_perfmon.so",
      "parameters": {
        "enable-monitoring": true,
        "interval-width-secs": 5,
        "stats-mgr-reporting": true,
        "alarm-report-secs": 600,
        "alarms": [
          {
            "duration-key": {
              "query-type": "DHCPDISCOVER",
              "response-type": "DHCPOFFER",
              "start-event": "process-started",
              "stop-event": "process-completed",
              "subnet-id": 0
            },
            "enable-alarm": true,
            "high-water-ms": 500,
            "low-water-ms": 25
          }
        ]
      }
    }
  ]
}
```

17.21 libdhcp_ping_check.so: Ping Check

This hook library adds the ability to perform a "ping check" of a candidate IPv4 address prior to offering it to a DHCP client. This feature is similar to a behavior available in ISC DHCP and one suggested in [RFC 2131](#), section 3.1, item 2.

Note

libdhcp_ping_check.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

17.21.1 Overview

The Ping Check library, added in Kea 2.5.4, can be loaded by the *kea-dhcp4* daemon by adding it to the `hooks-libraries` element of the server's configuration:

```
{
  "hooks-libraries": [
    {
      "library": "/usr/local/lib/libdhcp_ping_check.so",
      "parameters": {
        ...
      }
    },
    ...
  ],
  ...
}
```

When the library is loaded, *kea-dhcp4* conducts a ping check prior to offering a lease to a client if all of the following conditions are true:

1. The ping check hook library is loaded.
2. Ping checking is enabled.
3. The server is responding to a DHCPDISCOVER.
4. The candidate lease is neither active nor reserved.
5. Any of the following are true:
 - a. This is the first offer of this lease to this client. This check can only be done if `offer-lifetime` is greater than zero (i.e. temporary allocation on DHCPDISCOVER is enabled). If `offer-lifetime` is zero ping checks are done for every DHCPOFFER, as the server has no way to know whether prior offers were made.
 - b. The lease is being offered to a client other than its previous owner.
 - c. The lease is being offered to its previous owner and more than the number of seconds specified by `ping-cltt-secs` have elapsed since the CLTT of the original lease.

When the ping check library is loaded, the *kea-dhcp4* server responds to a DHCPDISCOVER by:

1. Selecting a candidate IPv4 address through normal processes and using it to construct a DHCPOFFER.
2. Parking the DHCPOFFER and requesting a ping check from the ping check hook library, via its `lease4_offer` callout. Refer to *Parked-Packet Limit* for a basic introduction to packet parking.

3. Testing the callout conditions described above: if they are not satisfied, it returns without conducting a check, and the server sends the DHCP OFFER to the client. Otherwise, the callout initiates a ping check for the lease address.

4. Upon conclusion of the ping check, the server either sends the DHCP OFFER to the client, if the check concluded that the address is available, or discards the DHCP OFFER and creates a DECLINED lease for the address.

Each ping check consists of the following steps:

1. If the number of ECHO REPLYs sent is less than the configured minimum number to send, send an ICMP ECHO REQUEST to the lease address. Otherwise, conclude that the address is available.
2. If no ECHO REPLY is received within a configurable amount of time, return to step 1.
3. Upon receipt of an ICMP ECHO REPLY, conclude that the lease is NOT available.
4. If any of the following occur:
 - a. Receipt of an ICMP DESTINATION UNREACHABLE message
 - b. Send failure of an ICMP ECHO REQUEST due to a network error (e.g. network is unreachable)
 - c. Send failure of an ICMP ECHO REQUEST due to a permissions error (e.g. lease address is a broadcast address)
 - d. Send failure of an ICMP ECHO REQUEST with socket buffer full error
 skip the address check and treat it as available.

Note

"Socket buffer full" errors indicate that the OS rate limits on ICMP are being exceeded. The server does not retry them, as this would likely only exacerbate the situation. If this occurs continuously, then the client load on the server may be too high to accommodate ping checking. Ping checking is not recommended for systems with high throughput demands.

17.21.2 Configuration

The ping check hook library currently supports the following configuration parameters that may be set at the global and subnet levels; subnet values override global values:

- **enable-ping-check** - This parameter enables or disables ping checking at a given scope.
- **min-ping-requests** - This is the minimum number of ECHO REQUESTs sent without receiving a reply, before an address is declared available. The default is 1; it must be greater than zero.
- **reply-timeout** - This is the maximum amount of time to wait for a reply to a single ECHO REQUEST. Specified in milliseconds, it must be greater than zero; the default is 100.
- **ping-cltt-secs** - This specifies the number of seconds that must elapse after the lease's CLTT before a ping check is conducted, when the client is the lease's previous owner. The default value is 60 seconds.

The following parameter is only supported at the global level:

- **ping-channel-threads** - In multi-threaded mode, this is the number of threads in the channel's thread pool. The default is 0, which instructs the library to use the same number of threads as the Kea core. This value is ignored if given when Kea is in single-threaded mode.

The following configuration excerpt illustrates a global-level configuration:

```
{
  "hooks-libraries": [{
```

(continues on next page)

(continued from previous page)

```

    "library": "lib/kea/hooks/libdhcp_ping_check.so",
    "parameters": {
        "enable-ping-check" : true,
        "min-ping-requests" : 1,
        "reply-timeout" : 100,
        "ping-cltt-secs" : 60,
        "ping-channel-threads" : 0
    }
}
}]
}

```

The following excerpt demonstrates a subnet-level configuration:

```

{
  "subnet4": [{
    "subnet": "192.0.2.0/24",
    "pools": [{
      "pool": "192.0.2.10 - 192.0.2.20"
    }],
    "user-context": {
      "enable-ping-check" : true,
      "min-ping-requests" : 2,
      "reply-timeout" : 250,
      "ping-cltt-secs" : 120
    }
  }
}
}]
}

```

17.22 libdhcp_pgsql.so: Database Backend for PostgreSQL

This hook library implements PostgreSQL database storage for:

- Lease Backend
- Host Backend
- Configuration Backend

17.22.1 PostgreSQL Configuration Backend

This hook library works in conjunction with *libdhcp_cb_cmds.so* to implement the API to create, read, update, and delete (CRUD) the configuration in a PostgreSQL database. Please see *libdhcp_cb_cmds.so: Configuration Backend Commands* for more details.

i Note

libdhcp_pgsql.so is part of the open source code and is available to every Kea user. However, this hook library only works with *libdhcp_cb_cmds.so*, which allows easy configuration management with the use of API commands. *libdhcp_cb_cmds.so* is available only to ISC customers with a paid support contract; for more information on subscription options, please complete the form at <https://www.isc.org/contact>.

17.23 `libdhcp_radius.so`: RADIUS Server Support

This hook library allows the Kea DHCP servers to use the RADIUS protocol to authorize DHCP clients through the access service, or for DHCP lease journaling through the accounting service. For details on RADIUS in Kea, please see *RADIUS*.

Note

`libdhcp_radius.so` is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

17.24 `libdhcp_rbac.so`: Role-Based Access Control

17.24.1 Role-Based Access Control (RBAC) Overview

Before the processing of commands in received HTTP requests, `libdhcp_rbac.so` takes specific parameters, e.g. the common-name part of the client certificate subject name, to assign a role to the request. The configuration associated with this role is used to accept or reject the command. After processing, the response can be rewritten, e.g. parts can be removed.

Here is a summary of the steps in processing a request:

- The HTTP library records some information to be used later, e.g. the remote address.
- When TLS is required but the request was not protected by TLS, the request is rejected by sending an "unauthorized" response.
- The command is extracted from the request.
- A role is assigned using recorded information in the request.
- The role is used to accept (pass through) or reject (send a forbidden response to) the command.

Here is a summary of the steps in processing a response:

- The information attached to the request is retrieved during the request processing (when the request was accepted).
- Request filters are applied to the response.

Note

`libdhcp_rbac.so` is available only to ISC customers with a paid support contract. For more information on subscription options, please complete the form at <https://www.isc.org/contact>.

Note

Since Kea 2.7.2 the RBAC is no longer limited to the Control Agent: it can be used by any other server supporting HTTP/HTTPS control sockets e.g. DHCPv4, DHCPv6 and DDNS servers.

17.24.2 Role-Based Access Control Configuration

17.24.2.1 Role Assignment

Role assignment is governed by the configured role-assignment method.

Table 6: Role assignment methods

Name	Description
remote-address	remote/client IP address
cert-subject	common-name part of the client certificate subject name
cert-issuer	common-name part of the client certificate issuer name
basic-authentication	user ID of basic HTTP authentication
custom-value	another role can be designed in external hooks

17.24.2.2 Role Configuration

Table 7: Role configuration parameters

Name	Description
name	the role name (with the exception of the default and unknown roles)
accept-commands	the accept access list
reject-commands	the reject access list
other-commands	specifies what to do for commands not matching accept and reject lists (default: reject)
list-match-first	specifies what to do for commands matching both the accept and reject list by giving the list to check and apply first (default: accept)
response-filters	the filters to apply to responses

Note

The role assignment may fail with `cert-subject` when the client certificate was not required, or it may have no subject common name and instead have a DNS alternative subject name. In this case, the role assignment returns the empty role and the `default-role` entry is used.

The role assignment can return an unexpected value, e.g. with an unregistered role name or a typing error. In this case the `unknown-role` entry is used.

The default for both `default-role` and `unknown-role` is to reject all commands.

17.24.2.3 API Commands

All commands of the REST API are described in files in the source directory `src/share/api`, or in installed Kea in `./share/kea/api`. `libdhcp_rbac.so` reads these files to take the name, the access right (i.e. `read` or `write`), and the hook name. The access right can be modified in the file but changes are only applied after the server or agent restarts. Removing command definitions from `./share/kea/api` has consequences: if the access control list is based on `read` or `write` and the definition file is missing, the server or agent always rejects such a command. If the access controls list is using `commands` to specify the name of a command and the definition file from `./share/kea/api` of this particular command is missing, the Control Agent logs an error on startup and exit.

Table 8: Extra command-definition parameters

Name	Description
name	(mandatory) the command name
access	(mandatory) the access right i.e. <code>read</code> or <code>write</code>
hook	(optional) the hook name (empty or not-present for commands of servers or agents)

Note

These command description files are security-sensitive, e.g. with too-permissive access rights a local attacker may modify them and defeat the RBAC goal.

17.24.2.4 Access Control Lists

Access control lists can be specified using a name (string) or a single entry map.

Table 9: Predefined named access list

Name	Description
ALL	matches everything
NONE	matches nothing
READ	matches commands with the read-access right
WRITE	matches commands with the write-access right

Map access list specifications use a list type in the name of the single entry and parameter in the value.

Table 10: Access list types

Name	Description	Parameter
not	logical not	access list
and	logical and	list of access lists
or	logical or	list of access lists
command	explicit list	list of command names
access	by access right	access right (<code>read</code> or <code>write</code>)
hook	by hook	hook name (can be empty)

17.24.2.5 Response Filters

Table 11: Predefined response filters

Name	Description
list-commands	Removes not-allowed commands from the list-commands response

17.24.2.6 Global Parameters

The global parameters are:

- `assign-role-method`: the name of the method which is used for role assignment. This parameter is mandatory.
- `api-files`: the path of the directory where the API files describing commands can be found. This parameter is mandatory.

- `require-tls`: the specification of whether received requests on HTTP (vs HTTPS) are rejected. It defaults to false when the role-assignment method is not based on certificates.
- `commands`: the list of extra command configurations.
- `access-control-lists`: the named access control list definitions (each definition is a single entry map; the name of the entry is the name of the access list, and the value is the specification). The name is used in other parts of the configuration, such as "accept-commands".
- `roles`: the role configurations.
- `default-role`: the configuration of the default role (used when "" is assigned).
- `unknown-role`: the configuration of the unknown role (used when the not-empty assigned role has no configuration).

17.24.3 Sample Configuration

A sample configuration is available in `doc/examples/agent/rbac.json` in the Kea source and is copied below.

```

1 {
2   "Control-agent": {
3     // We need to specify where the agent should listen to incoming HTTP
4     // queries.
5     "http-host": "127.0.0.1",
6
7     // If enabling HA and multi-threading, the 8000 port is used by the HA
8     // hook library http listener. When using HA hook library with
9     // multi-threading to function, make sure the port used by dedicated
10    // listener is different (e.g. 8001) than the one used by CA. Note
11    // the commands should still be sent via CA. The dedicated listener
12    // is specifically for HA updates only.
13    "http-port": 8000,
14
15    // TLS trust anchor (Certificate Authority). This is a file name or
16    // (for OpenSSL only) a directory path.
17    "trust-anchor": "my-ca",
18
19    // TLS server certificate file name.
20    "cert-file": "my-cert",
21
22    // TLS server private key file name.
23    "key-file": "my-key",
24
25    // TLS require client certificates flag. Default is true and means
26    // require client certificates. False means they are optional.
27    "cert-required": true,
28
29    // Add hooks here.
30    "hooks-libraries": [
31      {
32        "library": "/opt/lib/libdhcp_rbac.so",
33        "parameters": {
34          // This section configures the RBAC hook library.
35          // Mandatory parameters.
36          "assign-role-method": "cert-subject",

```

(continues on next page)

(continued from previous page)

```

37     "api-files": "/opt/share/kea/api",
38     // Optional parameters.
39     "require-tls": true,
40     "commands": [
41     {
42         "name": "my-command",
43         "access": "read",
44         "hook": "my-hook"
45     } ],
46     "access-control-lists": [
47     {
48         "my-none": { "not": "ALL" }
49     }, {
50         "another-none": { "and": [ "ALL", "NONE" ] }
51     }, {
52         "my-read": { "access": "read" }
53     } ],
54     "roles": [
55     {
56         "name": "kea-client",
57         "accept-commands":
58         {
59             "commands": [ "list-commands", "status-get" ]
60         },
61         "reject-commands": "NONE",
62         "other-commands": "reject",
63         "list-match-first": "accept",
64         "response-filters": [ "list-commands" ]
65     }, {
66         "name": "admin",
67         "accept-commands": "ALL",
68         "reject-commands":
69         {
70             "hook": "cb_cmds"
71         },
72         "list-match-first": "reject"
73     } ],
74     "default-role":
75     {
76         "accept-commands": "NONE",
77         "reject-commands": "ALL"
78     },
79     "unknown-role":
80     {
81         "accept-commands": "READ",
82         "reject-commands": "WRITE"
83     }
84 } ]

```

```

86 // Additional parameters, such as logging and others
87 // omitted for clarity.
88

```

(continues on next page)

(continued from previous page)

```

89 }
90 }
91 }

```

17.24.4 Accept/Reject Algorithm

This is the pseudo-code of the accept/reject decision algorithm which returns true (accept) or false (reject).

```

bool match(command) {
    if (list-match-first == accept) {
        if (accept_list && accept_list->match(command)) {
            return (true);
        }
        if (reject_list && reject_list->match(command)) {
            return (false);
        }
    } else {
        if (reject_list && reject_list->match(command)) {
            return (false);
        }
        if (accept_list && accept_list->match(command)) {
            return (true);
        }
    }
    if (others == reject) {
        return (false);
    } else {
        return (true);
    }
}

```

17.24.5 Custom Hook Commands and Command Redefinition

It is possible to have a custom hook with new commands. In this case, Role Based Access Control can be used to manage a new command in two ways.

The `command` global parameter can be used to define its name, access type, and hook name:

```

{
  "commands": [
    {
      "name": "my-new-command",
      "access": "write",
      "hook": "my-custom-hook"
    }
  ],
  ...
}

```

The new command can then be specified in roles:

```
{
  "roles": [
    {
      "name": "user1",
      "accept-commands": {
        "commands": [ "my-new-command" ] },
      "reject-commands": "WRITE",
      "list-match-first": "accept"
    },
    {
      "name": "user2",
      "accept-commands": { "hook": "my-custom-hook" },
      "reject-commands": "ALL",
      "list-match-first": "accept"
    }
  ],
  ...
}
```

The second method is to create a custom file in `.../share/kea/api` and define the access type of the custom command(s).

It is also possible to redefine an existing command by removing its definition file from `.../share/kea/api` and defining it in the `commands` global parameter:

```
{
  "commands": [
    {
      "name": "dhcp-disable",
      "access": "read",
      "hook": "my-custom-hook-3"
    }
  ]
}
```

With this approach, an administrator can put the configurations of all existing commands inside the Control Agent's configuration file.

17.24.6 Extensive Example

Here is an extensive example for a role accepting all read commands, with the exception of `config-get`, e.g. for hiding passwords. For any remote user who is not recognized as "user1", all commands should be rejected.

The first option is to put the allowed commands in the "accept-commands" list and to reject anything else:

```
{
  "roles": [
    {
      "name": "user1",
      "accept-commands":
      {
        "and": [
          "READ",
          { "not":
```

(continues on next page)

(continued from previous page)

```

        { "commands": [ "config-get" ] }
      }
    ],
    "reject-commands": "ALL",
    // This is the default but as the config relies on it
    // it is explicitly set.
    "list-match-first": "accept"
  },
  ...
],
...
}

```

A common alternative is not to set the "reject-commands" list, i.e. leave it empty and rely on "other-commands" to reject anything else.

```

{
  "roles": [
    {
      "name": "user2",
      "accept-commands":
      {
        "and": [
          "READ",
          { "not":
            { "commands": [ "config-get" ] }
          }
        ]
      },
      // This is the default but as the config relies on it
      // it is explicitly set.
      "other-commands": "reject"
    },
    ...
  ],
  ...
}

```

It is also possible to do the opposite, i.e. to set only the "reject-commands" list:

```

{
  "roles": [
    {
      "name": "user3",
      "reject-commands":
      {
        "or": [
          "WRITE",
          { "commands": [ "config-get" ] }
        ]
      },
    },
  ],
}

```

(continues on next page)

(continued from previous page)

```

    "other-commands": "accept"
  },
  ...
],
...
}

```

Or use both lists with the exception in the "reject-commands" list, which must be checked first as "config-get" has the read-access right.

```

{
  "roles": [
    {
      "name": "user4",
      "accept-commands": "READ",
      "reject-commands": { "commands": [ "config-get" ] },
      "list-match-first": "reject"
    },
    ...
  ],
  ...
}

```

To check any configuration, it is a good idea to use the "list-commands" response filter, which shows errors such as missing (rejected) commands and extra (accepted) commands.

access-control-lists can be used for definitions of access control lists and later reused in roles:

```

{
  "access-control-lists": [
    {
      "my-list-one": {
        "or": [
          {
            "hook": "subnet_cmds"
          },
          {
            "commands": [ "list-commands" ]
          }
        ]
      },
      {
        "my-list-two": {
          "and": [
            "READ",
            {
              "not": {
                "commands": [ "config-get" ]
              }
            }
          ]
        }
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    },
    {
        "my-list-three":{
            "or":[
                { "hook":"subnet_cmds" },
                { "hook":"class_cmds" },
                { "hook":"lease_cmds" }
            ]
        }
    }
],
"roles":[
    {
        "name":"admin",
        "accept-commands":"my-list-one",
        "reject-commands":"ALL",
        "list-match-first":"accept"
    },
    {
        "name":"admin2",
        "accept-commands":"my-list-two",
        "reject-commands":"ALL",
        "list-match-first":"accept"
    }
],
"unknown-role":{
    "accept-commands":"my-list-three",
    "reject-commands":"ALL"
},
...
}

```

17.25 `libdhcp_run_script.so`: Run Script Support for External Hook Scripts

The Run Script hook library adds support for calling an external script for specific packet-processing hook points.

Note

`libdhcp_run_script.so` is part of the open source code and is available to every Kea user.

Note

This library can only be loaded by the `kea-dhcp4` or `kea-dhcp6` process.

```

{
    "hooks-libraries": [
        {

```

(continues on next page)

(continued from previous page)

```

    "library": "/usr/local/lib/libdhcp_run_script.so",
    "parameters": {
        "name": "/full_path_to/script_name.sh",
        "sync": false
    }
}
]
}

```

The parameters contain the name, which indicates the full path to the external script to be called on each hook point, and also the `sync` option, to be able to wait synchronously for the script to finish execution. If the `sync` parameter is `false`, then the script will launch and Kea will not wait for the execution to finish, causing all the OUT parameters of the script (including the next step) to be ignored.

The external script must be stored in the "[kea-install-dir]/share/kea/scripts/" directory which can be overridden at startup by setting the environment variable `KEA_HOOK_SCRIPTS_PATH` to a different path.

Note

The script inherits all privileges from the server which calls it.

Note

Currently, enabling synchronous calls to external scripts is not supported.

This library has several hook-point functions implemented, which are called at the specific packet-processing stage.

The dhcpv4 hook points:

```

lease4_renew
lease4_expire
lease4_recover
leases4_committed
lease4_release
lease4_decline

```

The dhcpv6 hook points:

```

lease6_renew
lease6_rebind
lease6_expire
lease6_recover
leases6_committed
lease6_release
lease6_decline
addr6_register

```

Each hook point extracts the Kea internal data and exports it as string environment variables. These parameters are shared with the target script using the child process environment. The only parameter passed to the call of the target script is the name of the hook point.

An example of a script implementing all hook points is presented below:

```
#!/bin/sh

unknown_handle() {
    echo "Unhandled function call ${*}"
    exit 123
}

lease4_renew () {
    ...
}

lease4_expire () {
    ...
}

lease4_recover () {
    ...
}

leases4_committed () {
    ...
}

lease4_release () {
    ...
}

lease4_decline () {
    ...
}

lease6_renew () {
    ...
}

lease6_rebind () {
    ...
}

lease6_expire () {
    ...
}

lease6_recover () {
    ...
}

leases6_committed () {
    ...
}

lease6_release () {
```

(continues on next page)

(continued from previous page)

```
    ...
}

lease6_decline () {
    ...
}

addr6_register () {
    ...
}

case "$1" in
    "lease4_renew")
        lease4_renew
        ;;
    "lease4_expire")
        lease4_expire
        ;;
    "lease4_recover")
        lease4_recover
        ;;
    "leases4_committed")
        leases4_committed
        ;;
    "lease4_release")
        lease4_release
        ;;
    "lease4_decline")
        lease4_decline
        ;;
    "lease6_renew")
        lease6_renew
        ;;
    "lease6_rebind")
        lease6_rebind
        ;;
    "lease6_expire")
        lease6_expire
        ;;
    "lease6_recover")
        lease6_recover
        ;;
    "leases6_committed")
        leases6_committed
        ;;
    "lease6_release")
        lease6_release
        ;;
    "lease6_decline")
        lease6_decline
        ;;
    "addr6_register")
```

(continues on next page)

(continued from previous page)

```

        addr6_register
        ;;
    *)
        unknown_handle "${@}"
        ;;
esac

```

Available parameters for each hook point are presented below.

DHCPv4:

lease4_renew

```

QUERY4_TYPE
QUERY4_TXID
QUERY4_LOCAL_ADDR
QUERY4_LOCAL_PORT
QUERY4_REMOTE_ADDR
QUERY4_REMOTE_PORT
QUERY4_IFACE_INDEX
QUERY4_IFACE_NAME
QUERY4_HOPS
QUERY4_SECS
QUERY4_FLAGS
QUERY4_CIADDR
QUERY4_SIADDR
QUERY4_YIADDR
QUERY4_GIADDR
QUERY4_RELAYED
QUERY4_HWADDR
QUERY4_HWADDR_TYPE
QUERY4_LOCAL_HWADDR
QUERY4_LOCAL_HWADDR_TYPE
QUERY4_REMOTE_HWADDR
QUERY4_REMOTE_HWADDR_TYPE
QUERY4_OPTION_82
QUERY4_OPTION_82_SUB_OPTION_1
QUERY4_OPTION_82_SUB_OPTION_2
SUBNET4_ID
SUBNET4_NAME
SUBNET4_PREFIX
SUBNET4_PREFIX_LEN
PKT4_CLIENT_ID
PKT4_HWADDR
PKT4_HWADDR_TYPE
LEASE4_ADDRESS
LEASE4_CLTT
LEASE4_HOSTNAME
LEASE4_HWADDR
LEASE4_HWADDR_TYPE
LEASE4_STATE
LEASE4_SUBNET_ID
LEASE4_VALID_LIFETIME

```

(continues on next page)

(continued from previous page)

LEASE4_CLIENT_ID

lease4_expire

```
LEASE4_ADDRESS
LEASE4_CLTT
LEASE4_HOSTNAME
LEASE4_HWADDR
LEASE4_HWADDR_TYPE
LEASE4_STATE
LEASE4_SUBNET_ID
LEASE4_VALID_LIFETIME
LEASE4_CLIENT_ID
REMOVE_LEASE
```

lease4_recover

```
LEASE4_ADDRESS
LEASE4_CLTT
LEASE4_HOSTNAME
LEASE4_HWADDR
LEASE4_HWADDR_TYPE
LEASE4_STATE
LEASE4_SUBNET_ID
LEASE4_VALID_LIFETIME
LEASE4_CLIENT_ID
```

leases4_committed

```
QUERY4_TYPE
QUERY4_TXID
QUERY4_LOCAL_ADDR
QUERY4_LOCAL_PORT
QUERY4_REMOTE_ADDR
QUERY4_REMOTE_PORT
QUERY4_IFACE_INDEX
QUERY4_IFACE_NAME
QUERY4_HOPS
QUERY4_SECS
QUERY4_FLAGS
QUERY4_CIADDR
QUERY4_SIADDR
QUERY4_YIADDR
QUERY4_GIADDR
QUERY4_RELAYED
QUERY4_HWADDR
QUERY4_HWADDR_TYPE
QUERY4_LOCAL_HWADDR
QUERY4_LOCAL_HWADDR_TYPE
QUERY4_REMOTE_HWADDR
QUERY4_REMOTE_HWADDR_TYPE
QUERY4_OPTION_82
QUERY4_OPTION_82_SUB_OPTION_1
```

(continues on next page)

(continued from previous page)

```

QUERY4_OPTION_82_SUB_OPTION_2
LEASES4_SIZE
DELETED_LEASES4_SIZE

```

If LEASES4_SIZE or DELETED_LEASES4_SIZE is non-zero, then each lease has its own unique identifier, as shown below. The first index starts at 0.

```

LEASES4_AT0_ADDRESS
LEASES4_AT0_CLTT
LEASES4_AT0_HOSTNAME
LEASES4_AT0_HWADDR
LEASES4_AT0_HWADDR_TYPE
LEASES4_AT0_STATE
LEASES4_AT0_SUBNET_ID
LEASES4_AT0_VALID_LIFETIME
LEASES4_AT0_CLIENT_ID
DELETED_LEASES4_AT0_ADDRESS
DELETED_LEASES4_AT0_CLTT
DELETED_LEASES4_AT0_HOSTNAME
DELETED_LEASES4_AT0_HWADDR
DELETED_LEASES4_AT0_HWADDR_TYPE
DELETED_LEASES4_AT0_STATE
DELETED_LEASES4_AT0_SUBNET_ID
DELETED_LEASES4_AT0_VALID_LIFETIME
DELETED_LEASES4_AT0_CLIENT_ID

```

lease4_release

```

QUERY4_TYPE
QUERY4_TXID
QUERY4_LOCAL_ADDR
QUERY4_LOCAL_PORT
QUERY4_REMOTE_ADDR
QUERY4_REMOTE_PORT
QUERY4_IFACE_INDEX
QUERY4_IFACE_NAME
QUERY4_HOPS
QUERY4_SECS
QUERY4_FLAGS
QUERY4_CIADDR
QUERY4_SIADDR
QUERY4_YIADDR
QUERY4_GIADDR
QUERY4_RELAYED
QUERY4_HWADDR
QUERY4_HWADDR_TYPE
QUERY4_LOCAL_HWADDR
QUERY4_LOCAL_HWADDR_TYPE
QUERY4_REMOTE_HWADDR
QUERY4_REMOTE_HWADDR_TYPE
QUERY4_OPTION_82
QUERY4_OPTION_82_SUB_OPTION_1

```

(continues on next page)

(continued from previous page)

```

QUERY4_OPTION_82_SUB_OPTION_2
LEASE4_ADDRESS
LEASE4_CLTT
LEASE4_HOSTNAME
LEASE4_HWADDR
LEASE4_HWADDR_TYPE
LEASE4_STATE
LEASE4_SUBNET_ID
LEASE4_VALID_LIFETIME
LEASE4_CLIENT_ID

```

lease4_decline

```

QUERY4_TYPE
QUERY4_TXID
QUERY4_LOCAL_ADDR
QUERY4_LOCAL_PORT
QUERY4_REMOTE_ADDR
QUERY4_REMOTE_PORT
QUERY4_IFACE_INDEX
QUERY4_IFACE_NAME
QUERY4_HOPS
QUERY4_SECS
QUERY4_FLAGS
QUERY4_CIADDR
QUERY4_SIADDR
QUERY4_YIADDR
QUERY4_GIADDR
QUERY4_RELAYED
QUERY4_HWADDR
QUERY4_HWADDR_TYPE
QUERY4_LOCAL_HWADDR
QUERY4_LOCAL_HWADDR_TYPE
QUERY4_REMOTE_HWADDR
QUERY4_REMOTE_HWADDR_TYPE
QUERY4_OPTION_82
QUERY4_OPTION_82_SUB_OPTION_1
QUERY4_OPTION_82_SUB_OPTION_2
LEASE4_ADDRESS
LEASE4_CLTT
LEASE4_HOSTNAME
LEASE4_HWADDR
LEASE4_HWADDR_TYPE
LEASE4_STATE
LEASE4_SUBNET_ID
LEASE4_VALID_LIFETIME
LEASE4_CLIENT_ID

```

DHCPv6:

lease6_renew

```

QUERY6_TYPE

```

(continues on next page)

(continued from previous page)

```
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
QUERY6_REMOTE_HWADDR_TYPE
QUERY6_PROTO
QUERY6_CLIENT_ID
LEASE6_ADDRESS
LEASE6_CLTT
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
PKT6_IA_IAID
PKT6_IA_IA_TYPE
PKT6_IA_IA_T1
PKT6_IA_IA_T2
```

lease6_rebind

```
QUERY6_TYPE
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
QUERY6_REMOTE_HWADDR_TYPE
QUERY6_PROTO
QUERY6_CLIENT_ID
LEASE6_ADDRESS
LEASE6_CLTT
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
```

(continues on next page)

(continued from previous page)

```
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
PKT6_IA_IAID
PKT6_IA_IA_TYPE
PKT6_IA_IA_T1
PKT6_IA_IA_T2
```

lease6_expire

```
LEASE6_ADDRESS
LEASE6_CLTT
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
REMOVE_LEASE
```

lease6_recover

```
LEASE6_ADDRESS
LEASE6_CLTT
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
```

leases6_committed

```
QUERY6_TYPE
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
```

(continues on next page)

(continued from previous page)

```

QUERY6_REMOTE_HWADDR_TYPE
QUERY6_PROTO
QUERY6_CLIENT_ID
LEASES6_SIZE
DELETED_LEASES6_SIZE

```

If LEASES6_SIZE or DELETED_LEASES6_SIZE is non-zero, then each lease has its own unique identifier, as shown below. The first index starts at 0.

```

LEASES6_AT0_ADDRESS
LEASES6_AT0_CLTT
LEASES6_AT0_HOSTNAME
LEASES6_AT0_HWADDR
LEASES6_AT0_HWADDR_TYPE
LEASES6_AT0_STATE
LEASES6_AT0_SUBNET_ID
LEASES6_AT0_VALID_LIFETIME
LEASES6_AT0_DUID
LEASES6_AT0_IAID
LEASES6_AT0_PREFERRED_LIFETIME
LEASES6_AT0_PREFIX_LEN
LEASES6_AT0_TYPE
DELETED_LEASES6_AT0_ADDRESS
DELETED_LEASES6_AT0_CLTT
DELETED_LEASES6_AT0_HOSTNAME
DELETED_LEASES6_AT0_HWADDR
DELETED_LEASES6_AT0_HWADDR_TYPE
DELETED_LEASES6_AT0_STATE
DELETED_LEASES6_AT0_SUBNET_ID
DELETED_LEASES6_AT0_VALID_LIFETIME
DELETED_LEASES6_AT0_DUID
DELETED_LEASES6_AT0_IAID
DELETED_LEASES6_AT0_PREFERRED_LIFETIME
DELETED_LEASES6_AT0_PREFIX_LEN
DELETED_LEASES6_AT0_TYPE

```

lease6_release

```

QUERY6_TYPE
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
QUERY6_REMOTE_HWADDR_TYPE
QUERY6_PROTO
QUERY6_CLIENT_ID
LEASE6_ADDRESS
LEASE6_CLTT

```

(continues on next page)

(continued from previous page)

```
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
```

lease6_decline

```
QUERY6_TYPE
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
QUERY6_REMOTE_HWADDR_TYPE
QUERY6_PROTO
QUERY6_CLIENT_ID
LEASE6_ADDRESS
LEASE6_CLTT
LEASE6_HOSTNAME
LEASE6_HWADDR
LEASE6_HWADDR_TYPE
LEASE6_STATE
LEASE6_SUBNET_ID
LEASE6_VALID_LIFETIME
LEASE6_DUID
LEASE6_IAID
LEASE6_PREFERRED_LIFETIME
LEASE6_PREFIX_LEN
LEASE6_TYPE
```

addr6_register

```
QUERY6_TYPE
QUERY6_TXID
QUERY6_LOCAL_ADDR
QUERY6_LOCAL_PORT
QUERY6_REMOTE_ADDR
QUERY6_REMOTE_PORT
QUERY6_IFACE_INDEX
QUERY6_IFACE_NAME
QUERY6_REMOTE_HWADDR
QUERY6_REMOTE_HWADDR_TYPE
```

(continues on next page)

(continued from previous page)

```

QUERY6_PROTO
QUERY6_CLIENT_ID
ADDRESS6
OLD_LEASE6_ADDRESS
OLD_LEASE6_CLTT
OLD_LEASE6_HOSTNAME
OLD_LEASE6_HWADDR
OLD_LEASE6_HWADDR_TYPE
OLD_LEASE6_STATE
OLD_LEASE6_SUBNET_ID
OLD_LEASE6_VALID_LIFETIME
OLD_LEASE6_DUID
OLD_LEASE6_IAID
OLD_LEASE6_PREFERRED_LIFETIME
OLD_LEASE6_PREFIX_LEN
OLD_LEASE6_TYPE
NEW_LEASE6_ADDRESS
NEW_LEASE6_CLTT
NEW_LEASE6_HOSTNAME
NEW_LEASE6_HWADDR
NEW_LEASE6_HWADDR_TYPE
NEW_LEASE6_STATE
NEW_LEASE6_SUBNET_ID
NEW_LEASE6_VALID_LIFETIME
NEW_LEASE6_DUID
NEW_LEASE6_IAID
NEW_LEASE6_PREFERRED_LIFETIME
NEW_LEASE6_PREFIX_LEN
NEW_LEASE6_TYPE

```

The OLD_LEASE6 does not always exist.

The leases4_committed hook point needs for loops to handle the list of addresses. This can be achieved in the following way:

```

leases4_committed() {
    for i in $(seq 0 $((LEASES4_SIZE-1))); do
        LEASE4_ADDRESS=$(eval "echo \${LEASES4_AT${i}_ADDRESS}")
        ...
    done

    for i in $(seq 0 $((DELETED_LEASES4_SIZE-1))); do
        DELETED_LEASE4_ADDRESS=$(eval "echo \${DELETED_LEASES4_AT${i}_ADDRESS}")
        ...
    done
    exit 0
}

```

The leases6_committed hook point needs for loops to handle the list of addresses. This can be achieved in the following way:

```

leases6_committed() {
    for i in $(seq 0 $((LEASES6_SIZE-1))); do

```

(continues on next page)

(continued from previous page)

```

LEASE6_ADDRESS=$(eval "echo \${LEASES6_AT${i}_ADDRESS}")
...
done

for i in $(seq 0 $((DELETED_LEASES6_SIZE-1))); do
    DELETED_LEASE6_ADDRESS=$(eval "echo \${DELETED_LEASES6_AT${i}_ADDRESS}")
    ...
done
exit 0
}

```

17.26 libdhcp_stat_cmds.so: Statistics Commands for Supplemental Lease Statistics

This library provides additional commands for retrieving lease statistics from Kea DHCP servers. These commands were added to address an issue with obtaining accurate lease statistics in deployments running multiple Kea servers that use a shared lease backend. The in-memory statistics kept by individual servers only track lease changes made by that server; thus, in a deployment with multiple servers (e.g. two *kea-dhcp6* servers using the same PostgreSQL database for lease storage), these statistics are incomplete. The MySQL and PostgreSQL backends in Kea track lease allocation changes as they occur via database triggers. Additionally, all the lease backends were extended to support retrieving lease statistics for a single subnet, a range of subnets, or all subnets. Finally, this library provides commands for retrieving these statistics.

Note

libdhcp_stat_cmds.so is part of the open source code and is available to every Kea user.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

The commands provided by this library are:

- *stat-lease4-get* - fetches DHCPv4 lease statistics.
- *stat-lease6-get* - fetches DHCPv6 lease statistics.

The Statistics Commands library is part of the open source code and is available to every Kea user.

All commands use JSON syntax and can be issued directly to the servers via either the control channel (see *Management API*) or the Control Agent (see *The Kea Control Agent*).

This library is loaded in the same way as other libraries and currently has no parameters:

```

"Dhcp6": {
  "hooks-libraries": [
    {
      "library": "/path/libdhcp_stat_cmds.so"
    },
    ...
  ]
}

```

(continues on next page)

(continued from previous page)

```
]
}
```

In a deployment with multiple Kea DHCP servers sharing a common lease storage, this hook library can be loaded by any or all of the servers. However, a server's response to a *stat-lease4-get* / *stat-lease6-get* command will only contain data for subnets known to that server. In other words, if a subnet does not appear in a server's configuration, Kea will not retrieve statistics for it.

17.26.1 The *stat-lease4-get*, *stat-lease6-get* Commands

The *stat-lease4-get* and *stat-lease6-get* commands fetch lease statistics for a range of known subnets. The range of subnets is determined through the use of optional command input parameters:

- *subnet-id* - the ID of the subnet for which lease statistics should be fetched; used to get statistics for a single subnet. If the subnet does not exist, the command result code is 3 (i.e. *CONTROL_RESULT_EMPTY*).
- *subnet-range* - a pair of subnet IDs which describe an inclusive range of subnets for which statistics should be retrieved. The range may include one or more IDs that correspond to no subnet; in this case, the command only outputs lease statistics for those that exist. However, if the range does not include any known subnets, the command result code is 3 (i.e. *CONTROL_RESULT_EMPTY*).
 - *first-subnet-id* - the ID of the first subnet in the range.
 - *last-subnet-id* - the ID of the last subnet in the range.

The use of *subnet-id* and *subnet-range* are mutually exclusive. If no parameters are given, the result will contain data for all known subnets. Note that in configurations with many subnets, this can result in a large response.

The following command fetches lease statistics for all known subnets from a *kea-dhcp4* server:

```
{
  "command": "stat-lease4-get"
}
```

The following command fetches lease statistics for subnet ID 10 from a *kea-dhcp6* server:

```
{
  "command": "stat-lease6-get",
  "arguments": {
    "subnet-id" : 10
  }
}
```

The following command fetches lease statistics for all subnets with IDs in the range 10 through 50 from a *kea-dhcp4* server:

```
{
  "command": "stat-lease4-get",
  "arguments": {
    "subnet-range": {
      "first-subnet-id": 10,
      "last-subnet-id": 50
    }
  }
}
```

The response to either command will contain three elements:

- **result** - a numeric value indicating the outcome of the command where:
 - 0 - the command was successful;
 - 1 - an error occurred, and an explanation is the "text" element; or
 - 2 - the fetch found no matching data.
- **text** - an explanation of the command outcome. When the command succeeds, it contains the command name along with the number of rows returned.
- **arguments** - a map containing the data returned by the command as the element "result-set", which is patterned after SQL statement responses:
 - **columns** - a list of text column labels.

The columns returned for DHCPv4 are:

- * **subnet-id** - the ID of the subnet.
- * **total-addresses** - the total number of addresses available for DHCPv4 management in the subnet. In other words, this is the count of all addresses in all the configured pools in the subnet.
- * **cumulative-assigned-addresses** - the cumulative number of addresses in the subnet that have been assigned to a client by the server since it started.
- * **assigned-addresses** - the number of addresses in the subnet that are currently assigned to a client.
- * **declined-addresses** - the number of addresses in the subnet that are currently declined and are thus unavailable for assignment.

The columns returned for DHCPv6 are:

- * **subnet-id** - the ID of the subnet.
- * **total-nas** - the number of NA addresses available for DHCPv6 management in the subnet. In other words, this is the count of all the NA addresses in all the configured NA pools in the subnet.
- * **cumulative-assigned-nas** - the cumulative number of NA addresses in the subnet that have been assigned to a client by the server since it started.
- * **assigned-nas** - the number of NA addresses in the subnet that are currently assigned to a client.
- * **declined-addresses** - the number of NA addresses that are currently declined and are thus unavailable for assignment.
- * **total-pds** - the total number of PD prefixes available of DHCPv6 management in the subnet. In other words, this is the count of all prefixes in all the configured prefix pools in the subnet.
- * **cumulative-assigned-pds** - the cumulative number of PD prefixes in the subnet that have been assigned to a client by the server since it started.
- * **assigned-pds** - the number of PD prefixes in the subnet that are currently assigned to a client.
- **rows** - a list of rows, one per subnet ID. Each row contains a data value corresponding to and in the same order as each column listed in "columns" for a given subnet.
- **timestamp** - the textual date and time the data were fetched, expressed as GMT.

The response to a DHCPv4 command might look as follows:

```
{
  "result": 0,
  "text": "stat-lease4-get: 2 rows found",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "result-set": {
      "columns": [ "subnet-id", "total-addresses", "cumulative-assigned-addresses",
↪ "assigned-addresses", "declined-addresses" ],
      "rows": [
        [ 10, 256, 300, 111, 0 ],
        [ 20, 4098, 2034, 2034, 4 ]
      ],
      "timestamp": "2018-05-04 15:03:37.000000"
    }
  }
}

```

The response to a DHCPv6 command might look as follows, assuming subnet 10 has no prefix pools, subnet 20 has no NA pools, and subnet 30 has both NA and PD pools:

```

{
  "result": 0,
  "text": "stat-lease6-get: 2 rows found",
  "arguments": {
    "result-set": {
      "columns": [ "subnet-id", "total-nas", "cumulative-assigned-nas", "assigned-nas",
↪ "declined-addresses", "total-pds", "cumulative-assigned-pds", "assigned-pds" ],
      "rows": [
        [ 10, 4096, 5000, 2400, 3, 0, 0, 0 ],
        [ 20, 0, 0, 0, 0, 1048, 300, 233 ],
        [ 30, 256, 60, 60, 0, 1048, 15, 15 ]
      ],
      "timestamp": "2018-05-04 15:03:37.000000"
    }
  }
}

```

17.27 libdhcp_subnet_cmds.so: Subnet Commands to Manage Subnets and Shared Networks

This library offers commands used to query and manipulate subnet and shared network configurations in Kea. These can be very useful in deployments with a large number of subnets being managed by the DHCP servers, when those subnets are frequently updated. The commands offer a lightweight approach for manipulating subnets without needing to fully reconfigure the server, and without affecting existing servers' configurations. An ability to manage shared networks (listing, retrieving details, adding new ones, removing existing ones, and adding subnets to and removing them from shared networks) is also provided.

Note

libdhcp_subnet_cmds.so is part of the open source code and is available to every Kea user. It was previously available only to ISC customers with a paid support contract.

Note

This library can only be loaded by the *kea-dhcp4* or *kea-dhcp6* process.

The following commands are currently supported:

- *subnet4-list* / *subnet6-list* - lists all configured subnets.
- *subnet4-get* / *subnet6-get* - retrieves detailed information about a specified subnet.
- *subnet4-add* / *subnet6-add* - adds a new subnet into the server's configuration.
- *subnet4-update* / *subnet6-update* - updates (replaces) a single subnet in the server's configuration.
- *subnet4-del* / *subnet6-del* - removes a subnet from the server's configuration.
- *subnet4-delta-add* / *subnet6-delta-add* - updates (replaces) parts of a single subnet in the server's configuration.
- *subnet4-delta-del* / *subnet6-delta-del* - removes parts of a single subnet in the server's configuration.
- *network4-list* / *network6-list* - lists all configured shared networks.
- *network4-get* / *network6-get* - retrieves detailed information about a specified shared network.
- *network4-add* / *network6-add* - adds a new shared network to the server's configuration.
- *network4-del* / *network6-del* - removes a shared network from the server's configuration.
- *network4-subnet-add* / *network6-subnet-add* - adds an existing subnet to an existing shared network.
- *network4-subnet-del* / *network6-subnet-del* - removes a subnet from an existing shared network and demotes it to a plain subnet.

17.27.1 The subnet4-list Command

This command is used to list all currently configured subnets. Each subnet is returned with a subnet identifier, subnet prefix and shared network name, if available. To retrieve detailed information about the subnet, use the *subnet4-get* command.

This command has a simple structure:

```
{
  "command": "subnet4-list"
}
```

The list of subnets is returned in the following format:

```
{
  "result": 0,
  "text": "2 IPv4 subnets found",
  "arguments": {
    "subnets": [
      {
        "id": 10,
        "shared-network-name": "net1",
        "subnet": "10.0.0.0/8"
      },
      {
        "id": 100,
        "shared-network-name": null,
```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.2.0/24"
    }
}
}

```

If no IPv4 subnets are found, an error code is returned along with the error description.

17.27.2 The `subnet6-list` Command

This command is used to list all currently configured subnets. Each subnet is returned with a subnet identifier, subnet prefix and shared network name, if available. To retrieve detailed information about the subnet, use the [subnet6-get](#) command.

This command has a simple structure:

```

{
    "command": "subnet6-list"
}

```

The list of subnets is returned in the following format:

```

{
    "result": 0,
    "text": "2 IPv6 subnets found",
    "arguments": {
        "subnets": [
            {
                "id": 11,
                "shared-network-name": null,
                "subnet": "2001:db8:1::/64"
            },
            {
                "id": 233,
                "shared-network-name": "some-net",
                "subnet": "3000::/16"
            }
        ]
    }
}

```

If no IPv6 subnets are found, an error code is returned along with the error description.

17.27.3 The `subnet4-get` Command

This command is used to retrieve detailed information about the specified subnet. This command usually follows [subnet4-list](#), which is used to discover available subnets with their respective subnet identifiers and prefixes. Any of those parameters can then be used in [subnet4-get](#) to fetch subnet information:

```

{
    "command": "subnet4-get",
    "arguments": {
        "id": 10
    }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

or

```
{
  "command": "subnet4-get",
  "arguments": {
    "subnet": "10.0.0.0/8"
  }
}
```

If the subnet exists, the response will be similar to this:

```
{
  "result": 0,
  "text": "Info about IPv4 subnet 10.0.0.0/8 (id 10) returned",
  "arguments": {
    "subnets": [
      {
        "subnet": "10.0.0.0/8",
        "id": 1,
        "option-data": [
          {
            ...
          },
          ...
        ],
        ...
      }
    ]
  }
}
```

17.27.4 The subnet6-get Command

This command is used to retrieve detailed information about the specified subnet. This command usually follows *subnet6-list*, which is used to discover available subnets with their respective subnet identifiers and prefixes. Any of those parameters can be then used in *subnet6-get* to fetch subnet information:

```
{
  "command": "subnet6-get",
  "arguments": {
    "id": 11
  }
}
```

or

```
{
  "command": "subnet6-get",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "subnet": "2001:db8:1::/64"
  }
}

```

If the subnet exists, the response will be similar to this:

```

{
  "result": 0,
  "text": "Info about IPv6 subnet 2001:db8:1::/64 (id 11) returned",
  "arguments": {
    "subnets": [
      {
        "subnet": "2001:db8:1::/64",
        "id": 1,
        "option-data": [
          {
            ...
          },
          ...
        ],
        ...
      }
    ]
  }
}

```

17.27.5 The subnet4-add Command

This command is used to create and add a new subnet to the existing server configuration. This operation has no impact on other subnets. The subnet identifier must be specified and must be unique among all subnets. If the identifier or a subnet prefix is not unique, an error is reported and the subnet is not added.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within *subnet4-add*. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to add, remove, and modify static reservations.

```

{
  "command": "subnet4-add",
  "arguments": {
    "subnet4": [ {
      "id": 123,
      "subnet": "10.20.30.0/24",
      ...
    } ]
  }
}

```

The response to this command has the following structure:

```

{
  "result": 0,
  "text": "IPv4 subnet added",
  "arguments": {

```

(continues on next page)

(continued from previous page)

```

    "subnet4": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}

```

17.27.6 The subnet6-add Command

This command is used to create and add a new subnet to the existing server configuration. This operation has no impact on other subnets. The subnet identifier must be specified and must be unique among all subnets. If the identifier or a subnet prefix is not unique, an error is reported and the subnet is not added.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within `subnet6-add`. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to add, remove, and modify static reservations.

```

{
  "command": "subnet6-add",
  "arguments": {
    "subnet6": [ {
      "id": 234,
      "subnet": "2001:db8:1::/64",
      ...
    } ]
  }
}

```

The response to this command has the following structure:

```

{
  "result": 0,
  "text": "IPv6 subnet added",
  "arguments": {
    "subnet6": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}

```

As of Kea 2.6.0, Kea no longer automatically generates subnet IDs; they must be specified for `subnet4-add` and `subnet6-add`. If not specified, Kea tries to assign the next `subnet-id` value. This automatic ID value generator is simple; it returns the previous automatically assigned value, increased by 1. This works well, unless a subnet is manually created with a larger value than one previously used. For example, if `subnet4-add` is called five times, each without an ID, Kea will assign IDs 1, 2, 3, 4, and 5 and everything will work properly. However, if `subnet4-add` is called five times, with the first subnet having the `subnet-id` of value 3 and the remaining ones having no `subnet-id`, the operation will fail. The first command (with the explicit value) will use `subnet-id` 3; the second command will create a subnet with an ID of 1; the third will use a value of 2; and finally the fourth will have its `subnet-id` value

auto-generated as 3. However, since there is already a subnet with that ID, the process will fail.

The general recommendation is either never to use explicit values, so that auto-generated values will always work; or always use explicit values, so that auto-generation is never used. The two approaches can be mixed only if the administrator understands how internal automatic `subnet-id` generation works in Kea.

Note

Subnet IDs must be greater than zero and less than 4294967295.

17.27.7 The `subnet4-update` Command

This command is used to update (overwrite) a single subnet in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to replace; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within `subnet4-update`. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

```
{
  "command": "subnet4-update",
  "arguments": {
    "subnet4": [ {
      "id": 123,
      "subnet": "10.20.30.0/24",
      ...
    } ]
  }
}
```

The response to this command has the following structure:

```
{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnet4": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}
```

As with other update commands, this command overwrites all the contents of the entry. If the IPv4 subnet previously had a resource assigned to it and the `subnet4-update` command is missing the resource, it is deleted from the server configuration. If an incremental update of the subnet is desired, use `subnet4-delta-add`.

17.27.8 The subnet6-update Command

This command is used to update (overwrite) a single subnet in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to replace; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within `subnet6-update`. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

```
{
  "command": "subnet6-update",
  "arguments": {
    "subnet6": [ {
      "id": 234,
      "subnet": "2001:db8:1::/64",
      ...
    } ]
  }
}
```

The response to this command has the following structure:

```
{
  "result": 0,
  "text": "IPv6 subnet updated",
  "arguments": {
    "subnet6": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}
```

As with other update commands, this command overwrites all the contents of the entry. If the IPv6 subnet previously had a resource assigned to it and the `subnet6-update` command is missing the resource, it is deleted from the server configuration. If an incremental update of the subnet is desired, use `subnet6-delta-add`.

17.27.9 The subnet4-del Command

This command is used to remove a subnet from the server's configuration. This command has no effect on other configured subnets, but removing a subnet does have certain implications.

In most cases the server has assigned some leases to the clients belonging to the subnet. The server may also be configured with static host reservations which are associated with this subnet. The current implementation of the `subnet4-del` command removes neither the leases nor the host reservations associated with a subnet. This is the safest approach because the server does not lose track of leases assigned to clients from this subnet. However, removal of the subnet may still cause configuration errors and conflicts. For example: after removal of the subnet, the server administrator may update a new subnet with the ID used previously for the removed subnet. This means that the existing leases and static reservations will be in conflict with this new subnet. Thus, we recommend that this command be used with extreme caution.

This command can also be used to completely delete an IPv4 subnet that is part of a shared network. To simply remove the subnet from a shared network and keep the subnet configuration, use the `network4-subnet-del` command instead.

The command has the following structure:

```
{
  "command": "subnet4-del",
  "arguments": {
    "id": 123
  }
}
```

A successful response may look like this:

```
{
  "result": 0,
  "text": "IPv4 subnet 192.0.2.0/24 (id 123) deleted",
  "arguments": {
    "subnets": [
      {
        "id": 123,
        "subnet": "192.0.2.0/24"
      }
    ]
  }
}
```

17.27.10 The subnet6-del Command

This command is used to remove a subnet from the server's configuration. This command has no effect on other configured subnets, but removing a subnet does have certain implications.

In most cases the server has assigned some leases to the clients belonging to the subnet. The server may also be configured with static host reservations which are associated with this subnet. The current implementation of the `subnet6-del` command removes neither the leases nor the host reservations associated with a subnet. This is the safest approach because the server does not lose track of leases assigned to clients from this subnet. However, removal of the subnet may still cause configuration errors and conflicts. For example: after removal of the subnet, the server administrator may add a new subnet with the ID used previously for the removed subnet. This means that the existing leases and static reservations will be in conflict with this new subnet. Thus, we recommend that this command be used with extreme caution.

This command can also be used to completely delete an IPv6 subnet that is part of a shared network. To simply remove the subnet from a shared network and keep the subnet configuration, use the `network6-subnet-del` command instead.

The command has the following structure:

```
{
  "command": "subnet6-del",
  "arguments": {
    "id": 234
  }
}
```

A successful response may look like this:

```
{
  "result": 0,
  "text": "IPv6 subnet 2001:db8:1::/64 (id 234) deleted",
```

(continues on next page)

(continued from previous page)

```

    "subnets": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}

```

17.27.11 The subnet4-delta-add Command

This command is used to update a subnet by adding or overwriting its parts in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to update; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within *subnet4-delta-add*. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

The command uses keys to identify the respective object. The keys cannot be updated. Objects must first be deleted in order to add other objects with conflicting keys. The address pools are identified by the *pool* parameter; the options are identified by the *name* or *code*, and *space* parameters.

```

{
  "command": "subnet4-delta-add",
  "arguments": {
    "subnet4": [ {
      "valid-lifetime": 120,
      "id": 123,
      "subnet": "10.20.30.0/24",
      "option-data": [
        {
          "always-send": false,
          "code": 3,
          "csv-format": true,
          "data": "192.0.3.1",
          "name": "routers",
          "space": "dhcp4"
        }
      ],
    },
    "pools": [
      {
        "pool": "10.20.30.1-10.20.30.10",
        "option-data": [
          {
            "always-send": false,
            "code": 4,
            "csv-format": true,
            "data": "192.0.4.1",
            "name": "time-servers",
            "space": "dhcp4"
          }
        ]
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    ]
  } ]
}

```

The response to this command has the following structure:

```

{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnet4": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}

```

The command updates subnet "10.20.30.0/24" with ID 123 by changing the valid lifetime, by adding or changing the subnet level option 3 ("routers"), by adding or changing parameters of the pool "10.20.30.1-10.20.30.10", and by adding or changing the pool level option 4 ("time-servers").

17.27.12 The subnet6-delta-add Command

This command is used to update a subnet by adding or overwriting its parts in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to update; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within `subnet6-delta-add`. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

The command uses keys to identify the respective object. The keys cannot be updated. Objects must first be deleted in order to add other objects with conflicting keys. The address pools are identified by the *pool* parameter; the prefix pools are identified by the *prefix*, *prefix-len*, and *delegated-len* parameters; and the options are identified by the *name* or *code*, and *space* parameters.

```

{
  "command": "subnet6-delta-add",
  "arguments": {
    "subnet6": [ {
      "valid-lifetime": 120,
      "id": 243,
      "subnet": "2001:db8:1::/64",
      "option-data": [
        {
          "always-send": false,
          "code": 23,
          "csv-format": true,
          "data": "3000::3:1",
          "name": "dns-servers",
          "space": "dhcp6"
        }
      ]
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "pd-pools": [
    {
      "prefix": "2001:db8:2::",
      "prefix-len": 48,
      "delegated-len": 64,
      "option-data": [
        {
          "always-send": false,
          "code": 22,
          "csv-format": true,
          "data": "3000::4:1",
          "name": "sip-server-addr",
          "space": "dhcp6"
        }
      ]
    }
  ],
  "pools": [
    {
      "pool": "2001:db8:1::1-2001:db8:1::10",
      "option-data": [
        {
          "always-send": false,
          "code": 31,
          "csv-format": true,
          "data": "3000::5:1",
          "name": "snmp-servers",
          "space": "dhcp6"
        }
      ]
    }
  ]
}

```

The response to this command has the following structure:

```

{
  "result": 0,
  "text": "IPv6 subnet updated",
  "arguments": {
    "subnet6": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}

```

The command updates subnet "2001:db8:1::/64" with ID 243 by changing the valid lifetime, by adding or changing the subnet level option 23 ("dns-servers"), by adding or changing parameters of the pool "2001:db8:1::1-2001:db8:1::10", by adding or changing the pool level option 31 ("snmp-servers"), by adding or changing parameters of the pd-pool "2001:db8:2::" with prefix-len 48, and by adding or changing the pd-pool level option 22 ("sip-server-addr").

17.27.13 The `subnet4-delta-del` Command

This command is used to update a subnet by removing its parts in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to update; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within `subnet4-delta-del`. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

The command is flexible and can delete part of the subnet by either specifying the entire object that needs to be deleted, or just the keys identifying the respective object. The address pools are identified by the *pool* parameter; and the options are identified by the *name* or *code*, and *space* parameters. The *space* parameter can be omitted if the option belongs to the default *dhcp4* space.

```
{
  "command": "subnet4-delta-del",
  "arguments": {
    "subnet4": [ {
      "valid-lifetime": 0,
      "id": 123,
      "subnet": "10.20.30.0/24",
      "option-data": [
        { "name": "routers" }
      ],
      "pools": [
        {
          "option-data": [
            { "code": 4 }
          ],
          "pool": "10.20.30.11-10.20.30.20"
        },
        {
          "pool": "10.20.30.21-10.20.30.30"
        }
      ]
    } ]
  }
}
```

The response to this command has the following structure:

```
{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnet4": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

The command updates subnet "10.20.30.0/24" with id 123 by removing the valid lifetime, removing the subnet level option 3 ("routers"), by removing the pool "10.20.30.21-10.20.30.30" and by removing the pool level option 4 ("time-servers") in pool "10.20.30.11-10.20.30.20". The scalar values don't need to match what is configured, but still need to be present to maintain a valid json structure and to be a valid value to be able to be parsed.

17.27.14 The subnet6-delta-del Command

This command is used to update a subnet by removing its parts in the existing server configuration. This operation has no impact on other subnets. The subnet identifier is used to identify the subnet to update; it must be specified and must be unique among all subnets. The subnet prefix should not be updated.

The subnet information within this command has the same structure as the subnet information in the server configuration file, with the exception that static host reservations cannot be specified within *subnet6-delta-del*. The commands described in *libdhcp_host_cmds.so: Host Commands* should be used to update, remove, and modify static reservations.

The command is flexible and can delete the part of the subnet by either specifying the entire object that needs to be deleted, or just the keys identifying the respective object. The address pools are identified by the *pool* parameter; the prefix pools are identified by the *prefix*, *prefix-len* and *delegated-len* parameters; and the options are identified by the *name* or *code*, and *space* parameters. The *space* parameter can be omitted if the option belongs to the default *dhcp6* space.

```

{
  "command": "subnet6-delta-del",
  "arguments": {
    "subnet6": [ {
      "valid-lifetime": 0,
      "id": 234,
      "subnet": "2001:db8:1::/64",
      "option-data": [
        { "name": "dns-servers" }
      ],
      "pd-pools": [
        {
          "prefix": "2001:db8:3::",
          "prefix-len": 48,
          "delegated-len": 64,
          "option-data": [
            { "code": 22 }
          ]
        },
        {
          "prefix": "2001:db8:4::",
          "prefix-len": 48,
          "delegated-len": 64
        }
      ],
      "pools": [
        {

```

(continues on next page)

(continued from previous page)

```

        "option-data": [
            { "code": 31 }
        ],
        "pool": "2001:db8:1::11-2001:db8:1::20"
    },
    {
        "pool": "2001:db8:1::21-2001:db8:1::30"
    }
]
} ]
}
}

```

The response to this command has the following structure:

```

{
    "result": 0,
    "text": "IPv6 subnet updated",
    "arguments": {
        "subnet6": [
            {
                "id": 234,
                "subnet": "2001:db8:1::/64"
            }
        ]
    }
}

```

The command updates subnet "2001:db8:1::/64" with id 243 by removing the valid lifetime, removing the subnet level option 23 ("dns-servers"), by removing the pool "2001:db8:1::21-2001:db8:1::30", by removing the pool level option 31 ("ntp-servers") in pool "2001:db8:1::11-2001:db8:1::20", by removing the pd-pool "2001:db8:4::" with prefix-len 48, by removing the pd-pool level option 22 ("sip-server-addr") in pd-pool "2001:db8:3::" with prefix-len 48. The scalar values don't need to match what is configured, but still need to be present to maintain a valid json structure and to be a valid value to be able to be parsed.

17.27.15 The `network4-list`, `network6-list` Commands

These commands are used to retrieve the full list of currently configured shared networks. The list contains only very basic information about each shared network. If more details are needed, please use [network4-get](#) or [network6-get](#) to retrieve all information available. This command does not require any parameters and its invocation is very simple:

```

{
    "command": "network4-list"
}

```

An example response for `network4-list` looks as follows:

```

{
    "arguments": {
        "shared-networks": [
            { "name": "floor1" },
            { "name": "office" }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```

},
"result": 0,
"text": "2 IPv4 network(s) found"
}

```

The `network6-list` command uses exactly the same syntax for both the command and the response.

17.27.16 The `network4-get`, `network6-get` Commands

These commands are used to retrieve detailed information about shared networks, including subnets that are currently part of a given network. Both commands take one mandatory parameter, `name`, which specifies the name of the shared network. An example command to retrieve details about an IPv4 shared network with the name "floor13" looks as follows:

```

{
  "command": "network4-get",
  "arguments": {
    "name": "floor13"
  }
}

```

An example response could look as follows:

```

{
  "result": 0,
  "text": "Info about IPv4 shared network 'floor13' returned",
  "arguments": {
    "shared-networks": [
      {
        "match-client-id": true,
        "name": "floor13",
        "option-data": [ ],
        "rebind-timer": 90,
        "relay": {
          "ip-address": "0.0.0.0"
        },
        "renew-timer": 60,
        # Specify if the server should look up global reservations.
        "reservations-global": false,
        # Specify if the server should look up in-subnet reservations.
        "reservations-in-subnet": true,
        # Specify if the server can assume that all reserved addresses
        # are out-of-pool.
        "reservations-out-of-pool": false,
        "subnet4": [
          {
            "subnet": "192.0.2.0/24",
            "id": 5,
            ...
            # many other subnet-specific details here
          },
          {
            "id": 6,

```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.3.0/31",
        ...
        # many other subnet-specific details here
    }
  ],
  "valid-lifetime": 120
}
]
}

```

The actual response contains many additional fields that are omitted here for clarity. The response format is exactly the same as used in `config-get`, just limited to returning the shared network's information.

17.27.17 The `network4-add`, `network6-add` Commands

These commands are used to add a new shared network, which must have a unique name. This command requires one parameter, `shared-networks`, which is a list and should contain exactly one entry that defines the network. The only mandatory element for a network is its name. Although it does not make operational sense, it is possible to add an empty shared network that does not have any subnets in it. That is allowed for testing purposes, but having empty networks (or networks with only one subnet) is discouraged in production environments. For details regarding syntax, see *Shared Networks in DHCPv4* and *Shared Networks in DHCPv6*.

Note

As opposed to parameter inheritance during the processing of a full new configuration, this command does not fully handle parameter inheritance. Any missing parameters will be filled with default values, rather than inherited from the global scope.

An example that showcases how to add a new IPv4 shared network looks as follows:

```

{
  "command": "network4-add",
  "arguments": {
    "shared-networks": [ {
      "name": "floor13",
      "subnet4": [
        {
          "id": 100,
          "pools": [ { "pool": "192.0.2.2-192.0.2.99" } ],
          "subnet": "192.0.2.0/24",
          "option-data": [
            {
              "name": "routers",
              "data": "192.0.2.1"
            }
          ]
        }
      ]
    }
  ],
  {
    "id": 101,
    "pools": [ { "pool": "192.0.3.2-192.0.3.99" } ],

```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.3.0/24",
        "option-data": [
            {
                "name": "routers",
                "data": "192.0.3.1"
            }
        ]
    } ]
} ]
}
}

```

Assuming there was no shared network with a name "floor13" and no subnets with IDs 100 and 101 previously configured, the command will be successful and will return the following response:

```

{
  "arguments": {
    "shared-networks": [ { "name": "floor13" } ]
  },
  "result": 0,
  "text": "A new IPv4 shared network 'floor13' added"
}

```

The `network6-add` command uses the same syntax for both the query and the response. However, there are some parameters that are IPv4-only (e.g. `match-client-id`) and some that are IPv6-only (e.g. `interface-id`). The same applies to subnets within the network.

17.27.18 The `network4-del`, `network6-del` Commands

These commands are used to delete existing shared networks. Both commands take exactly one parameter, `name`, that specifies the name of the network to be removed. An example invocation of the `network4-del` command looks as follows:

```

{
  "command": "network4-del",
  "arguments": {
    "name": "floor13"
  }
}

```

Assuming such a network was configured, the response looks similar to the following:

```

{
  "arguments": {
    "shared-networks": [
      {
        "name": "floor13"
      }
    ]
  },
  "result": 0,
  "text": "IPv4 shared network 'floor13' deleted"
}

```

The `network6-del` command uses exactly the same syntax for both the command and the response.

If there are any subnets belonging to the shared network being deleted, they will be demoted to a plain subnet. There is an optional parameter called `subnets-action` that, if specified, takes one of two possible values: `keep` (which is the default) and `delete`. It controls whether the subnets are demoted to plain subnets or removed. An example usage in the `network6-del` command that deletes the shared network and all subnets in it could look as follows:

```
{
  "command": "network4-del",
  "arguments": {
    "name": "floor13",
    "subnets-action": "delete"
  }
}
```

Alternatively, to completely remove the subnets, it is possible to use the `subnet4-del` or `subnet6-del` commands.

17.27.19 The `network4-subnet-add`, `network6-subnet-add` Commands

These commands are used to add existing subnets to existing shared networks. There are several ways to add a new shared network. The system administrator can add the whole shared network at once, either by editing a configuration file or by calling the `network4-add` or `network6-add` command with the desired subnets in it. This approach works well for completely new shared subnets. However, there may be cases when an existing subnet is running out of addresses and needs to be extended with additional address space; in other words, another subnet needs to be added on top of it. For this scenario, a system administrator can use `network4-add` or `network6-add`, and then add an existing subnet to this newly created shared network using `network4-subnet-add` or `network6-subnet-add`.

The `network4-subnet-add` and `network6-subnet-add` commands take two parameters: `id`, which is an integer and specifies the ID of an existing subnet to be added to a shared network; and `name`, which specifies the name of the shared network to which the subnet will be added. The subnet must not belong to any existing network; to reassign a subnet from one shared network to another, use the `network4-subnet-del` or `network6-subnet-del` command first.

An example invocation of the `network4-subnet-add` command looks as follows:

```
{
  "command": "network4-subnet-add",
  "arguments": {
    "name": "floor13",
    "id": 5
  }
}
```

Assuming there is a network named "floor13", and there is a subnet with `subnet-id 5` that is not a part of the existing network, the command will return a response similar to the following:

```
{
  "result": 0,
  "text": "IPv4 subnet 10.0.0.0/8 (id 5) is now part of shared network 'floor13'"
}
```

The `network6-subnet-add` command uses exactly the same syntax for both the command and the response.

Note

As opposed to parameter inheritance during the processing of a full new configuration or when adding a new shared network with new subnets, this command does not fully handle parameter inheritance. Any missing parameters will be filled with default values, rather than inherited from the global scope or from the shared network.

17.27.20 The `network4-subnet-del`, `network6-subnet-del` Commands

These commands are used to remove a subnet that is part of an existing shared network and demote it to a plain, stand-alone subnet. To remove a subnet completely, use the `subnet4-del` or `subnet6-del` commands instead. The `network4-subnet-del` and `network6-subnet-del` commands take two parameters: `id`, which is an integer and specifies the ID of an existing subnet to be removed from a shared network; and `name`, which specifies the name of the shared network from which the subnet will be removed.

An example invocation of the `network4-subnet-del` command looks as follows:

```
{
  "command": "network4-subnet-del",
  "arguments": {
    "name": "floor13",
    "id": 5
  }
}
```

Assuming there was a subnet with `subnet-id 5`, that was part of a shared network named "floor13", the response would look similar to the following:

```
{
  "result": 0,
  "text": "IPv4 subnet 10.0.0.0/8 (id 5) is now removed from shared network 'floor13'"
}
```

The `network6-subnet-del` command uses exactly the same syntax for both the command and the response.

17.28 `libdhcp_user_chk.so`: User Check

This library serves several purposes:

- To assign "new" or "unregistered" users to a restricted subnet, while "known" or "registered" users are assigned to unrestricted subnets.
- To allow DHCP response options or vendor option values to be customized based on user identity.
- To provide a real-time record of user registration activity, which can be sampled by an external consumer.
- To serve as a demonstration of various capabilities possible using the hooks interface.

Note

`libdhcp_user_chk.so` is part of the open source code and is available to every Kea user.

Once loaded, the library allows the separation of incoming requests into known and unknown clients. For known clients, packets are processed as usual, although it is possible to override the sending of certain options on a per-host basis. Clients that are not on the known hosts list are treated as unknown and are assigned to the last subnet defined in the configuration file.

As an example of a use case, this behavior may be implemented to put unknown users into a separate subnet that leads to a "walled garden," where they can only access a registration portal. Once they fill in necessary data, their details are added to the known clients file and they get a proper address after their device is restarted.

Note

This library was developed several years before the host reservation mechanism became available. Host reservation is much more powerful and flexible, but the ability of `user_chk` to consult an external source of information about clients and alter Kea's behavior remains useful and of educational value.

The library reads the `/usr/local/share/kea/user_chk_registry.txt` file while being loaded and each time an incoming packet is processed. Each line of the file is expected to contain a self-contained JSON snippet which must have the following two entries:

- `type` - whose value is `"HW_ADDR"` for IPv4 users or `"DUID"` for IPv6 users.
- `id` - whose value is either the hardware address or the DUID from the request formatted as a string of hex digits, with or without ":" delimiters.

and may have zero or more of the following entries:

- `bootfile` - whose value is the pathname of the desired file.
- `tftp_server` - whose value is the hostname or IP address of the desired server.

A sample user registry file is shown below:

```
{ "type" : "HW_ADDR", "id" : "0c:0e:0a:01:ff:04", "bootfile" : "/usr/local/share/kea/
↪v4bootfile" }
{ "type" : "HW_ADDR", "id" : "0c:0e:0a:01:ff:06", "tftp_server" : "tftp.v4.example.com" }
{ "type" : "DUID", "id" : "00:01:00:01:19:ef:e6:3b:00:0c:01:02:03:04", "bootfile" : "/"
↪usr/local/share/kea/v6bootfile" }
{ "type" : "DUID", "id" : "00:01:00:01:19:ef:e6:3b:00:0c:01:02:03:06", "tftp_server" :
↪"tftp.v6.example.com" }
```

As with any other hook libraries provided by ISC, internals of the `user_chk` code are well-documented. Users may refer to the `user_chk` library section of the [Kea Developer's Guide](#) for information on how the code works internally. That, together with the [Hooks Framework](#) section of the [Kea Developer's Guide](#) should give users some pointers on how to extend this library and perhaps even write one from scratch.

STATISTICS

18.1 Statistics Overview

Both Kea DHCP servers support statistics gathering. A working DHCP server encounters various events that can cause certain statistics to be collected. For example, a DHCPv4 server may receive a packet (the `pkt4-received` statistic increases by one) that after parsing is identified as a DHCPDISCOVER (`pkt4-discover-received`). The server processes it and decides to send a DHCPOFFER representing its answer (the `pkt4-offer-sent` and `pkt4-sent statistics` increase by one). Such events happen frequently, so it is not uncommon for the statistics to have values in the high thousands. They can serve as an easy and powerful tool for observing a server's and a network's health. For example, if the `pkt4-received` statistic stops growing, it means that the clients' packets are not reaching the server.

There are four types of statistics:

- *integer* - this is the most common type. It is implemented as a 64-bit integer (`int64_t` in C++), so it can hold any value between -2^{63} to $2^{63}-1$.
- *big integer* - this type is intended to hold large numbers. It is implemented as a 128-bit integer (`boost::multiprecision::int128_t` in C++), so it can hold any value between -2^{127} to $2^{127}-1$.
- *floating point* - this type is intended to store floating-point precision. It is implemented as a C++ double type.
- *duration* - this type is intended for recording time periods. It uses the `boost::posix_time::time_duration` type, which stores hours, minutes, seconds, and microseconds.
- *string* - this type is intended for recording statistics in text form. It uses the C++ `std::string` type.

During normal operation, the DHCPv4 and DHCPv6 servers gather statistics. For a list of DHCPv4 and DHCPv6 statistics, see *Statistics in the DHCPv4 Server* and *Statistics in the DHCPv6 Server*, respectively.

To extract data from the statistics module, the control channel can be used. See *Management API* for details. It is possible to retrieve a single statistic or all statistics, reset the statistics (i.e. set them to a neutral value, typically zero), or even completely remove a single statistic or all statistics. See the section *Commands for Manipulating Statistics* for a list of statistics-oriented commands.

Statistics can be used by external tools to monitor Kea. One example of such a tool is Stork. See *Monitoring Kea With Stork* for details on how to use it and other data sources to retrieve statistics periodically to get better insight into Kea's health and operational status.

18.2 Statistics Lifecycle

All of the statistics supported by Kea's servers are initialized upon the servers' startup and are returned in response to the commands such as `statistic-get-all`. The runtime statistics concerning DHCP packets processed are initially set to 0 and are reset upon the server restart.

Per-subnet statistics are recalculated when reconfiguration takes place.

In general, once a statistic is initialized it is held in the manager until explicitly removed, via `statistic-remove` or `statistic-remove-all`, or when the server is shut down.

Removing a statistic that is updated frequently makes little sense, as it will be re-added when the server code next records that statistic. The `statistic-remove` and `statistic-remove-all` commands are intended to remove statistics that are not expected to be observed in the near future. For example, a misconfigured device in a network may cause clients to report duplicate addresses, so the server will report increasing values of `pkt4-decline-received`. Once the problem is found and the device is removed, the system administrator may want to remove the `pkt4-decline-received` statistic so that it is no longer reported, until and unless a duplicate address is again detected.

18.3 Commands for Manipulating Statistics

There are several commands defined that can be used for accessing (`-get`), resetting to zero or a neutral value (`-reset`), or removing a statistic completely (`-remove`). The statistics time-based limit (`-sample-age-set`) and size-based limit (`-sample-count-set`), which control how long or how many samples of a given statistic are retained, can also be changed.

The difference between `-reset` and `-remove` is somewhat subtle. The `-reset` command sets the value of the statistic to zero or a neutral value, so that after this operation, the statistic has a value of 0 (integer), 0.0 (float), 0h0m0s0us (duration), or "" (string). When requested, a statistic with the values mentioned is returned. `-remove` removes a statistic completely, so the statistic is no longer reported. However, the server code may add it back if there is a reason to record it.

Note

The following sections describe commands that can be sent to the server; the examples are not fragments of a configuration file. For more information on sending commands to Kea, see [Management API](#).

18.3.1 The `statistic-get` Command

The `statistic-get` command retrieves a single statistic. It takes a single-string parameter called `name`, which specifies the statistic name. An example command may look like this:

```
{
  "command": "statistic-get",
  "arguments": {
    "name": "pkt4-received"
  }
}
```

The server returns details of the requested statistic, with a result of 0 indicating success and the specified statistic as the value of the `arguments` parameter. If the requested statistic is not found, the response contains an empty map, i.e. only `{ }` as an argument, but the status code still indicates success (0).

Here is an example response:

```
{
  "command": "statistic-get",
  "arguments": {
    "pkt4-received": [ [ 125, "2019-07-30 10:11:19.498739" ], [ 100, "2019-07-30
↪10:11:19.498662" ] ]
  },
  "result": 0
}
```

18.3.2 The `statistic-reset` Command

The `statistic-reset` command sets the specified statistic to its neutral value: 0 for integer, 0.0 for float, 0h0m0s0us for time duration, and "" for string type. It takes a single-string parameter called `name`, which specifies the statistic name. An example command may look like this:

```
{
  "command": "statistic-reset",
  "arguments": {
    "name": "pkt4-received"
  }
}
```

If the specific statistic is found and the reset is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

18.3.3 The `statistic-remove` Command

The `statistic-remove` command deletes a single statistic. It takes a single-string parameter called `name`, which specifies the statistic name. An example command may look like this:

```
{
  "command": "statistic-remove",
  "arguments": {
    "name": "pkt4-received"
  }
}
```

If the specific statistic is found and its removal is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

18.3.4 The `statistic-get-all` Command

The `statistic-get-all` command retrieves all statistics recorded. An example command may look like this:

```
{
  "command": "statistic-get-all",
  "arguments": { }
}
```

The server responds with details of all recorded statistics, with a result set to 0 to indicate that it iterated over all statistics (even when the total number of statistics is zero).

Here is an example response returning all collected statistics:

```
{
  "command": "statistic-get-all",
  "arguments": {
    "cumulative-assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836166"
      ]
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```
],
"declined-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836156"
  ]
],
"pkt4-ack-received": [
  [
    0,
    "2023-06-13 20:42:46.616247"
  ]
],
"pkt4-ack-sent": [
  [
    0,
    "2023-06-13 20:42:46.616290"
  ]
],
"pkt4-decline-received": [
  [
    0,
    "2023-06-13 20:42:46.616296"
  ]
],
"pkt4-discover-received": [
  [
    0,
    "2023-06-13 20:42:46.616303"
  ]
],
"pkt4-inform-received": [
  [
    0,
    "2023-06-13 20:42:46.616308"
  ]
],
"pkt4-nak-received": [
  [
    0,
    "2023-06-13 20:42:46.616312"
  ]
],
"pkt4-nak-sent": [
  [
    0,
    "2023-06-13 20:42:46.616314"
  ]
],
"pkt4-offer-received": [
  [
    0,
```

(continues on next page)

(continued from previous page)

```

        "2023-06-13 20:42:46.616318"
    ],
    "pkt4-offer-sent": [
        [
            0,
            "2023-06-13 20:42:46.616323"
        ]
    ],
    "pkt4-parse-failed": [
        [
            0,
            "2023-06-13 20:42:46.616326"
        ]
    ],
    "pkt4-receive-drop": [
        [
            0,
            "2023-06-13 20:42:46.616330"
        ]
    ],
    "pkt4-received": [
        [
            0,
            "2023-06-13 20:42:46.616335"
        ]
    ],
    "pkt4-release-received": [
        [
            0,
            "2023-06-13 20:42:46.616339"
        ]
    ],
    "pkt4-request-received": [
        [
            0,
            "2023-06-13 20:42:46.616343"
        ]
    ],
    "pkt4-sent": [
        [
            0,
            "2023-06-13 20:42:46.616348"
        ]
    ],
    "pkt4-unknown-received": [
        [
            0,
            "2023-06-13 20:42:46.616354"
        ]
    ],
    "reclaimed-declined-addresses": [

```

(continues on next page)

(continued from previous page)

```

    [
      0,
      "2023-06-13 20:42:46.836159"
    ],
    "reclaimed-leases": [
      [
        0,
        "2023-06-13 20:42:46.836163"
      ]
    ],
    "subnet[1].assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836173"
      ]
    ],
    "subnet[1].cumulative-assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836098"
      ]
    ],
    "subnet[1].declined-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836178"
      ]
    ],
    "subnet[1].pool[0].assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836205"
      ]
    ],
    "subnet[1].pool[0].cumulative-assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836137"
      ]
    ],
    "subnet[1].pool[0].declined-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836213"
      ]
    ],
    "subnet[1].pool[0].reclaimed-declined-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836225"
      ]
    ]
  ]

```

(continues on next page)

(continued from previous page)

```

],
"subnet[1].pool[0].reclaimed-leases": [
  [
    0,
    "2023-06-13 20:42:46.836236"
  ]
],
"subnet[1].pool[0].total-addresses": [
  [
    11010049,
    "2023-06-13 20:42:46.836128"
  ]
],
"subnet[1].reclaimed-declined-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836186"
  ]
],
"subnet[1].reclaimed-leases": [
  [
    0,
    "2023-06-13 20:42:46.836194"
  ]
],
"subnet[1].total-addresses": [
  [
    11010049,
    "2023-06-13 20:42:46.836083"
  ]
],
"subnet[1].v4-lease-reuses": [
  [
    0,
    "2023-06-13 20:42:46.836105"
  ]
],
"subnet[1].v4-reservation-conflicts": [
  [
    0,
    "2023-06-13 20:42:46.836111"
  ]
],
"v4-allocation-fail": [
  [
    0,
    "2023-06-13 20:42:46.616358"
  ]
],
"v4-allocation-fail-classes": [
  [
    0,

```

(continues on next page)

(continued from previous page)

```

        "2023-06-13 20:42:46.616363"
    ],
    "v4-allocation-fail-no-pools": [
        [
            0,
            "2023-06-13 20:42:46.616368"
        ]
    ],
    "v4-allocation-fail-shared-network": [
        [
            0,
            "2023-06-13 20:42:46.616372"
        ]
    ],
    "v4-allocation-fail-subnet": [
        [
            0,
            "2023-06-13 20:42:46.616376"
        ]
    ],
    "v4-lease-reuses": [
        [
            0,
            "2023-06-13 20:42:46.616410"
        ]
    ],
    "v4-reservation-conflicts": [
        [
            0,
            "2023-06-13 20:42:46.616412"
        ]
    ]
},
"result": 0
}

```

or

```

{
  "command": "statistic-get-all",
  "arguments": {
    "cumulative-assigned-nas": [
      [
        0,
        "2023-06-13 21:28:57.196757"
      ]
    ],
    "cumulative-assigned-pds": [
      [
        0,
        "2023-06-13 21:28:57.196758"
      ]
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    ]
  ],
  "cumulative-registered": [
    [
      0,
      "2023-06-13 21:28:57.196758"
    ]
  ],
  "declined-addresses": [
    [
      0,
      "2023-06-13 21:28:57.196754"
    ]
  ],
  "pkt6-addr-reg-inform-received": [
    [
      0,
      "2023-06-13 21:28:57.177731"
    ]
  ],
  "pkt6-addr-reg-reply-received": [
    [
      0,
      "2023-06-13 21:28:57.177731"
    ]
  ],
  "pkt6-addr-reg-reply-sent": [
    [
      0,
      "2023-06-13 21:28:57.177731"
    ]
  ],
  "pkt6-advertise-received": [
    [
      0,
      "2023-06-13 21:28:57.177731"
    ]
  ],
  "pkt6-advertise-sent": [
    [
      0,
      "2023-06-13 21:28:57.177739"
    ]
  ],
  "pkt6-decline-received": [
    [
      0,
      "2023-06-13 21:28:57.177739"
    ]
  ],
  "pkt6-dhcpv4-query-received": [
    [

```

(continues on next page)

(continued from previous page)

```

        0,
        "2023-06-13 21:28:57.177740"
    ]
],
"pkt6-dhcpv4-response-received": [
    [
        0,
        "2023-06-13 21:28:57.177740"
    ]
],
"pkt6-dhcpv4-response-sent": [
    [
        0,
        "2023-06-13 21:28:57.177741"
    ]
],
"pkt6-infrequest-received": [
    [
        0,
        "2023-06-13 21:28:57.177742"
    ]
],
"pkt6-parse-failed": [
    [
        0,
        "2023-06-13 21:28:57.177742"
    ]
],
"pkt6-rebind-received": [
    [
        0,
        "2023-06-13 21:28:57.177743"
    ]
],
"pkt6-receive-drop": [
    [
        0,
        "2023-06-13 21:28:57.177743"
    ]
],
"pkt6-received": [
    [
        0,
        "2023-06-13 21:28:57.177744"
    ]
],
"pkt6-release-received": [
    [
        0,
        "2023-06-13 21:28:57.177744"
    ]
],
],

```

(continues on next page)

(continued from previous page)

```
"pkt6-renew-received": [  
  [  
    0,  
    "2023-06-13 21:28:57.177745"  
  ]  
],  
"pkt6-reply-received": [  
  [  
    0,  
    "2023-06-13 21:28:57.177745"  
  ]  
],  
"pkt6-reply-sent": [  
  [  
    0,  
    "2023-06-13 21:28:57.177746"  
  ]  
],  
"pkt6-request-received": [  
  [  
    0,  
    "2023-06-13 21:28:57.177747"  
  ]  
],  
"pkt6-sent": [  
  [  
    0,  
    "2023-06-13 21:28:57.177747"  
  ]  
],  
"pkt6-solicit-received": [  
  [  
    0,  
    "2023-06-13 21:28:57.177748"  
  ]  
],  
"pkt6-unknown-received": [  
  [  
    0,  
    "2023-06-13 21:28:57.177748"  
  ]  
],  
"reclaimed-declined-addresses": [  
  [  
    0,  
    "2023-06-13 21:28:57.196755"  
  ]  
],  
"reclaimed-leases": [  
  [  
    0,  
    "2023-06-13 21:28:57.196756"
```

(continues on next page)

(continued from previous page)

```

    ]
  ],
  "subnet[1].assigned-nas": [
    [
      0,
      "2023-06-13 21:28:57.196760"
    ]
  ],
  "subnet[1].assigned-pds": [
    [
      0,
      "2023-06-13 21:28:57.196761"
    ]
  ],
  "subnet[1].cumulative-assigned-nas": [
    [
      0,
      "2023-06-13 21:28:57.196727"
    ]
  ],
  "subnet[1].cumulative-assigned-pds": [
    [
      0,
      "2023-06-13 21:28:57.196729"
    ]
  ],
  "subnet[1].cumulative-registered": [
    [
      0,
      "2023-06-13 21:28:57.196727"
    ]
  ],
  "subnet[1].declined-addresses": [
    [
      0,
      "2023-06-13 21:28:57.196763"
    ]
  ],
  "subnet[1].pd-pool[0].assigned-pds": [
    [
      0,
      "2023-06-13 21:28:57.196785"
    ]
  ],
  "subnet[1].pd-pool[0].cumulative-assigned-pds": [
    [
      0,
      "2023-06-13 21:28:57.196744"
    ]
  ],
  "subnet[1].pd-pool[0].reclaimed-leases": [

```

(continues on next page)

(continued from previous page)

```

        0,
        "2023-06-13 21:28:57.196789"
    ]
},
"subnet[1].pd-pool[0].total-pds": [
    [
        256,
        "2023-06-13 21:28:57.196741"
    ]
},
"subnet[1].pool[0].assigned-nas": [
    [
        0,
        "2023-06-13 21:28:57.196773"
    ]
},
"subnet[1].pool[0].cumulative-assigned-nas": [
    [
        0,
        "2023-06-13 21:28:57.196739"
    ]
},
"subnet[1].pool[0].declined-addresses": [
    [
        0,
        "2023-06-13 21:28:57.196775"
    ]
},
"subnet[1].pool[0].reclaimed-declined-addresses": [
    [
        0,
        "2023-06-13 21:28:57.196779"
    ]
},
"subnet[1].pool[0].reclaimed-leases": [
    [
        0,
        "2023-06-13 21:28:57.196783"
    ]
},
"subnet[1].pool[0].total-nas": [
    [
        281474976710656,
        "2023-06-13 21:28:57.196736"
    ]
},
"subnet[1].reclaimed-declined-addresses": [
    [
        0,
        "2023-06-13 21:28:57.196766"
    ]
},
],

```

(continues on next page)

(continued from previous page)

```

"subnet[1].reclaimed-leases": [
  [
    0,
    "2023-06-13 21:28:57.196770"
  ]
],
"subnet[1].registered": [
  [
    0,
    "2023-06-13 21:28:57.196727"
  ]
],
"subnet[1].total-nas": [
  [
    281474976710656,
    "2023-06-13 21:28:57.196720"
  ]
],
"subnet[1].total-pds": [
  [
    256,
    "2023-06-13 21:28:57.196724"
  ]
],
"subnet[1].v6-ia-na-lease-reuses": [
  [
    0,
    "2023-06-13 21:28:57.196731"
  ]
],
"subnet[1].v6-ia-pd-lease-reuses": [
  [
    0,
    "2023-06-13 21:28:57.196733"
  ]
],
"v6-allocation-fail": [
  [
    0,
    "2023-06-13 21:28:57.177749"
  ]
],
"v6-allocation-fail-classes": [
  [
    0,
    "2023-06-13 21:28:57.177755"
  ]
],
"v6-allocation-fail-no-pools": [
  [
    0,
    "2023-06-13 21:28:57.177756"
  ]
]

```

(continues on next page)

(continued from previous page)

```

    ]
  ],
  "v6-allocation-fail-shared-network": [
    [
      0,
      "2023-06-13 21:28:57.177756"
    ]
  ],
  "v6-allocation-fail-subnet": [
    [
      0,
      "2023-06-13 21:28:57.177757"
    ]
  ],
  "v6-ia-na-lease-reuses": [
    [
      0,
      "2023-06-13 21:28:57.177757"
    ]
  ],
  "v6-ia-pd-lease-reuses": [
    [
      0,
      "2023-06-13 21:28:57.177758"
    ]
  ]
},
"result": 0
}

```

or

```

{
  "command": "statistic-get-all",
  "arguments": {
    "ncr-error": [
      [
        0,
        "2023-06-13 21:42:54.627751"
      ]
    ],
    "ncr-invalid": [
      [
        0,
        "2023-06-13 21:42:54.627749"
      ]
    ],
    "ncr-received": [
      [
        0,
        "2023-06-13 21:42:54.627737"
      ]
    ]
  ]
}

```

(continues on next page)

(continued from previous page)

```

    ],
    "queue-mgr-queue-full": [
        [
            0,
            "2023-06-13 21:42:54.627737"
        ]
    ],
    "update-error": [
        [
            0,
            "2023-06-13 21:42:54.627759"
        ]
    ],
    "update-sent": [
        [
            0,
            "2023-06-13 21:42:54.627752"
        ]
    ],
    "update-signed": [
        [
            0,
            "2023-06-13 21:42:54.627753"
        ]
    ],
    "update-success": [
        [
            0,
            "2023-06-13 21:42:54.627755"
        ]
    ],
    "update-timeout": [
        [
            0,
            "2023-06-13 21:42:54.627757"
        ]
    ],
    "update-unsigned": [
        [
            0,
            "2023-06-13 21:42:54.627754"
        ]
    ]
  },
  "result": 0
}

```

18.3.5 The `statistic-reset-all` Command

The `statistic-reset` command sets all statistics to their neutral values: 0 for integer, 0.0 for float, 0h0m0s0us for time duration, and "" for string type. An example command may look like this:

```
{
  "command": "statistic-reset-all",
  "arguments": { }
}
```

If the operation is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

18.3.6 The `statistic-remove-all` Command

The `statistic-remove-all` command attempts to delete all statistics. An example command may look like this:

```
{
  "command": "statistic-remove-all",
  "arguments": { }
}
```

If the removal of all statistics is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

18.3.7 The `statistic-sample-age-set` Command

The `statistic-sample-age-set` command sets a time-based limit on samples for a given statistic. It takes two parameters: a string called `name`, which specifies the statistic name, and an integer value called `duration`, which specifies the time limit for the given statistic in seconds. An example command may look like this:

```
{
  "command": "statistic-sample-age-set",
  "arguments": {
    "name": "pkt4-received",
    "duration": 1245
  }
}
```

If the command is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

18.3.8 The `statistic-sample-age-set-all` Command

The `statistic-sample-age-set-all` command sets time-based limits on samples for all statistics. It takes a single-integer parameter called `duration`, which specifies the time limit for the statistic in seconds. An example command may look like this:

```
{
  "command": "statistic-sample-age-set-all",
  "arguments": {
    "duration": 1245
  }
}
```


If the command is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

18.3.9 The `statistic-sample-count-set` Command

The `statistic-sample-count-set` command sets a size-based limit on samples for a given statistic. An example command may look like this:

```
{
  "command": "statistic-sample-count-set",
  "arguments": {
    "name": "pkt4-received",
    "max-samples": 100
  }
}
```

If the command is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

18.3.10 The `statistic-sample-count-set-all` Command

The `statistic-sample-count-set-all` command sets size-based limits on samples for all statistics. An example command may look like this:

```
{
  "command": "statistic-sample-count-set-all",
  "arguments": {
    "max-samples": 100
  }
}
```

If the command is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

18.4 Time Series

With certain statistics, a single isolated data point may be useful. However, some statistics, such as received packet size, packet processing time, or number of database queries needed to process a packet, are not cumulative and it is useful to keep many data points, perhaps to do some statistical analysis afterwards.

Each Kea statistic holds 20 data points; setting such a limit prevents unlimited memory growth. There are two ways to define the limits: time-based (e.g. keep samples from the last 5 minutes) and size-based. The size-based limit can be changed using one of two commands: `statistic-sample-count-set`, to set a size limit for a single statistic, and `statistic-sample-count-set-all`, to set size-based limits for all statistics. To set time-based limits for a single statistic, use `statistic-sample-age-set`; use `statistic-sample-age-set-all` to set time-based limits for all statistics. For a given statistic only one type of limit can be active; storage is limited by either time or size, not both.

MANAGEMENT API

A classic approach to daemon configuration assumes that the server's configuration is stored in configuration files and, when the configuration is changed, the daemon is restarted. This approach has the significant disadvantage of introducing periods of downtime when client traffic is not handled. Another risk is that if the new configuration is invalid for any reason, the server may refuse to start, which will further extend the downtime period until the issue is resolved.

To avoid such problems, the DHCPv4, DHCPv6, and D2 servers in Kea include support for a mechanism that allows online reconfiguration without requiring server shutdown. Both servers can be instructed to open control sockets, which is a communications channel. The server is able to receive commands on that channel, act on them, and report back status.

The DHCPv4, DHCPv6, and D2 servers receive commands over the UNIX domain sockets. For details on how to configure these sockets, see [UNIX Control Socket](#), [UNIX Control Socket](#) and [UNIX Control Socket](#). While it is possible to control the servers directly using UNIX domain sockets, that requires that the controlling client be running on the same machine as the server. SSH is usually used to connect remotely to the controlled machine.

Network administrators usually prefer using some form of a RESTful API to control the servers, rather than using UNIX domain sockets directly. Therefore, Kea includes a component called the Control Agent (CA), which exposes a RESTful API to the controlling clients and can forward commands to the respective Kea services over the UNIX domain sockets. The CA configuration is described in [Configuration](#).

The HTTP requests received by the CA contain the control commands encapsulated within HTTP requests. Simply speaking, the CA is responsible for stripping the HTTP layer from the received commands and forwarding the commands in a JSON format over the UNIX domain sockets to the respective services. Because the CA receives commands for all services, it requires additional "forwarding" information to be included in the client's messages. This forwarding information is carried within the `service` parameter of the received command. If the `service` parameter is not included, or if the parameter is a blank list, the CA assumes that the control command is targeted at the CA itself and attempts to respond.

Control connections over both HTTP and UNIX domain sockets are guarded with timeouts. The timeout value is set to 10 seconds and is not configurable.

This API can be used by external tools to manage and monitor Kea operation. An example of such a monitoring tool is ISC's Stork. For details, see [Monitoring Kea With Stork](#).

19.1 Data Syntax

Communication over the control channel is conducted using JSON structures. If configured, Kea opens a socket and listens for incoming connections. A process connecting to this socket is expected to send JSON commands structured as follows:

```
{  
  "command": "foo",
```

(continues on next page)

(continued from previous page)

```

    "service": [ "dhcp4" ],
    "arguments": {
        "param1": "value1",
        "param2": "value2",
        ...
    }
}

```

The same command sent over the RESTful interface to the CA has the following structure:

```

POST / HTTP/1.1\r\n
Content-Type: application/json\r\n
Content-Length: 147\r\n\r\n
{
    "command": "foo",
    "service": [ "dhcp4" ],
    "arguments": {
        "param1": "value1",
        "param2": "value2",
        ...
    }
}

```

The `command` parameter contains the name of the command to execute and it is mandatory. The `arguments` map contains the parameters required to carry out the given command. The exact content and format of the map are command-specific.

The `service` list contains the servers at which the control command is targeted. In the example above, the control command is targeted at the DHCPv4 server. In most cases, the CA simply forwards this command to the DHCPv4 server for processing via a UNIX domain socket. Sometimes, the command including a service value may also be processed by the CA, if the CA is running a hook library which handles such a command for the given server. As an example, the hook library loaded by the CA may perform some operations on the database, such as adding host reservations, modifying leases, etc. An advantage of performing DHCPv4-specific administrative operations in the CA, rather than forwarding it to the DHCPv4 server, is the ability to perform these operations without disrupting the DHCPv4 service, since the DHCPv4 server does not have to stop processing DHCP messages to apply changes to the database. Nevertheless, these situations are rather rare; in most cases, when the `service` parameter contains a name of the service, the commands are simply forwarded by the CA. The forwarded command includes the `service` parameter, but this parameter is ignored by the receiving server. This parameter is only meaningful to the CA.

If the command received by the CA does not include a `service` parameter or this list is empty, the CA simply processes this message on its own. For example, a `config-get` command which includes no service parameter returns the Control Agent's own configuration. The `config-get` command with a service value "dhcp4" is forwarded to the DHCPv4 server and returns the DHCPv4 server's configuration.

The following list shows the mapping of the values carried within the `service` parameter to the servers to which the commands are forwarded:

- `dhcp4` - the command is forwarded to the `kea-dhcp4` server.
- `dhcp6` - the command is forwarded to the `kea-dhcp6` server.
- `d2` - the command is forwarded to the `kea-dhcp-ddns` server.

The server processing the incoming command sends a response of the form:

```
{
  "result": 0, // 0|1|2|3|4
  "text": "textual description",
  "arguments": {
    "argument1": "value1",
    "argument2": "value2",
    ...
  }
}
```

The `result` value is a status code indicating a result of the command. The following general status codes are currently supported:

- 0 - the command has been processed successfully.
- 1 - a general error or failure has occurred during the command processing.
- 2 - the specified command is unsupported by the server receiving it.
- 3 - the requested operation has been completed but the requested resource was not found. This status code is returned when a command returns no resources or affects no resources.
- 4 - the well-formed command has been processed but the requested changes could not be applied, because they were in conflict with the server state or its notion of the configuration.

For example, a well-formed command that requests a subnet that exists in a server's configuration returns the result 0. If the server encounters an error condition, it returns 1. If the command asks for an IPv6 subnet, but was sent to a DHCPv4 server, it returns 2. If the query asks for a subnet with `subnet-id` that has matches, the result is 3. If the command attempts to update a lease but the specified `subnet-id` does not match the identifier in the server's configuration, the result is 4.

Hook libraries can sometimes return additional status codes specific to their use cases.

The `text` field typically appears when the result is non-zero and contains a description of the error encountered, but it often also appears for successful outcomes. The exact text is command-specific, but in general uses plain English to describe the outcome of the command. The `arguments` map contains additional data values returned by the server which are specific to the command issued. The map may be present, but that depends on the specific command.

Note

Since Kea 1.9.7, it is possible to put comments in commands as in the configuration file. For instance:

```
{
  "command": "foo",
  // service is a list
  "service": [ "dhcp4" ],
  # command arguments are here.
  "arguments": {
    "param1": "value1",
    ...
    /*
    "param2": "value2",
    ...
    */
  }
}
```

Note

Since Kea 2.7.5 it is possible to specify extra HTTP headers which are added to HTTP responses. Each header is specified by its name and value with optionally a user context. For instance:

```
{
  "http-headers": [
    {
      "name": "Strict-Transport-Security",
      "value": "max-age=31536000"
    }
  ],
  ...
}
```

adds a HSTS header declaring that HTTPS (vs HTTP) must be used for one year.

19.2 Control Agent Command Response Format

When sending commands via the Control Agent, it is possible to specify multiple services at which the command is targeted. CA forwards this command to each service individually. Thus, the CA response to the controlling client is always wrapped in an array (JSON list) of individual responses. For example, the response for a command sent to one service would be structured as follows:

```
[
  {
    "result": 0, // 0|1|2|3|4
    "text": "textual description",
    "arguments": {
      "argument1": "value1",
      "argument2": "value2",
      ...
    }
  }
]
```

If the command is sent to more than one service, the array would contain responses from each service, in the order they were requested:

```
[
  {
    "result": 0, // 0|1|2|3|4
    "text": "textual description",
    "arguments": {
      "argument1": "value1",
      "argument2": "value2",
      ...
    }
  },
  {
    "result": 0, // 0|1|2|3|4
    "text": "textual description",
```

(continues on next page)

(continued from previous page)

```

    "arguments": {
        "argument1": "value1",
        "argument2": "value2",
        ...
    }
},
...
]

```

An exception to this are authentication or authorization errors which cause CA to reject the command entirely. The response to such an error is formatted as a single entry (JSON map) as follows:

```

{
    "result": 403,
    "text": "Forbidden"
}

```

These types of errors are possible on systems configured for either basic authentication or agents that load *libdhcp_rbac.so*.

19.3 Using the Control Channel

The easiest way to start interacting with the control API is to use common UNIX/Linux tools such as `socat` and `curl`.

In order to control the given Kea service via a UNIX domain socket, use `socat` in interactive mode as follows:

```
$ socat UNIX:/path/to/the/kea/socket -
```

or in batch mode, include the "ignoreeof" option as shown below to ensure `socat` waits long enough for the server to respond:

```
$ echo "{ some command...}" | socat UNIX:/path/to/the/kea/socket -,ignoreeof
```

where `/path/to/the/kea/socket` is the path specified in the `Dhcp4/control-socket/socket-name` parameter in the Kea configuration file. Text passed to `socat` is sent to Kea and the responses received from Kea are printed to standard output. This approach communicates with the specific server directly and bypasses the Control Agent.

It is also easy to open a UNIX socket programmatically. An example of a simple client written in C is available in the Kea Developer's Guide, in the Control Channel Overview chapter, in the [Using Control Channel](#) section.

To use Kea's RESTful API with `curl`, use the following:

```
$ curl -X POST -H "Content-Type: application/json" -d '{ "command": "config-get",
↪ "service": [ "dhcp4" ] }' http://ca.example.org:8000/
```

This assumes that the Control Agent is running on host `ca.example.org` and is running the RESTful service on port 8000.

19.4 Commands Supported by Both the DHCPv4 and DHCPv6 Servers

19.4.1 The build-report Command

The *build-report* command returns on the control channel what the command line *-W* argument displays, i.e. the embedded content of the *config.report* file. This command does not take any parameters.

```
{
  "command": "build-report"
}
```

19.4.2 The config-get Command

The *config-get* command retrieves the current configuration used by the server. This command does not take any parameters. The configuration returned is roughly equal to the configuration that was loaded using the *-c* command-line option during server start-up, or was later set using the *config-set* command. However, there may be certain differences, as comments are not retained. If the original configuration used file inclusion, the returned configuration includes all parameters from all included files. In Kea 2.4.0 and later, the successful response also contains a SHA-256 digest that can be used to easily determine whether a configuration has changed.

Warning

The returned configuration is not redacted, i.e. it contains database passwords in plain text, if those were specified in the original configuration. Care should be taken not to expose the command channel to unprivileged users.

An example command invocation looks like this:

```
{
  "command": "config-get"
}
```

19.4.3 The config-hash-get Command

The *config-hash-get* command retrieves the SHA-256 hash of the current configuration used by the server. This command does not take any parameters. The returned hash can be used to detect configuration changes.

An example command invocation looks like this:

```
{
  "command": "config-hash-get"
}
```

And the server's response:

```
{
  "result": 0,
  "arguments": {
    "hash": "5C3C90EF7035249E2FF74D003C19F34EE0B83A3D329E741B52B2EF95A2C9CC5C"
  }
}
```

In Kea 2.4.0 and later, `config-set` and `config-get` also return the SHA-256 hash of the new or current configuration. This may be used to determine whether a configuration has changed.

19.4.4 The `config-reload` Command

The `config-reload` command instructs Kea to load again the configuration file that was used previously. This operation is useful if the configuration file has been changed by some external source; for example, a system administrator can tweak the configuration file and use this command to force Kea to pick up the changes.

Care should be taken when using this in conjunction with the `config-set` command. Kea remembers the location of the configuration file it was started with, and this configuration can be significantly changed using the `config-set` command. When `config-reload` is issued after `config-set`, Kea attempts to reload its original configuration from the file, possibly losing all changes introduced using `config-set` or other commands.

The `config-reload` command does not take any parameters. An example command invocation looks like this:

```
{
  "command": "config-reload"
}
```

If the configuration file is incorrect, reloading it can raise an error which leaves the server in an unusable state. See *The `config-set` Command* to learn how to recover from a non-working server.

19.4.5 The `config-test` Command

The `config-test` command instructs the server to check whether the new configuration supplied in the command's arguments can be loaded. The supplied configuration is expected to be the full configuration for the target server, along with an optional logger configuration. The configuration is sanity-checked to the extent possible without the server actually attempting to load it; it is possible for a configuration which successfully passes this command to still fail in the `config-set` command or at launch time. The structure of the command is as follows:

```
{
  "command": "config-test",
  "arguments": {
    "<server>": {
    }
  }
}
```

where `<server>` is the configuration element name for a given server, such as "Dhcp4" or "Dhcp6". For example:

```
{
  "command": "config-test",
  "arguments": {
    "Dhcp6": {
      ...
    }
  }
}
```

The server's response contains a numeric code, `result` (0 for success, non-zero on failure), and a string, `text`, describing the outcome:

```
{"result": 0, "text": "Configuration seems sane..." }
```

(continues on next page)

(continued from previous page)

or

```
{"result": 1, "text": "unsupported parameter: BOGUS (<string>:16:26)" }
```

19.4.6 The config-write Command

The *config-write* command instructs the Kea server to write its current configuration to a file on disk. It takes one optional argument, called "filename", that specifies the name of the file to write the configuration to. If not specified, the name used when starting Kea (passed as a *-c* argument) is used. If a relative path is specified, Kea writes its files only in the directory where it is running.

An example command invocation looks like this:

```
{
  "command": "config-write",
  "arguments": {
    "filename": "config-modified-2017-03-15.json"
  }
}
```

Note

As of Kea 2.7.9, the config file may only be written to the same directory as the config file used when starting Kea (passed as a *-c* argument).

19.4.7 The leases-reclaim Command

The *leases-reclaim* command instructs the server to reclaim all expired leases immediately. The command has the following JSON syntax:

```
{
  "command": "leases-reclaim",
  "arguments": {
    "remove": true
  }
}
```

The *remove* boolean parameter is mandatory and indicates whether the reclaimed leases should be removed from the lease database (if *true*), or left in the *expired-reclaimed* state (if *false*). The latter facilitates lease affinity, i.e. the ability to re-assign an expired lease to a returning client that previously used that lease. See *Configuring Lease Affinity* for details. Also, see *Lease Reclamation* for general information about the processing of expired leases (lease reclamation).

19.4.8 The list-commands Command

The *list-commands* command retrieves a list of all commands supported by the server. It does not take any arguments. An example command may look like this:

```
{
  "command": "list-commands",
  "arguments": { }
}
```

The server responds with a list of all supported commands. The arguments element is a list of strings, each of which conveys one supported command.

19.4.9 The config-set Command

The `config-set` command instructs the server to replace its current configuration with the new configuration supplied in the command's arguments. The supplied configuration is expected to be the full configuration for the target server, along with an optional logger configuration. While optional, the logger configuration is highly recommended, as without it the server reverts to its default logging configuration. The structure of the command is as follows:

```
{
  "command": "config-set",
  "arguments": {
    "<server>": {
    }
  }
}
```

where <server> is the configuration element name for a given server, such as "Dhcp4" or "Dhcp6". For example:

```
{
  "command": "config-set",
  "arguments": {
    "Dhcp6": {
      ...
    }
  }
}
```

If the new configuration proves to be invalid, the server retains its current configuration; however, in some cases a fatal error message is logged indicating that the server is no longer providing any service: a working configuration must be loaded as soon as possible. If the control channel is dead, the configuration file can still be reloaded using the SIGHUP signal. If that is unsuccessful, restart the server.

Please note that the new configuration is retained in memory only; if the server is restarted or a configuration reload is triggered via a signal, the server uses the configuration stored in its configuration file. The server's response contains a numeric code, `result` (0 for success, non-zero on failure), and a string, `text`, describing the outcome:

```
{"result": 0, "text": "Configuration successful." }

or

{"result": 1, "text": "unsupported parameter: BOGUS (<string>:16:26)" }
```

In Kea 2.4.0 and later, the successful response from a DHCPv4, DHCPv6, or DHCP-DDNS daemon also contains a SHA-256 digest of the newly set configuration. The digest can be used to easily determine whether a configuration has changed.

19.4.10 The shutdown Command

The `shutdown` command instructs the server to initiate its shutdown procedure; it is the equivalent of sending a SIGTERM signal to the process. This command does not take any arguments. An example command may look like this:

```
{
  "command": "shutdown",
  "arguments": {
    "exit-value": 3
  }
}
```

The server responds with a confirmation that the shutdown procedure has been initiated. The optional parameter, `exit-value`, specifies the numeric value with which the server process exits to the system. The default value is zero.

The DDNS daemon supports an extra parameter, `type`, which controls the way the process cleans up on exit. The supported shutdown types are:

- "normal" - stops the queue manager and finishes all current transactions before exiting. This is the default.
- "drain_first" - stops the queue manager but continues processing requests from the queue until it is empty.
- "now" - exits immediately.

An example command may look like this:

```
{
  "command": "shutdown",
  "arguments": {
    "exit-value": 3,
    "type": "drain_first"
  }
}
```

19.4.11 The `dhcp-disable` Command

The `dhcp-disable` command globally disables the DHCP service. The server continues to operate, but it drops all received DHCP messages. This command is useful when the server's maintenance requires that the server temporarily stop allocating new leases and renew existing leases. It is also useful in failover-like configurations during a synchronization of the lease databases at startup, or recovery after a failure. The optional parameter `max-period` specifies the time in seconds after which the DHCP service should be automatically re-enabled, if the `dhcp-enable` command is not sent before this time elapses.

Since Kea 1.9.4, there is an additional `origin` parameter that specifies the command source. A server administrator should typically omit this parameter because the default value "user" indicates that the administrator sent the command. In Kea 2.5.5 through 2.5.7, this parameter was also used in communication between the HA partners to specify the identifier of an HA service sending the command in a numeric format. However, due to compatibility issues with older Kea versions that did not properly parse numeric values, it was necessary to introduce the new parameter, `origin-id`, in Kea 2.5.8.

It holds a numeric value representing the origin of the command. The same value can still be passed using the `origin` parameter, but it can cause the aforementioned compatibility issues between HA partners running different Kea versions; if both are used, `origin-id` takes precedence. New Kea versions favor using `origin-id` in communication between the HA partners, but overall, it is recommended that both parameters be omitted and the default value used.

The following codes represent the supported origins in numeric format:

- 1 - a user command; the same as specifying "origin": "user".
- 2000, 2001, 2002, etc. - origins specified by HA partners where the increments above 2000 are distinct HA service identifiers used when the partners have many relationships.

In the following example:

```
{
  "command": "dhcp-disable",
  "arguments": {
    "max-period": 20,
    "origin-id": 2002,
    "origin": "user"
  }
}
```

the effective origin is 2002, which indicates it is an HA partner sending the command for the service with ID of 2. The `origin` parameter will be ignored in this case.

19.4.12 The `dhcp-enable` Command

The `dhcp-enable` command globally enables the DHCP service.

Since Kea 1.9.4, there is an additional `origin` parameter that specifies the command source. A server administrator should typically omit this parameter because the default value "user" indicates that the administrator sent the command. In Kea 2.5.5 through 2.5.7, this parameter was also used in communication between the HA partners to specify the identifier of an HA service sending the command in a numeric format. However, due to compatibility issues with older Kea versions that did not properly parse numeric values, it was necessary to introduce the new parameter, `origin-id`, in Kea 2.5.8.

It holds a numeric value representing the origin of the command. The same value can still be passed using the `origin` parameter, but it can cause the aforementioned compatibility issues between HA partners running different Kea versions.; if both are used, `origin-id` takes precedence. New Kea versions favor using `origin-id` in communication between the HA partners, but overall, it is recommended that both

The following codes represent the supported origins in numeric format:

- 1 - a user command; the same as specifying "origin": "user".
- 2000, 2001, 2002, etc. - origins specified by HA partners where the increments above 2000 are distinct HA service identifiers used when the partners have many relationships.

In the following example:

```
{
  "command": "dhcp-enable",
  "arguments": {
    "origin-id": 2002,
    "origin": "user"
  }
}
```

the effective origin is 2002, which indicates it is an HA partner sending the command for the service with ID of 2. The `origin` parameter will be ignored in this case.

19.4.13 The `status-get` Command

The `status-get` command returns the server's runtime information:

- `pid`: the process ID.
- `uptime`: the number of seconds since the start of the server.
- `reload`: the number of seconds since the last configuration (re)load.

- **high-availability:** HA-specific status information about the DHCP servers configured to use the HA hook library:
 - **local:** the state, the role (primary, secondary, ...), and the scopes (i.e. what the server is actually processing) of the local server.
 - **remote:** the remote server's last known state, its served HA scopes, and the role of the remote server in the HA relationship.
- **multi-threading-enabled:** a flag indicating whether multi-threading is enabled.
- **thread-pool-size:** the number of DHCP service threads.
- **packet-queue-size:** the maximum size of the packet queue. There is one queue, regardless of the number of running threads.
- **packet-queue-statistics:** the average queue size for the last 10, 100, and 1000 packets, using an approach similar to the UNIX `top` command. The average queue size for the last 10 packets can be considered an instantaneous value, while the average for the last 1000 packets shows a longer-term trend.

The **high-availability** information is returned only when the command is sent to the DHCP servers in an HA setup. This parameter is never returned when the `status-get` command is sent to the Control Agent or DDNS daemon.

The **thread-pool-size**, **packet-queue-size** and **packet-queue-statistics** parameters are returned only when the command is sent to DHCP servers with multi-threading enabled. These three parameters and **multi-threading-enabled** are never returned when the `status-get` command is sent to the Control Agent or DDNS daemon.

To learn more about the HA status information returned by the `status-get` command, please refer to the *The status-get Command* section.

19.4.14 The server-tag-get Command:

The `server-tag-get` command returns the configured server tag of the DHCPv4 or DHCPv6 server (*Configuration Sharing and Server Tags* explains the server tag concept).

19.4.15 The config-backend-pull Command:

The `config-backend-pull` command triggers the polling of configuration backends (which must be configured for this command to have an effect), explained in *Enabling the Configuration Backend*.

19.4.16 The version-get Command

The `version-get` command returns extended information about the Kea version; it is the same information available via the `-V` command-line argument. This command does not take any parameters.

```
{  
  "command": "version-get"  
}
```

19.5 Commands Supported by the DHCPv4 Server

19.5.1 The subnet4-select-test Command

The `subnet4-select-test` provides a way to test DHCPv4 subnet selection based on a set of input parameters typically supplied in a client packet. Recognized parameters take strings and are:

- **interface** - the incoming interface name

- `address` - the client address
- `relay` - the relay/gateway address
- `local` - the local/destination address
- `remote` - the remote/source address
- `link` - the RAI link-selection address
- `subnet` - the subnet-selection address
- `classes` - (list of strings) client classes (allowing to select a guarded subnet)

The RAI link-selection is ignored when the `ignore-rai-link-selection` compatibility flag is `true`. When it is not ignored it has precedence over the subnet-selection.

Outside of errors possible results are:

- (empty) "no selected subnet"
- "selected shared network '<name>' starting with subnet '<subnet>' id <id>"
- "selected subnet '<subnet>' id <id>"

19.5.2 The `subnet4o6-select-test` Command

The `subnet4o6-select-test` provides a way to test DHCPv4-over-DHCPv6 subnet selection based on a set of input parameters typically supplied in a client packet. Recognized parameters take strings and are:

- `interface` - the incoming interface name of the DHCPv6 server
- `interface-id` - the binary content of the interface-id relay option
- `address` - the client address
- `relay` - the relay/gateway address
- `local` - the local/destination address
- `remote` - the remote/source IPv6 address of the DHCPv6 server
- `link` - the first relay link IPv6 address
- `subnet` - the subnet-selection address
- `classes` - (list of strings) client classes (allowing to select a guarded subnet)

According to the code only `remote`, `interface-id` and `interface` selectors are used. In DHCPv4-over-DHCPv6 implementation `interface` and `remote` values are transmitted from the DHCPv6 server, `interface-id` and `link` are carried in the relay info part of the DHCPv6 packet so are the same as for the DHCPv6 server.

Outside of errors possible results are:

- (empty) "no selected subnet"
- "selected shared network '<name>' starting with subnet '<subnet>' id <id>"
- "selected subnet '<subnet>' id <id>"

19.6 Commands Supported by the DHCPv6 Server

19.6.1 The `subnet6-select-test` Command

The `subnet6-select-test` provides a way to test DHCPv6 subnet selection based on a set of input parameters typically supplied in a client packet. Recognized parameters take strings and are:

- `interface` - the incoming interface name
- `interface-id` - the binary content of the interface-id relay option
- `remote` - the remote/source address
- `link` - the first relay link address
- `classes` - (list of strings) client classes (allowing to select a guarded subnet)

Outside of errors possible results are:

- (empty) "no selected subnet"
- "selected shared network '<name>' starting with subnet '<subnet>' id <id>"
- "selected subnet '<subnet>' id <id>"

19.7 Commands Supported by the D2 Server

The D2 server supports only a subset of the DHCPv4/DHCPv6 server commands:

- `build-report`
- `config-get`
- `config-hash-get`
- `config-reload`
- `config-set`
- `config-test`
- `config-write`
- `list-commands`
- `shutdown`
- `status-get`
- `version-get`

19.8 Commands Supported by the Control Agent

The following commands, listed in *Commands Supported by Both the DHCPv4 and DHCPv6 Servers*, are also supported by the Control Agent; when the `service` parameter is blank, the commands are handled by the CA and they relate to the CA process itself:

- `build-report`
- `config-get`
- `config-hash-get`
- `config-reload`
- `config-set`
- `config-test`
- `config-write`
- `list-commands`
- `shutdown`

- `status-get`
- `version-get`

19.9 Migration from the Control Agent

Since Kea version 2.7.2 DHCP servers support HTTP/HTTPS control channels so the Control Agent (CA) is no longer needed.

The DHCPv4, DHCPv6, and D2 servers extend the `control-socket` entry to `control-sockets` list. To migrate a CA configuration add an element to this list with:

- `socket-type` set to `http` or `https`
- `socket-address` with the content of CA `http-host`
- `socket-port` with the content of CA `http-port`
- `trust-anchor` (unchanged)
- `cert-file` (unchanged)
- `key-file` (unchanged)
- `cert-required` (unchanged)
- `authentication` (unchanged)

User context is supported too. Please look at respective HTTP control socket sections for defaults and other details (beware that two servers must use different address / port pairs): [HTTP/HTTPS Control Socket](#), [HTTP/HTTPS Control Socket](#) and [HTTP/HTTPS Control Socket](#)

For compatibility the JSON result of these HTTP/HTTPS control sockets is still encapsulated into a list.

LOGGING

20.1 Logging Configuration

During its operation Kea may produce many log messages. They differ in severity (some are more important than others) and source (different components, like hooks, produce different messages). It is useful to understand which log messages are critical and which are not, and to configure logging appropriately. For example, debug-level messages can be safely ignored in a typical deployment. They are, however, very useful when debugging a problem.

The logging system in Kea is configured through the `loggers` entry in the server section of the configuration file.

20.1.1 Loggers

Within Kea, a message is logged through an entity called a "logger." Different components log messages through different loggers, and each logger can be configured independently of the others. Some components, in particular the DHCP server processes, may use multiple loggers to log messages pertaining to different logical functions of the component. For example, the DHCPv4 server uses one logger for messages about packet reception and transmission, another logger for messages related to lease allocation, and so on. Some of the libraries used by the Kea server, such as `libdhcpd`, use their own loggers.

Users implementing hook libraries (code attached to the server at runtime) are responsible for creating the loggers used by those libraries. Such loggers should have unique names, different from the logger names used by Kea. That way, the messages produced by the hook library can be distinguished from messages issued by the core Kea code. Unique names also allow the hook loggers to be configured independently of loggers used by Kea. Whenever it makes sense, a hook library can use multiple loggers to log messages pertaining to different logical parts of the library.

In the server section of a configuration file, the configuration for zero or more loggers (including loggers used by the proprietary hook libraries) can be specified. If there are no loggers specified, the code uses default values; these cause Kea to log messages of INFO severity or greater to standard output. There is a small time window after Kea has been started but before it has read its configuration; logging in this short period can be controlled using environment variables. For details, see *Logging During Kea Startup*.

The three main elements of a logger configuration are: `name` (the component that is generating the messages), `severity` (what to log), and `output_commands` (where to log). There is also a `debuglevel` element, which is only relevant if debug-level logging has been selected.

20.1.1.1 The `name` (string) Logger

Each logger in the system has a name: that of the component binary file using it to log messages. For instance, to configure logging for the DHCPv4 server, add an entry for a logger named "kea-dhcp4". This configuration will then be used by the loggers in the DHCPv4 server and all the libraries used by it, unless a library defines its own logger and there is a specific logger configuration that applies to that logger.

When tracking down an issue with the server's operation, use of DEBUG logging is required to obtain the verbose output needed for problem diagnosis. However, the high verbosity is likely to overwhelm the logging system in cases where the server is processing high-volume traffic. To mitigate this problem, Kea can use multiple loggers, for different

functional parts of the server, that can each be configured independently. If the user is reasonably confident that a problem originates in a specific function of the server, or that the problem is related to a specific type of operation, they may enable high verbosity only for the relevant logger, thereby limiting the DEBUG messages to the required minimum.

The loggers are associated with a particular library or binary of Kea. However, each library or binary may (and usually does) include multiple loggers. For example, the DHCPv4 server binary contains separate loggers for packet parsing, dropped packets, callouts, etc.

The loggers form a hierarchy. For each program in Kea, there is a "root" logger, named after the program (e.g. the root logger for *kea-dhcp4*, the DHCPv4 server, is named *kea-dhcp4*). All other loggers are children of this logger and are named accordingly, e.g. the allocation engine in the DHCPv4 server logs messages using a logger called *kea-dhcp4.alloc-engine*.

This relationship is important, as each child logger derives its default configuration from its parent root logger. In the typical case, the root logger configuration is the only logging configuration specified in the configuration file and so applies to all loggers. If an entry is made for a given logger, any attributes specified override those of the root logger, whereas any not specified are inherited from it.

To illustrate this, suppose we are using the DHCPv4 server with the root logger *kea-dhcp4* logging at the INFO level. In order to enable DEBUG verbosity for DHCPv4 packet drops, we must create a configuration entry for the logger with "name": "*kea-dhcp4.bad-packets*", "severity": "DEBUG", and an explicit debug level. All other configuration parameters may be omitted for this logger if it should use the default values specified in the root logger's configuration.

`debuglevel` is inherited only if `severity` is missing as well. For predictable results, if `severity` is "DEBUG", these two attributes should always be explicitly specified or omitted together. An entry with an explicit "DEBUG" severity does not inherit `debuglevel` from the root logger and defaults to 0 if missing, resulting in no debug messages being logged. This is a consequence of relying on the log4cplus inheritance mechanism.

If there are multiple logger specifications in the configuration that might match a particular logger, the specification with the more specific logger name takes precedence. For example, if there are entries for both *kea-dhcp4* and *kea-dhcp4.dhcpd*, the main DHCPv4 server program — and all libraries it uses other than the *dhcpd* library (*libdhcpd*) — logs messages according to the configuration in the first entry (*kea-dhcp4*). Messages generated by the *dhcpd* library are logged according to the configuration set by the second entry.

The currently defined loggers are listed in the following table. The "Software Package" column of this table specifies whether the particular loggers belong to the core Kea code (open source Kea binaries and libraries), or hook libraries (open source or subscriber).

Table 1: List of loggers supported by Kea servers and hook libraries shipped with Kea/subscriber packages

Logger Name	Software Package	Description
<i>kea-ctrl-agent</i>	core	The root logger for the Control Agent exposing the RESTful control API. All components used by the Control Agent inherit the settings from this logger.
<i>kea-ctrl-agent.auth</i>	core	A logger which covers access control details, such as the result of basic HTTP authentication.
<i>kea-ctrl-agent.ctrl-agent</i>	core	Used to log results of configuration checks, information about services starting or failing to start, command reception, and forwarding.
<i>kea-ctrl-agent.http</i>	core	A logger which outputs log messages related to receiving, parsing, and sending HTTP messages.
<i>kea-ctrl-agent.rbac-hooks</i>	<i>libdhcp_rbac.so</i> subscriber hook library	Used to log messages related to the operation of the RBAC hook library.

continues on next page

Table 1 – continued from previous page

Logger Name	Software Package	Description
kea-dhcp4	core	The root logger for the DHCPv4 server. All components used by the DHCPv4 server inherit the settings from this logger.
kea-dhcp6	core	The root logger for the DHCPv6 server. All components used by the DHCPv6 server inherit the settings from this logger.
kea-dhcp4. alloc-engine, kea-dhcp6. alloc-engine	core	Used by the lease allocation engine, which is responsible for managing leases in the lease database, i.e. creating, modifying, and removing DHCP leases as a result of processing messages from clients.
kea-ctrl-agent. auth, kea-dhcp4. auth, kea-dhcp6. auth	core	Used to log malformed HTTP packets when using basic authentication.
kea-dhcp4. bad-packets, kea-dhcp6. bad-packets	core	Used by the DHCP servers for logging inbound client packets that were dropped or to which the server responded with a DHCPNAK. It allows administrators to configure a separate log output that contains only packet drop and reject entries.
kea-dhcp4. bootp-hooks	libdhcp_bootp.so open-source hook library	This logger is used to log messages related to the operation of the BOOTP hook library.
kea-ctrl-agent. callouts, kea-dhcp4. callouts, kea-dhcp6. callouts, kea-dhcp-ddns. callouts	core	Used to log messages pertaining to the callouts registration and execution for a particular hook point.
kea-dhcp4. cb-cmds-hooks, kea-dhcp6. cb-cmds-hooks	libdhcp_cb_cmds.so subscriber hook library	Used to log messages related to the operation of the Config Backend Commands hook library.
kea-dhcp4. class-cmds-hooks, kea-dhcp6. class-cmds-hooks	libdhcp_class_cmds.so open-source hook library	Used to log messages related to the operation of the Class Commands hook library.
kea-dhcp4. commands, kea-dhcp6. commands	core	Used to log messages related to the handling of commands received by the DHCP server over the command channel.
kea-dhcp4. database, kea-dhcp6. database	core	Used to log messages related to general operations of the relational databases.
kea-dhcp4. ddns-tuning-hooks, kea-dhcp6. ddns-tuning-hooks	libdhcp_ddns_tuning.so open-source hook library	Used to log messages related to the operation of the DDNS Tuning hook library.
kea-dhcp4.ddns, kea-dhcp6.ddns	core	Used by the DHCP server to log messages related to Client FQDN and Hostname option processing. It also includes log messages related to the relevant DNS updates.

continues on next page

Table 1 – continued from previous page

Logger Name	Software Package	Description
kea-dhcp4.dhcp4, kea-dhcp6.dhcp6	core	Used to log basic operations.
kea-dhcp4.dhcpdsvr, kea-dhcp6.dhcpdsvr	core	The base loggers for the libkea-dhcpdsvr.so library.
kea-dhcp4.eval, kea-dhcp6.eval	core	Used to log messages related to the client classification expression evaluation code.
kea-dhcp4. flex-id-hooks, kea-dhcp6. flex-id-hooks	libdhcp_flex_id.so open-source hook library	Used to log messages related to the operation of the Flexible Identifier hook library.
kea-dhcp4. flex-option-hooks, kea-dhcp6. flex-option-hooks	libdhcp_flex_option.so open-source hook library	Used to log messages related to the operation of the Flexible Option hook library.
kea-dhcp4. ha-hooks, kea-dhcp6. ha-hooks	libdhcp_ha.so open-source hook library	Used to log messages related to the operation of the High Availability hook library.
kea-ctrl-agent. hooks, kea-dhcp4.hooks, kea-dhcp6.hooks, kea-dhcp-ddns. hooks	core	Used to log messages related to the management of hook libraries, e.g. registration and deregistration of the libraries, and to the initialization of the callouts execution for various hook points within the DHCP server.
kea-dhcp4. host-cache-hooks, kea-dhcp6. host-cache-hooks	libdhcp_host_cache.so open-source hook library	Used to log messages related to the operation of the Host Cache hook library.
kea-dhcp4. host-cmds-hooks, kea-dhcp6. host-cmds-hooks	libdhcp_host_cmds.so open-source hook library	Used to log messages related to the operation of the Host Commands hook library. In general, these pertain to the loading and unloading of the library and the execution of commands by the library.
kea-dhcp4.hosts, kea-dhcp6.hosts	core	Used within the libkea-dhcpdsvr.so to log messages related to the management of DHCP host reservations, i.e. retrieving reservations and adding new reservations.
kea-dhcp4. lease-cmds-hooks, kea-dhcp6. lease-cmds-hooks	libdhcp_lease_cmds.so open-source hook library	Used to log messages related to the operation of the Lease Commands hook library. In general, these pertain to the loading and unloading of the library and the execution of commands by the library.
kea-dhcp4. limits-hooks, kea-dhcp6. limits-hooks	libdhcp_limits.so open-source hook library	Used to log messages related to the operation of the Limits hook library.
kea-dhcp4. lease-query-hooks, kea-dhcp6. lease-query-hooks	libdhcp_lease_query.so open-source hook library	Used to log messages related to the operation of the Leasequery hook library.
kea-dhcp4.leases, kea-dhcp6.leases	core	Used by the DHCP server to log messages related to lease allocation. The messages include detailed information about the allocated or offered leases, errors during the lease allocation, etc.

continues on next page

Table 1 – continued from previous page

Logger Name	Software Package	Description
kea-dhcp4. legal-log-hooks, kea-dhcp6. legal-log-hooks	<i>libdhcp_legal_log.so</i> open-source hook library	Used to log messages related to the operation of the Forensic Logging hook library.
kea-dhcp4. mysql-cb-hooks, kea-dhcp6. mysql-cb-hooks	<i>libdhcp_mysql.so</i> open-source hook library	Used to log messages related to the operation of the MySQL Configuration Backend hook library.
kea-dhcp4. mysql-hb-hooks, kea-dhcp6. mysql-hb-hooks	<i>libdhcp_mysql.so</i> open-source hook library	Used to log messages related to the operation of the MySQL Host Backend hook library.
kea-dhcp4. mysql-lb-hooks, kea-dhcp6. mysql-lb-hooks	<i>libdhcp_mysql.so</i> open-source hook library	Used to log messages related to the operation of the MySQL Lease Backend hook library.
kea-dhcp4.options, kea-dhcp6.options	core	Used by the DHCP server to log messages related to the processing of options in the DHCP messages, i.e. parsing options, encoding options into on-wire format, and packet classification using options contained in the received packets.
kea-dhcp4.packets, kea-dhcp6.packets	core	Mostly used to log messages related to transmission of DHCP packets, i.e. packet reception and the sending of a response. Such messages include information about the source and destination IP addresses and interfaces used to transmit packets. The logger is also used to log messages related to subnet selection, as this selection is usually based on the IP addresses, relay addresses, and/or interface names, which can be retrieved from the received packet even before the DHCP message carried in the packet is parsed.
kea-dhcp4. perfmon-hooks, kea-dhcp6. perfmon-hooks	<i>libdhcp_perfmon.so</i> open-source hook library	Used to log messages related to the operation of the Performance Monitoring hook library.
kea-dhcp4. ping-check-hooks	<i>libdhcp_ping_check.so</i> open-source hook library	Used to log messages related to the operation of the Ping Check hook library.
kea-dhcp4. pgsql-cb-hooks, kea-dhcp6. pgsql-cb-hooks	<i>libdhcp_pgsql.so</i> open-source hook library	Used to log messages related to the operation of the PostgreSQL Configuration Backend hook library.
kea-dhcp4. pgsql-hb-hooks, kea-dhcp6. pgsql-hb-hooks	<i>libdhcp_pgsql.so</i> open-source hook library	Used to log messages related to the operation of the PostgreSQL Host Backend hook library.
kea-dhcp4. pgsql-lb-hooks, kea-dhcp6. pgsql-lb-hooks	<i>libdhcp_pgsql.so</i> open-source hook library	Used to log messages related to the operation of the PostgreSQL Lease Backend hook library.

continues on next page

Table 1 – continued from previous page

Logger Name	Software Package	Description
kea-dhcp4. run-script-hooks, kea-dhcp6. run-script-hooks	<i>libdhcp_run_script.so</i> open-source hook library	Used to log messages related to the operation of the Run Script hook library.
kea-dhcp4. radius-hooks, kea-dhcp6. radius-hooks	<i>libdhcp_radius.so</i> open-source hook library	Used to log messages related to the operation of the RADIUS hook library.
kea-dhcp4. stat-cmds-hooks, kea-dhcp6. stat-cmds-hooks	<i>libdhcp_stat_cmds.so</i> open-source hook library	Used to log messages related to the operation of the Statistics Commands hook library. In general, these pertain to loading and unloading the library and the execution of commands by the library.
kea-dhcp4. subnet-cmds-hooks, kea-dhcp6. subnet-cmds-hooks	<i>libdhcp_subnet_cmds.so</i> open-source hook library	Used to log messages related to the operation of the Subnet Commands hook library. In general, these pertain to loading and unloading the library and the execution of commands by the library.
kea-dhcp4.tcp, kea-dhcp6.tcp	core	Used to log messages related to TCP traffic.
kea-dhcp4. user_chk, kea-dhcp6. user_chk	<i>libdhcp_user_chk.so</i> open-source hook library	Used to log messages related to the operation of the User Check hook library.
kea-dhcp-ddns	core	The root logger for the <i>kea-dhcp-ddns</i> daemon. All components used by this daemon inherit the settings from this logger unless there are configurations for more specialized loggers.
kea-ctrl-agent. dctl, kea-dhcp-ddns. dctl	core	Used to log basic information about the process, received signals, and triggered reconfigurations.
kea-dhcp-ddns. asiodns	core	Used to log messages about network events in DDNS operations.
kea-dhcp-ddns. dhcpddns	core	Used by the <i>kea-dhcp-ddns</i> daemon to log events related to DDNS operations.
kea-dhcp-ddns. dhcp-to-d2	core	Used by the <i>kea-dhcp-ddns</i> daemon to log information about events dealing with receiving messages from the DHCP servers and adding them to the queue for processing.
kea-dhcp-ddns. d2-to-dns	core	Used by the <i>kea-dhcp-ddns</i> daemon to log information about events dealing with sending and receiving messages to and from the DNS servers.
kea-dhcp-ddns. gss-tsig-hooks	<i>libddns_gss_tsig.so</i> open-source hook library	Used to log messages related to the operation of the GSS-TSIG hook library.
kea-dhcp-ddns. libdhcp-ddns	core	Used to log events related to DDNS operations.
kea-netconf	core	The root logger for the NETCONF agent. All components used by NETCONF inherit the settings from this logger if there is no specialized logger provided.

Note that user-defined hook libraries should not use any of the loggers mentioned above, but should instead define new loggers with names that correspond to the libraries using them. Suppose that a user created a library called “libdhcp-packet-capture” to dump packets received and transmitted by the server to a file. An appropriate name for the logger

could be `kea-dhcp4.packet-capture-hooks`. (Note that the hook library implementer only specifies the second part of this name, i.e. “packet-capture”. The first part is a root-logger name and is prepended by the Kea logging system.) It is also important to note that since this new logger is a child of a root logger, it inherits the configuration from the root logger, something that can be overridden by an entry in the configuration file.

The easiest way to find a logger name is to configure all logging to go to a single destination and look there for specific logger names. See *Logging Message Format* for details.

20.1.1.2 The severity (string) Logger

This specifies the category of messages logged. Each message is logged with an associated severity, which may be one of the following (in descending order of severity):

- **FATAL** - associated with messages generated by a condition that is so serious that the server cannot continue executing.
- **ERROR** - associated with messages generated by an error condition. The server continues executing, but the results may not be as expected.
- **WARN** - indicates an out-of-the-ordinary condition. However, the server continues executing normally.
- **INFO** - an informational message marking some event.
- **DEBUG** - messages produced for debugging purposes.

When the severity of a logger is set to one of these values, it only logs messages of that severity and above (e.g. setting the logging severity to **INFO** logs **INFO**, **WARN**, **ERROR**, and **FATAL** messages). The severity may also be set to **NONE**, in which case all messages from that logger are inhibited.

Note

The *keactrl* tool, described in *Managing Kea with keactrl*, can be configured to start the servers in verbose mode. If this is the case, the settings of the logging severity in the configuration file have no effect; the servers use a logging severity of **DEBUG** regardless of the logging settings specified in the configuration file. To control severity via the configuration file, please make sure that the `kea_verbose` value is set to “no” within the *keactrl* configuration.

20.1.1.3 The debuglevel (integer) Logger

When a logger's severity is set to **DEBUG**, this value specifies the level of debug messages to be printed. It ranges from 0 (least verbose) to 99 (most verbose). If severity for the logger is not **DEBUG**, this value is ignored.

20.1.1.4 The output-options (list) Logger

Each logger can have zero or more `output-options`. These specify where log messages are sent and are explained in detail below.

Note

The alias `output-options` was added in Kea 2.5.1, to be used interchangeably with the previous `output_options` configuration key. This was done to keep all configuration keys consistent, i.e. using a hyphen (-) instead of an underscore (_) in the key name.

As of Kea 2.5.2, `output-options` becomes the default configuration key and `output_options` can be used as an alias.

As of Kea 2.7.4, `output_options` configuration key has been deprecated and will be removed in future versions. Please consider updating your configuration by switching to using `output-options` instead.

20.1.1.4.1 The `output` (string) Option

This value determines the type of output. There are several special values allowed here: `stdout` (messages are printed on standard output), `stderr` (messages are printed on stderr), `syslog` (messages are logged to syslog using the default name), `syslog:name` (messages are logged to syslog using a specified name). Any other value is interpreted as a filename to which messages should be written.

Note

As of Kea 2.7.9, log files may only be written to the output directory determined during compilation as: "[kea-install-dir]/var/log/kea". This path may be overridden at startup by setting the environment variable `KEA_LOG_FILE_DIR` to the desired path. If a path other than this value is used in `output`, Kea will emit an error and refuse to start or, if already running, log an unrecoverable error. For ease of use simply omit the path component from `output` and specify only the file name.

20.1.1.4.2 The `flush` (boolean) Option

This flushes the buffers after each log message. Doing this reduces performance but ensures that if the program terminates abnormally, all messages up to the point of termination are recorded. The default is `true`.

20.1.1.4.3 The `maxsize` (integer) Option

This option is only relevant when the destination is a file; this is the maximum size in bytes that a log file may reach. When the maximum size is reached, the file is renamed and a new file created. Initially, a ".1" is appended to the name; if a ".1" file exists, it is renamed ".2", etc. This is referred to as rotation.

The default value is 10240000 (10MB). The smallest value that can be specified without disabling rotation is 204800. Any value less than this, including 0, disables rotation. The greatest possible value is `INT_MAX` MB, which is approximately 2PB.

Note

Due to a limitation of the underlying logging library (log4cplus), rolling over the log files (from ".1" to ".2", etc.) may show odd results; there can be multiple small files at the timing of rollover. This can happen when multiple processes try to roll over the files simultaneously. Version 1.1.0 of log4cplus solved this problem, so if this version or later of log4cplus is used to build Kea, the issue should not occur. Even with older versions, it is normally expected to happen rarely unless the log messages are produced very frequently by multiple different processes.

20.1.1.4.4 The `maxver` (integer) Option

This option is only relevant when the destination is a file and rotation is enabled (i.e. `maxsize` is large enough). This is the maximum number of rotated versions that will be kept. Once that number of files has been reached, the oldest file, "log-name.maxver", is discarded each time the log rotates. In other words, at most there will be the active log file plus `maxver` rotated files. The minimum and default value is 1.

20.1.1.4.5 The `pattern` (string) Option

This option can be used to specify the layout pattern of messages for a logger. Kea logging is implemented using the log4cplus library and its output formatting is based, conceptually, on the `printf` formatting from C; this is discussed in detail in the next section, *Logging Message Format*.

Each output type (`stdout`, `file`, or `syslog`) has a default `pattern` which describes the content of its log messages. This parameter can be used to specify a desired pattern. The pattern for each logger is governed individually, so each

configured logger can have its own pattern. Omitting the `pattern` parameter or setting it to an empty string, "", causes Kea to use the default pattern for that logger's output type.

In addition to the log text itself, the default patterns used for `stdout` and files contain information such as date and time, logger level, and process information. The default pattern for `syslog` is limited primarily to log level, source, and the log text. This avoids duplicating information which is usually supplied by `syslog`.

Warning

Users are strongly encouraged to test their pattern(s) on a local, non-production instance of Kea, running in the foreground and logging to `stdout`.

20.1.2 Logging Message Format

As mentioned above, Kea log message content is controlled via a scheme similar to the C language's `printf` formatting. The "pattern" used for each message is described by a string containing one or more format components as part of a text string. In addition to the components, the string may contain any other useful text for the administrator.

The behavior of Kea's format strings is determined by `log4cplus`. The following time format options are possible enclosed in `%D{}` or `%d{}`:

Table 2: List of supported time format string components by Kea's logger

Component	Value
%a	Abbreviated weekday name
%A	Full weekday name
%b	Abbreviated month name
%B	Full month name
%c	Standard date and time string
%d	Day of month as a decimal(1-31)
%H	Hour(0-23)
%I	Hour(1-12)
%j	Day of year as a decimal(1-366)
%m	Month as decimal(1-12)
%M	Minute as decimal(0-59)
%p	Locale's equivalent of AM or PM
%q	milliseconds as decimal(0-999)
%Q	microseconds as decimal(0-999.999)
%S	Second as decimal(0-59)
%U	Week of year, Sunday being first day(0-53)
%w	Weekday as a decimal(0-6, Sunday being 0)
%W	Week of year, Monday being first day(0-53)
%x	Standard date string
%X	Standard time string
%y	Year in decimal without century(0-99)
%Y	Year including century as decimal
%Z	Time zone name
%%	The percent sign

Refer to the documentation for the `strftime()` function found in the `<ctime>` header or the `strftime(3)` UNIX manual page for more information.

It is probably easiest to understand this by examining the default pattern for `stdout` and files; currently they are the same. That pattern is shown below:

```
"%D{%Y-%m-%d %H:%M:%S.%q} %-5p [%c/%i.%t] %m\n"
```

and a typical log produced by this pattern looks something like this:

```
2019-08-05 14:27:45.871 DEBUG [kea-dhcp4.dhcpsrv/8475.12345] DHCPDRV_TIMERMGR_START_
↳TIMER starting timer: reclaim-expired-leases
```

That breaks down to:

- `%D{%Y-%m-%d %H:%M:%S.%q}` `"%D"` is the local date and time when the log message is generated, while everything between the curly braces, `"{}"`, are date and time components. From the example log above this produces: `2019-08-05 14:27:45.871`
- `%-5p` The severity of the message, output as a minimum of five characters, using right-padding with spaces. In our example log: `DEBUG`
- `%c` The log source. This includes two elements: the Kea process generating the message, in this case, `kea-dhcp4`; and the component within the program from which the message originated, `dhcpsrv` (e.g. the name of the library used by DHCP server implementations).
- `%i` The process ID. From the example log: `8475`.
- `%t` The thread ID. From the example log: `12345`. The format of the thread ID is OS-dependent: e.g. on some systems it is an address, so it is displayed in hexadecimal.
- `%m` The log message itself. Kea log messages all begin with a message identifier followed by arbitrary log text. Every message in Kea has a unique identifier, which can be used as an index to the *Kea Messages Manual*, where more information can be obtained. In our example log above, the identifier is `DHCPDRV_TIMERMGR_START_TIMER`. The log text is typically a brief description detailing the condition that caused the message to be logged. In our example, the information logged, `starting timer: reclaim-expired-leases`, explains that the timer for the expired lease reclamation cycle has been started.

Warning

Omitting `%m` omits the log message text from the output, making it rather useless. `%m` should be considered mandatory.

Finally, note that spacing between components, the square brackets around the log source and PID, and the final carriage return `\n` are all literal text specified as part of the pattern.

Warning

To ensure that each log entry is a separate line, patterns must end with an `\n`. There may be use cases where it is not desired so we do not enforce its inclusion. If it is omitted from the pattern, the log entries will run together in one long "line".

The default pattern for syslog output is:

```
"%-5p [%c.%t] %m\n"
```

It omits the date and time as well as the process ID, as this information is typically output by syslog. Note that Kea uses the pattern to construct the text it sends to syslog (or any other destination). It has no influence on the content syslog may add or formatting it may do.

Consult the OS documentation for syslog behavior, as there are multiple implementations.

A complete list of logging parameters supported by Kea is shown in the table below:

Table 3: List of supported format string components by Kea's logger

Component	Value
%b	Outputs file that called the log e.g., logger_impl.cc
%c	Outputs the logger of the event e.g., kea-dhcp4.hosts
%d	%d{} formats UTC time output e.g., %d{%Y-%m-%d %H:%M:%S.%q}
%D	%D{} formats LOCAL time output e.g., %D{%Y-%m-%d %H:%M:%S.%q}
%E	Outputs environment variables e.g., %E{PATH}
%F	Outputs filename where logging request was issued e.g., logger_impl.cc
%h	Outputs hostname of the system e.g., host-1
%H	Outputs fully qualified domain name e.g., host-1.example.com
%l	Equivalent to %F:%L e.g., logger_impl.cc:179
%L	Outputs the line number where the log was called e.g., 179
%m	Outputs the actual log message
%M	Outputs caller (function name) of the log message
%n	Outputs line separator supported by platform e.g., \n in Linux
%p	Outputs log severity e.g., INFO
%r	Outputs milliseconds since program start e.g., 1235
%t	Outputs thread ID that generated the log message e.g., 281472855306256
%T	Outputs thread name that generated the log message e.g., 168005
%i	Outputs process ID that generated the log message e.g., 168030
%%	Outputs a literal percent sign

Padding and truncation are also possible with modifiers preceding the component; this is done by placing a number and other modifier characters between the component and the % sign. There are five ways of modifying the output:

- %20p Left pads with spaces (align right) if the severity is shorter than 20 characters.
- %-15r Right pads with spaces (align left) if the milliseconds since program start is shorter than 15 characters.
- %.30m Truncates from the beginning of the message if the message is longer than 30 characters.
- %10.35E{PATH} Left pad with spaces (align right) if the environment variable "PATH" is shorter than 10 characters. If the content is longer than 35 characters, then truncate from the beginning of the string.
- %-15.40m Right pad with spaces (align left) if the log message is shorter than 15 characters. If the message is longer than 40 characters, truncate from the beginning.

Supported parameters depend on the liblog4cplus version used to compile Kea, which can be checked by executing `kea-dhcp4 -W | grep -i log4cplus`. This produces a result such as: `LOG4CPLUS_VERSION: 2.0.5`. Consult the documentation in the [log4cplus wiki](#) for further information about the installed version.

20.1.2.1 Example Logger Configurations

In this example, we want to set the server logging to write to the console using standard output.

```
"Server": {
  "loggers": [
    {
      "name": "kea-dhcp4",
      "output-options": [
        {
          "output": "stdout"
        }
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "severity": "WARN"
  }
]
}

```

As a second example, we want to store DEBUG log messages in a file that is at most 2MB and keep up to eight copies of old log files. Once the logfile grows to 2MB, it should be renamed and a new file should be created.

```

"Server": {
  "loggers": [
    {
      "name": "kea-dhcp6",
      "output-options": [
        {
          "output": "/var/log/kea/kea-debug.log",
          "maxver": 8,
          "maxsize": 204800,
          "flush": true,
          "pattern": "%d{%j %H:%M:%S.%q} %c %m\n"
        }
      ],
      "severity": "DEBUG",
      "debuglevel": 99
    }
  ]
}

```

Notice that the above configuration uses a custom pattern which produces output like this:

```

220 13:50:31.783 kea-dhcp4.dhcp4 DHCP4_STARTED Kea DHCPv4 server version 1.6.0-beta2-git_
↳ started

```

20.1.3 Logging During Kea Startup

The logging configuration is specified in the configuration file. However, when Kea starts, the configuration file is not read until partway into the initialization process. Prior to that, the logging settings are set to default values, although it is possible to modify some aspects of the settings by means of environment variables. In the absence of any logging configuration in the configuration file, the settings of the (possibly modified) default configuration persist while the program is running.

The following environment variables can be used to control the behavior of logging during startup:

KEA_LOCKFILE_DIR

Specifies a directory where the logging system should create its lock file. If not specified, it is prefix/var/run/kea, where "prefix" defaults to /usr/local. This variable must not end with a slash. There is one special value: "none", which instructs Kea not to create a lock file at all. This may cause issues if several processes log to the same file.

KEA_LOGGER_DESTINATION

Specifies logging output. There are several special values:

- stdout Log to standard output.

`stderr` Log to standard error.

`syslog[:fac]` Log via syslog. The optional "fac" (which is separated from the word "syslog" by a colon) specifies the facility to be used for the log messages. Unless specified, messages are logged using the facility "local0".

Any other value is treated as a name of the output file. If not otherwise specified, Kea logs to standard output.

20.2 Logging Levels

All Kea servers follow the overall intention to let the user know what is going on while not overloading the logging system with too much information, as that could easily be used as a denial-of-service attack.

Unlike the FATAL, ERROR, WARN and INFO levels, DEBUG has additional parameters. The following list details the basic information that is logged on each level. Sometimes the circumstances determine whether a piece of information is logged on a higher or lower level. For example, if a packet is being dropped due to configured classification, that is an execution of the configured policy and would be logged on debuglevel 15. However, if the packet is dropped due to an exception being thrown, it is much more important, as it may indicate a software bug, serious problems with memory, or database connectivity problems. As such it may be logged on much higher levels, such as WARN or even ERROR.

- 0 - singular messages printed during startup or shutdown of the server.
- 10 - log information about received API commands.
- 15 - information about reasons why a packet was dropped.
- 40 - tracing information, including processing decisions, results of expression evaluations, and more.
- 45 - similar to level 40, but with more details, e.g. the subnet being selected for an incoming packet.
- 50 - evaluations of expressions, status received from hook points, lease processing, packet processing details, including unpacking, packing, sending, etc.
- 55 - includes all details available, including full packet contents with all options printed.

The debug levels apply only to messages logged on DEBUG, and are configured using the `debuglevel` option. See the *The `debuglevel (integer)` Logger* section for details.

THE KEA SHELL

21.1 Overview of the Kea Shell

The Kea Control Agent (CA, see *The Kea Control Agent*) provides a RESTful control interface over HTTP. That API is typically expected to be used by various IPAMs and similar management systems. Nevertheless, there may be cases when an administrator wants to send a command to the CA directly, and the Kea shell provides a way to do this. It is a simple command-line, scripting-friendly, text client that is able to connect to the CA, send it commands with parameters, retrieve the responses, and display them.

As the primary purpose of the Kea shell is as a tool in a scripting environment, it is not interactive. However, by following simple guidelines it can be run manually.

Kea 1.9.0 introduced basic HTTP authentication support.

Note

Kea version 2.7.2 introduces direct support of HTTP/HTTPS control channel by Kea servers (*kea-dhcp4*, *kea-dhcp6* and *kea-dhcp-ddns*). The *kea-shell* can be used with these servers, the only difference is the `--service` argument becomes useless and is ignored by Kea servers. Note the result is encapsulated in a list for *kea-ctrl-agent* compatibility.

21.2 Shell Usage

kea-shell is run as follows:

```
$ kea-shell [--host hostname] [--port number] [--path path] [--auth-user] [--auth-  
↪password] [--auth-password-file] [--timeout seconds] [--service service-name] [command]
```

where:

- `--host hostname` specifies the hostname of the CA. If not specified, "localhost" is used.
- `--port number` specifies the TCP port on which the CA listens. If not specified, 8000 is used.
- `--path path` specifies the path in the URL to connect to. If not specified, an empty path is used. As the CA listens at the empty path, this parameter is useful only with a reverse proxy.
- `--auth-user` specifies the user ID for basic HTTP authentication. If not specified or specified as the empty string, authentication is not used.
- `--auth-password` specifies the password for basic HTTP authentication. If not specified but the user ID is specified, an empty password is used. If used together with `--auth-password-file`, it is disregarded.

- `--auth-password-file` specifies a text file whose first line contains the password to be used for basic HTTP authentication. This flag takes precedence over `--auth-password` flag. Caution is advised when setting file permissions for the file containing the password. Access to the file should be as strict as possible and limited only to the Kea administrators.
- `--timeout seconds` specifies the timeout (in seconds) for the connection. If not given, 10 seconds is used.
- `--service service-name` specifies the target of a command. If not given, the CA is used as the target. This may be used more than once to specify multiple targets.
- `command` specifies the command to be sent. If not specified, the `list-commands` command is used.

Other switches are:

- `-h` - prints a help message.
- `-v` - displays the Kea version.
- `-V` - displays the extended Kea version.

See [TLS Support](#) for new command-line arguments associated with TLS/HTTPS support.

Once started, the shell reads the parameters for the command from standard input, which are expected to be in JSON format. When all have been read, the shell establishes a connection with the CA using HTTP, sends the command, and awaits a response. Once that is received, it is displayed on standard output.

For a list of available commands, see [Management API](#); additional commands may be provided by hook libraries. For a list of all supported commands from the CA, use the `list-commands` command.

The following shows a simple example of usage:

```
$ kea-shell --host 192.0.2.1 --port 8001 --auth-user foo --auth-password-file secret_
↪file --service dhcp4 list-commands
^D
```

After the command line is entered, the program waits for command parameters to be entered. Since `list-commands` does not take any arguments, Ctrl-D (represented in the above example by "^D") indicates end-of-file and terminates the parameter input. The shell then contacts the CA and prints out the list of available commands returned for the service named `dhcp4`.

The Kea shell will likely be most frequently used in scripts; the next example shows a simple scripted execution. It sends the command `config-write` to the CA (the `--service` parameter has not been used), along with the parameters specified in `param.json`. The result will be stored in `result.json`.

```
$ cat param.json
{"filename": "my-config-file.json"}
$ cat param.json | kea-shell --host 192.0.2.1 config-write > result.json
```

When a reverse proxy is used to de-multiplex requests to different servers, the default empty path in the URL is not enough, so the `--path` parameter should be used. For instance, if requests to the `/kea` path are forwarded to the CA this can be used:

```
$ kea-shell --host 192.0.2.1 --port 8001 --path kea ...
```

The Kea shell requires Python to be installed on the system. It has been tested with various versions of Python 3.

Note

Beginning with Kea 2.4.0, `kea-shell` no longer supports Python 2.7.

The Kea shell is intended to serve more as a demonstration of the RESTful interface's capabilities (and, perhaps, an illustration for people interested in integrating their management environments with Kea) than as a serious management client. It is not likely to be significantly expanded in the future; it is, and will remain, a simple tool.

Note

When using this tool with basic HTTP authentication, please keep in mind that command-line arguments are not hidden from local users.

21.3 TLS Support

kea-shell supports HTTPS connections; the TLS/HTTPS support requires Python 3. The additional command-line arguments are:

- `--ca` specifies the file or directory name of the Certification Authority. If not specified, HTTPS is not used.
- `--cert` specifies the file name of the user end-entity public key certificate. If specified, the file name of the user key must also be specified.
- `--key` specifies the file name of the user key file. If specified, the file name of the user certificate must also be specified. Encrypted key files are not supported.

For example, a basic HTTPS request to get a list of commands could look like this:

```
$ kea-shell --host 127.0.0.1 --port 8000 --ca ./kea-ca.crt list-commands
```


INTEGRATION WITH EXTERNAL SYSTEMS

Kea provides optional support for a variety of external systems, such as RADIUS, NETCONF, YANG, and GSS-TSIG. The following sections describe how to compile Kea with those additional capabilities and how to configure them.

22.1 YANG/NETCONF

22.1.1 Overview

The Network Configuration Protocol, or NETCONF, is a network management protocol defined in [RFC 4741](#). It uses the YANG modeling language, defined in [RFC 6020](#), to provide a uniform way of handling the configuration syntax of various networking devices. Kea provides optional support for a YANG/NETCONF interface with the *kea-netconf* agent.

22.1.2 Installing NETCONF

To get its NETCONF capabilities, Kea requires the v2 versions of libyang and Sysrepo. The specific versions that have been thoroughly tested with Kea are:

- libyang v2.1.4
- sysrepo v2.2.12
- libyang-cpp v1.1.0 (ae7d649ea75da081725c119dd553b2ef3121a6f8)
- sysrepo-cpp v1.1.0 (02634174ffc60568301c3d9b9b7cf710cff6a586)

Note

For users who are unable to upgrade to one of the versions of libyang and Sysrepo listed above, these are the oldest versions known to work reliably with current Kea releases:

- libyang v2.0.256 (56d4e07ef1cdeab3eb2e6700247f83ec9148edcc)
- sysrepo v2.1.84
- libyang-cpp v1.1.0 (7824d9a862f2dc1d8ad4f6a90ab6cee9200f7c81)
- sysrepo-cpp v1.1.0 (e66b2f0c53a428eeb743d355cf86fb30e8e491f1)

Note

kea-netconf may be compatible with later versions of libyang and Sysrepo, but only the versions listed above have been thoroughly tested by ISC.

Installing from packages is recommended, if they are provided by the system. If not, users can build from sources following the directions below, which should work on all popular operating systems.

22.1.2.1 Installing libyang From Sources

```
$ git clone https://github.com/CESNET/libyang.git
$ cd libyang
$ git checkout v2.1.4
$ mkdir build
$ cd build
$ cmake ..
$ make
$ make install
```

22.1.2.2 Installing sysrepo From Sources

```
$ git clone https://github.com/sysrepo/sysrepo.git
$ cd sysrepo
$ git checkout v2.2.12
$ mkdir build
$ cd build
$ cmake -DREPO_PATH=/etc/sysrepo ..
$ make
$ make install # without sudo if you're doing development and want to run unit tests
```

22.1.2.3 Installing libyang-cpp From Sources

```
$ git clone https://github.com/CESNET/libyang-cpp.git
$ cd libyang-cpp
$ git checkout ae7d649ea75da081725c119dd553b2ef3121a6f8
$ mkdir build
$ cd build
$ cmake -DBUILD_TESTING=OFF ..
$ make
$ make install
```

22.1.2.4 Installing sysrepo-cpp From Sources

```
$ git clone https://github.com/sysrepo/sysrepo-cpp.git
$ cd sysrepo-cpp
$ git checkout 02634174ffc60568301c3d9b9b7cf710cff6a586
$ mkdir build
$ cd build
$ cmake -DBUILD_TESTING=OFF ..
$ make
$ make install
```

22.1.3 Compiling With NETCONF

1. Obtain the Kea sources.

```
$ git clone gitlab.isc.org/isc-projects/kea.git
$ cd kea
```

2. Set up the build.

```
$ meson setup build -D netconf=enabled -D cpp_std=c++20
```

Note

If any of the libraries are installed in a custom location, the `PKG_CONFIG_PATH` can point to the `.pc` files.

```
PKG_CONFIG_PATH=/opt/libyang/lib/pkgconfig:/opt/libyang-cpp/lib/pkgconfig:/opt/sysrepo/
lib/pkgconfig:/opt/sysrepo-cpp/lib/pkgconfig
```

3. Check `build/config.report` to verify NETCONF support.

```
NETCONF:
YANG:           2.1.4
YANG_PREFIX:    /opt/libyang
YANGCPP:        1.1.0
YANGCPP_PREFIX: /opt/libyang-cpp
SYSREPO:        2.2.12
SYSREPO_PREFIX: /opt/sysrepo
SYSREPOCPP:     1.1.0
SYSREPOCPP_PREFIX: /opt/sysrepo-cpp
```

4. Compile as usual.

```
$ meson compile -C build
```

22.1.4 Quick Sysrepo Overview

This section offers a brief overview of a subset of available functions in Sysrepo. For more complete information, see the [Sysrepo homepage](#).

In YANG, configurations and state data are described in YANG syntax in module files named `<module-name>[@<revision>].yang`

The revision part is optional and follows the YYYY-MM-DD format. An alternate XML syntax YIN is defined but less user-friendly. Top-level modules are named in Kea models (a short version of schema models).

There are two major modules that Kea is able to support: `kea-dhcp4-server` and `kea-dhcp6-server`. While there is an active effort in the DHC working group at IETF to develop a DHCPv6 YANG model, a similar initiative in the past for DHCPv4 failed. Therefore, Kea uses its own dedicated models for DHCPv4 and DHCPv6 but partially supports the IETF model for DHCPv6.

All of the models have extra modules as dependencies, which are also provided. All of the modules can be found in `src/share/yang/modules` in sources and in `share/kea/yang/modules` in the installation directory. This directory is referred to as `share_directory` in the commands below.

To install modules from sources or upgrade them from older revisions, run the following command. In the case of a revision upgrade, YANG data will be migrated automatically to the new module schema.

```
$ ${share_directory}/yang/modules/utils/reinstall.sh
```

However, if there are any issues during the upgrade process, and data can be recreated from a NETCONF client or through other means, Kea modules can be easily uninstalled before installing again via this command:

```
$ ${share_directory}/yang/modules/utils/reinstall.sh -u
```

This script should be able to reinstall Sysrepo. However, the `-s` flag can also be used to specify a path:

```
$ ./build/src/share/yang/modules/utils/reinstall.sh -s /path/to/sysrepo
```

To individually install all modules:

```
$ cd ./src/share/yang/modules
$ sysrepoctl -i ./ietf-dhcpv6-server*.yang
$ sysrepoctl -i ./kea-dhcp4-server*.yang
$ sysrepoctl -i ./kea-dhcp6-server*.yang
...
```

The installation should look similar to the following:

```
$ ./build/src/share/yang/modules/utils/reinstall.sh
[INF] Connection 2 created.
[INF] Module "keatest-module" was installed.
[INF] File "keatest-module@2022-11-30.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 4 created.
[ERR] Module "ietf-interfaces@2018-02-20" already installed.
[INF] No datastore changes to apply.
[INF] Connection 7 created.
[ERR] Module "ietf-dhcpv6-client" is already in sysrepo.
[INF] No datastore changes to apply.
[INF] Connection 9 created.
[ERR] Module "ietf-dhcpv6-relay" is already in sysrepo.
[INF] No datastore changes to apply.
[INF] Connection 11 created.
[ERR] Module "ietf-dhcpv6-server" is already in sysrepo.
[INF] No datastore changes to apply.
[INF] Connection 13 created.
[ERR] Write permission "ietf-yang-types" check failed.
[INF] No datastore changes to apply.
[INF] Connection 15 created.
[ERR] Module "ietf-dhcpv6-options" is already in sysrepo.
[INF] No datastore changes to apply.
[INF] Connection 17 created.
[ERR] Module "ietf-dhcpv6-types" is already in sysrepo.
[INF] No datastore changes to apply.
[INF] Connection 21 created.
[INF] Module "kea-types" was installed.
[INF] File "kea-types@2019-08-12.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 23 created.
[INF] Module "kea-dhcp-types" was installed.
[INF] File "kea-dhcp-types@2022-11-30.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 25 created.
```

(continues on next page)

(continued from previous page)

```
[INF] Module "kea-dhcp-ddns" was installed.
[INF] File "kea-dhcp-ddns@2022-07-27.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 27 created.
[INF] Module "kea-ctrl-agent" was installed.
[INF] File "kea-ctrl-agent@2019-08-12.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 29 created.
[INF] Module "kea-dhcp4-server" was installed.
[INF] File "kea-dhcp4-server@2022-11-30.yang" was installed.
[INF] No datastore changes to apply.
[INF] Connection 31 created.
[INF] Module "kea-dhcp6-server" was installed.
[INF] File "kea-dhcp6-server@2022-11-30.yang" was installed.
[INF] No datastore changes to apply.
```

To confirm whether the modules have been imported correctly, check the list of currently installed YANG modules. It should be similar to this:

```
$ sysrepoctl -l
Sysrepo repository: /etc/sysrepo

Module Name          | Revision   | Flags | Owner   | Startup Perms | Submodules
--|-----
-----
ietf-datastores       | 2018-02-14 | I     | user:user | 444            |
ietf-dhcpv6-client    | 2018-09-04 | I     | user:user | 600            |
ietf-dhcpv6-options   | 2018-09-04 | I     | user:user | 600            |
ietf-dhcpv6-relay     | 2018-09-04 | I     | user:user | 600            |
ietf-dhcpv6-server    | 2018-09-04 | I     | user:user | 600            |
ietf-dhcpv6-types     | 2018-09-04 | I     | user:user | 600            |
ietf-inet-types       | 2013-07-15 | I     | user:user | 444            |
ietf-interfaces       | 2018-02-20 | I     | user:user | 600            |
ietf-netconf          | 2013-09-29 | I     | user:user | 644            |
ietf-netconf-acm      | 2018-02-14 | I     | user:user | 600            |
ietf-netconf-notifications | 2012-02-06 | I     | user:user | 644            |
ietf-netconf-with-defaults | 2011-06-01 | I     | user:user | 444            |
ietf-origin           | 2018-02-14 | I     | user:user | 444            |
```

(continues on next page)

(continued from previous page)

ietf-yang-library	2019-01-04	I	user:user	644		U
↪						
ietf-yang-metadata	2016-08-05	i				U
↪						
ietf-yang-schema-mount	2019-01-14	I	user:user	644		U
↪						
ietf-yang-types	2013-07-15	I	user:user	444		U
↪						
kea-ctrl-agent	2019-08-12	I	user:user	600		U
↪						
kea-dhcp-ddns	2022-07-27	I	user:user	600		U
↪						
kea-dhcp-types	2022-11-30	I	user:user	600		U
↪						
kea-dhcp4-server	2022-11-30	I	user:user	600		U
↪						
kea-dhcp6-server	2022-11-30	I	user:user	600		U
↪						
kea-types	2019-08-12	I	user:user	600		U
↪						
keatest-module	2022-11-30	I	user:user	600		U
↪						
sysrepo-monitoring	2022-04-08	I	user:user	600		U
↪						
sysrepo-plugind	2022-03-10	I	user:user	644		U
↪						
yang	2022-06-16	I	user:user	444		U
↪						
Flags meaning: I - Installed/i - Imported; R - Replay support						

To reinstall a module, if the revision YANG entry was bumped, simply installing it will update it automatically. Otherwise, it must first be uninstalled:

```
$ sysrepoctl -u kea-dhcp4-server
```

If the module is used (i.e. imported) by other modules, it can be uninstalled only after the dependent modules have first been uninstalled. Installation and uninstallation must be done in dependency order and reverse-dependency order, as appropriate.

22.1.5 Supported YANG Models

The currently supported models are `kea-dhcp4-server` and `kea-dhcp6-server`. There is partial support for `ietf-dhcpv6-server`, but the primary focus of testing has been on Kea DHCP servers. Other models (*`kea-dhcp-ddns`* and *`kea-ctrl-agent`*) are currently not supported.

22.1.6 Using the NETCONF Agent

The NETCONF agent follows this algorithm:

- For each managed server, get the initial configuration from the server through the control socket.
- Open a connection with the Sysrepo environment and establish two sessions with the startup and running datastores.

- Check that the used (not-essential) and required (essential) modules are installed in the Sysrepo repository at the right revision. If an essential module - that is, a module where the configuration schema for a managed server is defined - is not installed, raise a fatal error.
- For each managed server, get the YANG configuration from the startup datastore, translate it to JSON, and load it onto the server being configured.
- For each managed server, subscribe a module change callback using its model name.
- When a running configuration is changed, try to validate or load the updated configuration via the callback to the managed server.

22.1.7 Configuration

The behavior described in *Using the NETCONF Agent* is controlled by several configuration flags, which can be set in the global scope or in a specific managed-server scope. If the latter, the value defined in the managed-server scope takes precedence. These flags are:

- **boot-update** - controls the initial configuration phase; when **true** (the default), the initial configuration retrieved from the classic Kea server JSON configuration file is loaded first, and then the startup YANG model is loaded. This setting lets administrators define a control socket in the local JSON file and then download the configuration from YANG. When set to **false**, this phase is skipped.
- **subscribe-changes** - controls the module change subscription; when **true** (the default), a module change callback is subscribed, but when **false** the phase is skipped and running configuration updates are disabled. When set to **true**, the running datastore is used to subscribe for changes.
- **validate-changes** - controls how Kea monitors changes in the Sysrepo configuration. Sysrepo offers two stages where Kea can interact: validation and application. At the validation (or **SR_EV_CHANGE** event, in the Sysrepo naming convention) stage, Kea retrieves the newly committed configuration and verifies it. If the configuration is incorrect for any reason, the Kea servers reject it and the error is propagated back to the Sysrepo, which then returns an error. This step only takes place if **validate-changes** is set to **true**. In the application (or **SR_EV_UPDATE** event in the Sysrepo naming convention) stage, the actual configuration is applied. At this stage Kea can receive the configuration, but it is too late to signal back any errors as the configuration has already been committed.

The idea behind the initial configuration phase is to boot Kea servers with a minimal configuration which includes only a control socket, making them manageable. For instance, for the DHCPv4 server:

```
{
  "Dhcp4": {
    "control-socket": {
      "socket-name": "kea-dhcp4-ctrl.sock",
      "socket-type": "unix"
    }
  }
}
```

With module change subscriptions enabled, the *kea-netconf* daemon monitors any configuration changes as they appear in the Sysrepo. Such changes can be done using the *sysrepocfg* tool or remotely using any NETCONF client. For details, please see the Sysrepo documentation or *A Step-by-Step NETCONF Agent Operation Example*. Those tools can be used to modify YANG configurations in the running datastore. Note that committed configurations are only updated in the running datastore; to keep them between server reboots they must be copied to the startup datastore.

When module changes are tracked (using **subscribe-changes** set to **true**) and the running configuration has changed (e.g. using *sysrepocfg* or any NETCONF client), the callback validates the modified configuration (if **validate-changes** was not set to **false**) and runs a second time to apply the new configuration. If the validation fails, the callback is still called again but with an **SR_EV_ABORT** (vs. **SR_EV_DONE**) event with rollback changes.

The returned code of the callback on an `SR_EV_DONE` event is ignored, as it is too late to refuse a bad configuration.

There are four ways in which a modified YANG configuration might be incorrect:

1. It could be non-compliant with the schema, e.g. an unknown entry, missing a mandatory entry, a value with a bad type, or not matching a constraint.
2. It could fail to be translated from YANG to JSON, e.g. an invalid user context.
3. It could fail Kea server sanity checks, e.g. an out-of-subnet-pool range or an unsupported database type.
4. The syntax may be correct and pass server sanity checks but the configuration could fail to run, e.g. the configuration specifies database credentials but the database refuses the connection.

The first case is handled by Sysrepo. The second and third cases are handled by *kea-netconf* in the validation phase (if not disabled by setting `validate-changes` to `true`). The last case causes the application phase to fail without a sensible report to Sysrepo.

The managed Kea servers and agents are described in the `managed-servers` section. Each sub-section begins with the service name: `dhcp4`, `dhcp6`, `d2` (the DHCP-DDNS server does not support the control-channel feature yet), and `ca` (the control agent).

Each managed server entry may contain:

- `control-flags` - `boot-update`, `subscribe-changes`, and/or `validate-changes`.
- `model` - specifies the YANG model/module name. For each service, the default is the corresponding Kea YANG model, e.g. for `"dhcp4"` it is `"kea-dhcp4-server"`.
- `control-socket` - specifies the control socket for managing the service configuration.

A control socket is specified by:

- `socket-type` - the socket type is either `stdout`, `unix`, or `http`. `stdout` is the default; it is not really a socket, but it allows *kea-netconf* to run in debugging mode where everything is printed on stdout, and it can also be used to redirect commands easily. `unix` is the standard direct server control channel, which uses UNIX sockets; `http` uses a control agent, which accepts HTTP connections.
- `socket-name` - the local socket name for the `unix` socket type (default empty string).
- `socket-url` - the HTTP URL for the `http` socket type (default `http://127.0.0.1:8000/`).

User contexts can store arbitrary data as long as they are in valid JSON syntax and their top-level element is a map (i.e. the data must be enclosed in curly brackets). They are accepted at the NETCONF entry, i.e. below the top-level, managed-service entry, and control-socket entry scopes.

Hook libraries can be loaded by the NETCONF agent just as with other servers or agents; however, currently no hook points are defined. The `hooks-libraries` list contains the list of hook libraries that should be loaded by *kea-netconf*, along with their configuration information specified with `parameters`.

Please consult *Logging* for details on how to configure logging. The name of the NETCONF agent's main logger is *kea-netconf*, as given in the example above.

22.1.8 A kea-netconf Configuration Example

The following example demonstrates the basic NETCONF configuration. More examples are available in the `doc/examples/netconf` directory in the Kea sources.

```
// This is a simple example of a configuration for the NETCONF agent.
// This server provides a YANG interface for all Kea servers and the agent.
{
  "Netconf":
```

(continues on next page)

(continued from previous page)

```

{
    // Control flags can be defined in the global scope or
    // in a managed server scope. Precedences are:
    // - use the default value (true)
    // - use the global value
    // - use the local value.
    // So this overwrites the default value:
    "boot-update": false,

    // This map specifies how each server is managed. For each server there
    // is a name of the YANG model to be used and the control channel.
    // Currently three control channel types are supported:
    // "stdout" which outputs the configuration on the standard output,
    // "unix" which uses the local control channel supported by the
    // "dhcp4" and "dhcp6" servers ("d2" support is not yet available),
    // and "http" which uses the Control Agent "ca" to manage itself or
    // to forward commands to "dhcp4" or "dhcp6".
    "managed-servers":
    {
        // This is how kea-netconf can communicate with the DHCPv4 server.
        "dhcp4":
        {
            "comment": "DHCPv4 server",
            "model": "kea-dhcp4-server",
            "control-socket":
            {
                "socket-type": "unix",
                "socket-name": "kea4-ctrl-socket"
            }
        },

        // DHCPv6 parameters.
        "dhcp6":
        {
            "model": "kea-dhcp6-server",
            "control-socket":
            {
                "socket-type": "unix",
                "socket-name": "kea6-ctrl-socket"
            }
        },

        "d2":
        {
            "model": "kea-dhcp-ddns",
            "control-socket":
            {
                "socket-type": "unix",
                "socket-name": "kea-ddns-ctrl-socket",
                "user-context": { "in-use": true }
            }
        },
    },
}

```

(continues on next page)

(continued from previous page)

```

// Of course the Control Agent (CA) supports HTTP.
"ca":
{
    "model": "kea-ctrl-agent",
    "control-socket":
    {
        "socket-type": "http",
        "socket-url": "http://127.0.0.1:8000/"
    }
}
},

// kea-netconf is able to load hook libraries that augment its operation.
// Currently there are no hook points defined in kea-netconf
// processing.
"hooks-libraries": [
    // The hook libraries list may contain more than one library.
    {
        // The only necessary parameter is the library filename.
        "library": "/opt/local/custom_hooks_example.so",

        // Some libraries may support parameters. Make sure you
        // type this section carefully, as kea-netconf does not
        // validate it (because the format is library-specific).
        "parameters": {
            "param1": "foo"
        }
    }
],

// Similar to other Kea components, NETCONF also uses logging.
"loggers": [
    {
        "name": "kea-netconf",
        "output-options": [
            {
                "output": "/var/log/kea/kea-netconf.log",
                // Several additional parameters are possible in
                // addition to the typical output.
                // Flush determines whether logger flushes output
                // to a file.
                // Maxsize determines maximum filesize before
                // the file is rotated.
                // Maxver specifies the maximum number of
                // rotated files to be kept.
                "flush": true,
                "maxsize": 204800,
                "maxver": 4
            }
        ],
        "severity": "INFO",
    }
]

```

(continues on next page)

(continued from previous page)

```

        "debuglevel": 0
    }
}
}

```

22.1.9 Starting and Stopping the NETCONF Agent

kea-netconf accepts the following command-line switches:

- **-c file** - specifies the configuration file.
- **-d** - specifies whether the agent logging should be switched to debug/verbose mode. In verbose mode, the logging severity and debuglevel specified in the configuration file are ignored and "debug" severity and the maximum debuglevel (99) are assumed. The flag is convenient for temporarily switching the server into maximum verbosity, e.g. when debugging.
- **-t file** - specifies the configuration file to be tested. *kea-netconf* attempts to load it and conducts sanity checks; certain checks are possible only while running the actual server. The actual status is reported with exit code (0 = configuration appears valid, 1 = error encountered). Kea prints out log messages to standard output and error to standard error when testing the configuration.
- **-v** - displays the version of *kea-netconf* and exits.
- **-V** - displays the extended version information for *kea-netconf* and exits. The listing includes the versions of the libraries dynamically linked to Kea.
- **-W** - displays the Kea configuration report and exits. The report is a copy of the `config.report` file produced by `meson setup`; it is embedded in the executable binary.

The contents of the `config.report` file may also be accessed by examining certain libraries in the installation tree or in the source tree.

```

# from installation using libkea-process.so
$ strings ${prefix}/lib/libkea-process.so | sed -n 's/;;; //p'

# from sources using libkea-process.so
$ strings src/lib/process/.libs/libkea-process.so | sed -n 's/;;; //p'

# from sources using libkea-process.a
$ strings src/lib/process/.libs/libkea-process.a | sed -n 's/;;; //p'

# from sources using libcfgrpt.a
$ strings src/lib/process/cfgrpt/.libs/libcfgrpt.a | sed -n 's/;;; //p'

```

22.1.10 A Step-by-Step NETCONF Agent Operation Example

Note

Copies of example configurations presented within this section can be found in the Kea source code, under `doc/examples/netconf/kea-dhcp6-operations`.

22.1.10.1 Setup of NETCONF Agent Operation Example

The test box has an Ethernet interface named `eth1`. On some systems it is possible to rename interfaces; for instance, on Linux with an `ens38` interface:

```
# ip link set down dev ens38
# ip link set name eth1 dev ens38
# ip link set up dev eth1
```

The interface must have an address in the test prefix:

```
# ip -6 addr add 2001:db8::1/64 dev eth1
```

The Kea DHCPv6 server must be launched with the configuration specifying a control socket used to receive control commands. The `kea-netconf` process uses this socket to communicate with the DHCPv6 server, i.e. it pushes translated configurations to that server using control commands. The following is an example control socket specification for the Kea DHCPv6 server:

```
{
  "Dhcp6": {
    "control-socket": {
      "socket-name": "kea-dhcp6-ctrl.sock",
      "socket-type": "unix"
    }
  }
}
```

In order to launch the Kea DHCPv6 server using the configuration contained within the `boot.json` file, run:

```
# kea-dhcp6 -d -c boot.json
```

The current configuration of the server can be fetched via a control socket by running:

```
# echo '{ "command": "config-get" }' | socat UNIX:/opt/kea/var/run/kea/kea-dhcp6-ctrl.
↪sock '_,ignoreeof'
```

The following is the example `netconf.json` configuration for `kea-netconf`, to manage the Kea DHCPv6 server:

```
{
  "Netconf": {
    "loggers": [
      {
        "debuglevel": 99,
        "name": "kea-netconf",
        "output-options": [
          {
            "output": "stderr"
          }
        ],
        "severity": "DEBUG"
      }
    ],
    "managed-servers": {
      "dhcp6": {
        "control-socket": {
```

(continues on next page)

(continued from previous page)

```

        "socket-name": "kea-dhcp6-ctrl.sock",
        "socket-type": "unix"
    }
}
}
}
}

```

Note that in production there should not be a need to log at the DEBUG level.

The Kea NETCONF agent is launched by:

```
# kea-netconf -d -c netconf.json
```

Now that both *kea-netconf* and *kea-dhcp6* are running, it is possible to populate updates to the configuration to the DHCPv6 server. The following is the configuration extracted from *startup.xml*:

```

<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet6>
    <id>1</id>
    <pool>
      <start-address>2001:db8::1:0</start-address>
      <end-address>2001:db8::1:ffff</end-address>
      <prefix>2001:db8::1:0/112</prefix>
    </pool>
    <subnet>2001:db8::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
</config>

```

To populate this new configuration:

```
$ sysrepoctl -d startup -f xml -m kea-dhcp6-server --import=startup.xml
```

kea-netconf pushes the configuration found in the Sysrepo startup datastore to all Kea servers during its initialization phase, after it subscribes to module changes in the Sysrepo running datastore. This action copies the configuration from the startup datastore to the running datastore and enables the running datastore, making it available.

Changes to the running datastore are applied after validation to the Kea servers. Note that they are not by default copied back to the startup datastore, i.e. changes are not permanent.

Note

kea-netconf fetches the entire configuration from any Sysrepo datastore in a single *get-config* NETCONF operation. Prior to Kea 2.3.2, a *get-config* operation was done for each leaf and leaf-list node.

22.1.10.2 Example of Error Handling in NETCONF Operation

There are four classes of issues with configurations applied via NETCONF:

1. The configuration does not comply with the YANG schema.
2. The configuration cannot be translated from YANG to the Kea JSON.
3. The configuration is rejected by the Kea server.
4. The configuration was validated by the Kea server but cannot be applied.

In the first case, consider the following `BAD-schema.xml` configuration file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet4>
    <id>1</id>
    <pool>
      <start-address>2001:db8::1:0</start-address>
      <end-address>2001:db8::1:ffff</end-address>
      <prefix>2001:db8::1:0/112</prefix>
    </pool>
    <subnet>2001:db8::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
</config>
```

It is directly rejected by `sysrepcfg`:

```
$ sysrepcfg -d running -f xml -m kea-dhcp6-server --import=BAD-schema.xml
```

In the second case, the configuration is rejected by `kea-netconf`. For example, consider this `BAD-translator.xml` file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet6>
    <id>1</id>
    <pool>
      <start-address>2001:db8::1:0</start-address>
      <end-address>2001:db8::1:ffff</end-address>
      <prefix>2001:db8::1:0/112</prefix>
    </pool>
    <subnet>2001:db8::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
```

(continues on next page)

(continued from previous page)

```
<user-context>bad</user-context>
</config>
```

In the third case, the configuration is presented to the Kea DHCPv6 server and fails to validate, as in this BAD-config.xml file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet6>
    <id>1</id>
    <pool>
      <start-address>2001:db8:1::0</start-address>
      <end-address>2001:db8:1::ffff</end-address>
      <prefix>2001:db8:1::0/112</prefix>
    </pool>
    <subnet>2001:db8::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
</config>
```

In the last case, the misconfiguration is detected too late and the change must be reverted in Sysrepo, e.g. using the startup datastore as a backup.

22.1.10.3 NETCONF Operation Example with Two Pools

This example adds a second pool to the initial (i.e. startup) configuration in the twopools.xml file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet6>
    <id>1</id>
    <pool>
      <start-address>2001:db8::1:0</start-address>
      <end-address>2001:db8::1:ffff</end-address>
      <prefix>2001:db8::1:0/112</prefix>
    </pool>
    <pool>
      <start-address>2001:db8::2:0</start-address>
      <end-address>2001:db8::2:ffff</end-address>
      <prefix>2001:db8::2:0/112</prefix>
    </pool>
    <subnet>2001:db8::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
</config>
```

(continues on next page)

(continued from previous page)

```
</control-socket>
</config>
```

This configuration is installed by:

```
$ sysrepcfg -d running -f xml -m kea-dhcp6-server --import=twopools.xml
```

22.1.10.4 NETCONF Operation Example with Two Subnets

This example specifies two subnets in the twosubnets.xml file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <subnet6>
    <id>1</id>
    <pool>
      <start-address>2001:db8:1::</start-address>
      <end-address>2001:db8:1::ffff</end-address>
      <prefix>2001:db8:1::/112</prefix>
    </pool>
    <subnet>2001:db8:1::/64</subnet>
  </subnet6>
  <subnet6>
    <id>2</id>
    <pool>
      <start-address>2001:db8:2::</start-address>
      <end-address>2001:db8:2::ffff</end-address>
      <prefix>2001:db8:2::/112</prefix>
    </pool>
    <subnet>2001:db8:2::/64</subnet>
  </subnet6>
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <control-socket>
    <socket-name>kea-dhcp6-ctrl.sock</socket-name>
    <socket-type>unix</socket-type>
  </control-socket>
</config>
```

This configuration is installed by:

```
$ sysrepcfg -d running -f xml -m kea-dhcp6-server --import=twosubnets.xml
```

22.1.10.5 NETCONF Operation Example With Logging

This example adds a logger entry to the initial (i.e. startup) configuration in the logging.xml file:

```
<config xmlns="urn:ietf:params:xml:ns:yang:kea-dhcp6-server">
  <interfaces-config>
    <interfaces>eth1</interfaces>
  </interfaces-config>
  <subnet6>
```

(continues on next page)

(continued from previous page)

```

<id>1</id>
<pool>
  <start-address>2001:db8::1:0</start-address>
  <end-address>2001:db8::1:ffff</end-address>
  <prefix>2001:db8::1:0/112</prefix>
</pool>
<subnet>2001:db8::/64</subnet>
</subnet6>
<control-socket>
  <socket-name>kea-dhcp6-ctrl.sock</socket-name>
  <socket-type>unix</socket-type>
</control-socket>
<logger>
  <name>kea-dhcp6</name>
  <output-option>
    <output>stderr</output>
  </output-option>
  <debuglevel>99</debuglevel>
  <severity>DEBUG</severity>
</logger>
</config>

```

The corresponding Kea configuration in JSON is:

```

{
  "Dhcp6": {
    "control-socket": {
      "socket-name": "kea-dhcp6-ctrl.sock",
      "socket-type": "unix"
    },
    "interfaces-config": {
      "interfaces": [ "eth1" ]
    },
    "subnet6": [
      {
        "id": 1,
        "pools": [
          {
            "pool": "2001:db8::1:0/112"
          }
        ],
        "subnet": "2001:db8::/64"
      }
    ],
    "loggers": [
      {
        "name": "kea-dhcp6",
        "output-options": [
          {
            "output": "stderr"
          }
        ]
      }
    ],
  }
}

```

(continues on next page)

(continued from previous page)

```

        "severity": "DEBUG",
        "debuglevel": 99
    }
]
}
}

```

Finally, any of the previous examples can be replayed by using `sysrepocfg` in edit mode as follows:

```
$ sysrepocfg -d running -f xml -m kea-dhcp6-server --edit
```

or by using a NETCONF client like `netopeer2-cli` from the [Netopeer2](#) NETCONF Toolset.

22.1.10.6 Migrating YANG Data From a Prior Sysrepo Version

1. Shut down `kea-netconf`. This ensures that backups for both datastores are done at the same configuration state and that no change happens between exporting them.

2. Make data backups for all YANG modules, with one XML for each datastore.

```
$ sysrepocfg --datastore running --export=save.xml --format=xml
$ sysrepocfg --datastore startup --export=save.xml --format=xml
```

Note

Sysrepo v0 does not support import/export of all YANG modules; this capability was added in Sysrepo v1. Users that are migrating from Sysrepo v0 will need to do per-module backups. This has the added benefit of isolating potential failures and preventing them from affecting all modules. The command is the same, except it has the module name added to it at the end.

```
$ sysrepocfg --datastore running --export=save.xml --format=xml kea-dhcp6-server
$ sysrepocfg --datastore startup --export=save.xml --format=xml kea-dhcp6-server
```

3. Upgrade Sysrepo to the newer version and then run:

```
$ sysrepocfg --datastore running --import=save.xml
$ sysrepocfg --datastore startup --import=save.xml
```

22.2 GSS-TSIG

22.2.1 GSS-TSIG Overview

Kea provides support for DNS updates, which can be protected using Transaction Signatures (or TSIG). This protection is often adequate. However, some systems, in particular Active Directory (AD) on Microsoft Windows servers, have chosen to adopt a more complex GSS-TSIG approach that offers additional capabilities, such as using negotiated dynamic keys.

Kea supports GSS-TSIG to protect DNS updates sent by the Kea DHCP-DDNS (D2) server in a hook, called `libddns_gss_tsig.so`.


GSS-TSIG is defined in [RFC 3645](#). The GSS-TSIG protocol itself is an implementation of generic GSS-API v2 services, defined in [RFC 2743](#).

Many protocols are involved in this mechanism:

- Kerberos 5 - [RFC 4120](#), which provides the security framework;
- GSS-API (Generic Security Services Application Program Interface) - [RFC 2743](#) for the API, [RFC 2744](#) for the C bindings, and [RFC 4121](#) for the application to Kerberos 5;
- SPNEGO (Simple and Protected GSS-API Negotiation Mechanism) - [RFC 4178](#) for the negotiation;
- DNS update [RFC 2136](#);
- TSIG (Secret Key Transaction Authentication for DNS) - [RFC 8945](#), which protects DNS exchanges;
- Secure Domain Name System (DNS) Dynamic Update - [RFC 3007](#), which is the application of TSIG to DNS update protection;
- TKEY (Secret Key Establishment for DNS) - [RFC 2930](#), which establishes secret keys for TSIG by transmitting crypto payloads between DNS parties; and
- GSS-TSIG - [RFC 3645](#), which is the application of GSS-API to TSIG.


To summarize, GSS-API for Kerberos 5 with SPNEGO and TKEY are used to negotiate a security context between the Kea D2 server and a DNS server:

TKEY Exchange (GSS-TSIG hook)



The security context is then used by GSS-TSIG to protect updates:

DNS Update Exchange (GSS-TSIG hook)



The Kea implementation of GSS-TSIG uses a GSS-API for Kerberos 5 with the SPNEGO library. Two implementations meet this criteria: MIT Kerberos 5 and Heimdal.

22.2.2 GSS-TSIG Compilation

The following procedure was tested on Debian 12. A similar approach can be applied to other systems.

1. Install either the MIT (`libkrb5-dev`) or the Heimdal (`heimdal-dev`) library, for instance:

```
sudo apt install libkrb5-dev
```

2. Obtain the Kea sources, extract the Kea sources.
3. Set up the build:

```
meson setup build -D krb5=enabled
```

4. Make sure `build/config.report` shows the Kerberos5 GSS-API entry.

```
Kerberos5 GSS-API:
VERSION:          1.20.1
VENDOR:           MIT
```

Either the `krb5.pc` file or the `krb5-config` tool are required to be present. They are provided by both MIT Kerberos 5 and Heimdal; however, on some systems where both Kerberos 5 and Heimdal are installed, it is a symbolic link to one of them. If the tool is not in the standard location, it can be selected by setting `PKG_CONFIG_PATH=/path/to/krb5.pc:${PKG_CONFIG_PATH}` or `PATH=/path/to/krb5-config:${PATH}` prior to `meson setup`.

6. Compile with `meson compile -C build`.
7. After compilation, `libddns_gss_tsig.so` is available in the `build/src/hooks/d2/gss_tsig` directory. It can be loaded by `kea-dhcp-ddns`.

`libddns_gss_tsig.so` supports both MIT Kerberos 5 and Heimdal implementations. Heimdal is picky about security-sensitive file permissions and is known to emit an unclear error message. It is a good idea to keep these files plain, with one link and no access for the group or other users.

The `krb5-config` script should provide an `--vendor` option which identifies the implementation.

22.2.3 GSS-TSIG Deployment

Before using GSS-TSIG, a GSS-TSIG capable DNS server, such as BIND 9 or Microsoft Active Directory (AD), must be deployed. Other GSS-TSIG capable implementations may work, but have not been tested.

22.2.3.1 Kerberos 5 Setup

There are two kinds of key tables (keytab files): the system one used by servers, and client tables used by clients. For Kerberos 5, Kea is a **client**.

Install the Kerberos 5 client library and `kadmin` tool:

```
sudo apt install krb5-kdc krb5-admin-server
```

The following examples use the `EXAMPLE.ORG` realm to demonstrate required configuration steps and settings.

The Kerberos 5 client library must be configured to accept incoming requests for the realm `EXAMPLE.ORG` by updating the `krb5.conf` file (e.g. on Linux: `/etc/krb5.conf`):

```
[libdefaults]
default_realm = EXAMPLE.ORG
kdc_timesync = 1
ccache_type = 4
```

(continues on next page)

(continued from previous page)

```

forwardable = true
proxiabile = true

[realms]
EXAMPLE.ORG = {
    kdc = kdc.example.org
    admin_server = kdc.example.org
}

```

In addition to the `krb5.conf` file, the `kdc.conf` file can be used (e.g. on Linux: `/etc/krb5kdc/kdc.conf`):

```

[kdcdefaults]
kdc_ports = 750,88

[realms]
EXAMPLE.ORG = {
    database_name = /var/lib/krb5kdc/principal
    admin_keytab = FILE:/etc/krb5kdc/kadm5.keytab
    acl_file = /etc/krb5kdc/kadm5.acl
    key_stash_file = /etc/krb5kdc/stash
    kdc_ports = 750,88
    max_life = 10h 0m 0s
    max_renewable_life = 7d 0h 0m 0s
    master_key_type = des3-hmac-sha1
    #supported_encetypes = aes256-cts:normal aes128-cts:normal
    default_principal_flags = +preauth
}

```

The `kadmind` daemon Access Control List (ACL) must be configured to give permissions to the DNS client principal to access the Kerberos 5 database (e.g. on Linux: `/etc/krb5kdc/kadm5.acl`):

```
DHCP/admin.example.org@EXAMPLE.ORG *
```

The administrator password for the default realm must be set:

```
krb5_newrealm
```

After the following message is displayed, enter the password for the default realm:

```

This script should be run on the master KDC/admin server to initialize
a Kerberos realm. It will ask you to type in a master key password.
This password will be used to generate a key that is stored in
/etc/krb5kdc/stash. You should try to remember this password, but it
is much more important that it be a strong password than that it be
remembered. However, if you lose the password and /etc/krb5kdc/stash,
you cannot decrypt your Kerberos database.
Loading random data
Initializing database '/var/lib/krb5kdc/principal' for realm 'EXAMPLE.ORG',
master key name 'K/M@EXAMPLE.ORG'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
Enter KDC database master key:

```

Then retype the password:


```
Re-enter KDC database master key to verify:
```

If successfully applied, the following message is displayed:

```
Now that your realm is set up you may wish to create an administrative
principal using the addprinc subcommand of the kadmin.local program.
Then, this principal can be added to /etc/krb5kdc/kadm5.acl so that
you can use the kadmin program on other computers. Kerberos admin
principals usually belong to a single user and end in /admin. For
example, if jruser is a Kerberos administrator, then in addition to
the normal jruser principal, a jruser/admin principal should be
created.
```

```
Don't forget to set up DNS information so your clients can find your
KDC and admin servers. Doing so is documented in the administration
guide.
```

The next step is to create the principals for the BIND 9 DNS server (the service protected by the GSS-TSIG TKEY) and for the DNS client (the Kea DHCP-DDNS server).

The BIND 9 DNS server principal (used for authentication) is created the following way:

```
kadmin.local -q "addprinc -randkey DNS/server.example.org"
```

If successfully created, the following message is displayed:

```
No policy specified for DNS/server.example.org@EXAMPLE.ORG; defaulting to no policy
Authenticating as principal root/admin@EXAMPLE.ORG with password.
Principal "DNS/server.example.org@EXAMPLE.ORG" created.
```

The DNS server principal must be exported so that it can be used by the BIND 9 DNS server. Only this principal is required, and it is exported to the keytab file with the name `dns.keytab`.

```
kadmin.local -q "ktadd -k /usr/local/share/kea/dns.keytab DNS/server.example.org"
```

If successfully exported, the following message is displayed:

```
Authenticating as principal root/admin@EXAMPLE.ORG with password.
Entry for principal DNS/server.example.org with kvno 2, encryption type aes256-cts-hmac-
↪sha1-96 added to keytab WRFILE:/usr/local/share/kea/dns.keytab.
Entry for principal DNS/server.example.org with kvno 2, encryption type aes128-cts-hmac-
↪sha1-96 added to keytab WRFILE:/usr/local/share/kea/dns.keytab.
```

The DHCP client principal (used by the Kea DHCP-DDNS server) is created the following way:

```
kadmin.local -q "addprinc -randkey DHCP/admin.example.org"
```

If successfully created, the following message is displayed:

```
No policy specified for DHCP/admin.example.org@EXAMPLE.ORG; defaulting to no policy
Authenticating as principal root/admin@EXAMPLE.ORG with password.
Principal "DHCP/admin.example.org@EXAMPLE.ORG" created.
```

The DHCP client principal must be exported so that it can be used by the Kea DHCP-DDNS server and the GSS-TSIG hook library. It is exported to the client keytab file with the name `dhcp.keytab`.

```
kadmin.local -q "ktadd -k /usr/local/share/kea/dhcp.keytab DHCP/admin.example.org"
```

Finally, the `krb5-admin-server` must be restarted:

```
systemctl restart krb5-admin-server.service
```

22.2.3.2 BIND 9 with GSS-TSIG Configuration

The BIND 9 DNS server must be configured to use GSS-TSIG, and to use the previously exported DNS server principal from the keytab file `dns.keytab`. Updating the `named.conf` file is required:

```
options {
    ...
    directory "/var/cache/bind";
    dnssec-validation auto;
    listen-on-v6 { any; };
    tkey-gssapi-keytab "/etc/bind/dns.keytab";
};
zone "example.org" {
    type master;
    file "/var/lib/bind/db.example.org";
    update-policy {
        grant "DHCP/admin.example.org@EXAMPLE.ORG" zonesub any;
    };
};
zone "84.102.10.in-addr.arpa" {
    type master;
    file "/etc/bind/db.10";
};
```

The zone files should have an entry for the server principal FQDN `server.example.org`.

The `/etc/bind/db.10` file needs to be created or updated:

```
;
; BIND reverse data file for local loopback interface
;
$TTL      604800                ; 1 week
@         IN      SOA      server.example.org. root.example.org. (
                                2          ; Serial
                                604800     ; Refresh
                                86400      ; Retry
                                2419200    ; Expire
                                604800     ; Negative Cache TTL
                                )
;
@         IN      NS       ns.
40        IN      PTR      ns.example.org.
```

The `/var/lib/bind/db.example.org` file needs to be created or updated:

```
$ORIGIN .
$TTL      604800                ; 1 week
example.org IN SOA  server.example.org. root.example.org. (
```

(continues on next page)

(continued from previous page)

```

        8          ; serial
        604800     ; refresh (1 week)
        86400      ; retry (1 day)
        2419200    ; expire (4 weeks)
        604800     ; minimum (1 week)
    )
    NS      example.org.
    A       ${BIND9_IP_ADDR}
    AAAA    ::1
$ORIGIN example.org.
kdc        A       ${KDC_IP_ADDR}
server     A       ${BIND9_IP_ADDR}

```

After any configuration change the server must be reloaded or restarted:

```
systemctl restart named.service
```

It is possible to get the status or restart the logs:

```
systemctl status named.service
journalctl -u named | tail -n 30
```

22.2.3.3 Windows Active Directory Configuration

This sub-section is based on an Amazon AWS provided Microsoft Windows Server 2016 with Active Directory pre-installed, so it describes only the steps used for GSS-TSIG deployment. (For the complete configuration process, please refer to Microsoft's documentation or other external resources. We found [this](#) tutorial very useful during configuration of our internal QA testing systems.)

Two Active Directory (AD) user accounts are needed:

- the first account is used to download AD information, such as the client key table of Kea
- the second account is mapped to the Kea DHCP client principal

Kea needs to know:

- the server IP address
- the domain/realm name: the domain is in lower case, the realm in upper case, both without a final dot
- the server name

The second account (named **kea** below) is used to create a Service Principal Name (SPN):

```
setspn -S DHCP/kea.<domain> kea
```

After a shared secret key is generated and put in a key table file:

```
ktpass -princ DHCP/kea.<domain>@<REALM> -mapuser kea +rndpass -mapop set -ptype KRB5_NT_
↪PRINCIPAL -out dhcp.keytab
```

The `dhcp.keytab` takes the same usage as for UNIX Kerberos.

22.2.3.4 GSS-TSIG Troubleshooting

While testing GSS-TSIG integration with Active Directory we came across one very cryptic error:

```
INFO [kea-dhcp-ddns.gss-tsig-hooks/4678.139690935890624] GSS_TSIG_VERIFY_FAILED GSS-
↳TSIG verify failed: gss_verify_mic failed with GSSAPI error:
Major = 'A token had an invalid Message Integrity Check (MIC)' (393216), Minor = 'Packet_
↳was replayed in wrong direction' (100002).
```

In our case, the problem was that the Kea D2 server was trying to perform an update of a reverse DNS zone while it was not configured. An easy solution is to add a reverse DNS zone similar to the one configured in Kea. To do that, open the "DNS Manager" and choose "DNS" from the list; from the dropdown list, choose "Reverse Lookup Zones"; then click "Action" and "New Zone"; finally, follow the New Zone Wizard to add a new zone.

The standard requires both anti-replay and sequence services. Experiences with the BIND 9 nsupdate showed the sequence service led to problems so it is disabled by default in the hook. It seems the anti-replay service can also lead to problems with Microsoft DNS servers so it is now configurable. Note that these security services are useless for DNS dynamic update which was designed to run over UDP so with out of order and duplicated messages.

22.2.4 Using GSS-TSIG

There are a number of steps required to enable the GSS-TSIG mechanism:

1. *libddns_gss_tsig.so* must be loaded by *kea-dhcp-ddns*.
2. The GSS-TSIG-capable DNS servers must be specified with their parameters.

An excerpt from a D2 server configuration is provided below; more examples are available in the `doc/examples/ddns` directory in the Kea sources.

```
1 {
2   "DhcpDdns": {
3     // The following parameters are used to receive NCRs (NameChangeRequests)
4     // from the local Kea DHCP server. Make sure your kea-dhcp4 and kea-dhcp6
5     // matches this.
6     "ip-address": "127.0.0.1",
7     "port": 53001,
8     "dns-server-timeout" : 1000,
9
10    // Forward zone: secure.example.org. It uses GSS-TSIG. It is served
11    // by two DNS servers, which listen for DDNS requests at 192.0.2.1
12    // and 192.0.2.2.
13    "forward-ddns":
14    {
15      "ddns-domains":
16      [
17        // DdnsDomain for zone "secure.example.org."
18        {
19          "name": "secure.example.org.",
20          "comment": "DdnsDomain example",
21          "dns-servers":
22          [
23            { // This server has an entry in gss/servers and
24              // thus will use GSS-TSIG.
25              "ip-address": "192.0.2.1"
26            },
```

(continues on next page)

(continued from previous page)

```

27         { // This server also has an entry there, so will
28           // use GSS-TSIG, too.
29           "ip-address": "192.0.2.2",
30           "port": 5300
31         }
32       ]
33     }
34   ],
35 },
36
37 // Reverse zone: we want to update the reverse zone "2.0.192.in-addr.arpa".
38 "reverse-ddns":
39 {
40   "ddns-domains":
41   [
42     {
43       "name": "2.0.192.in-addr.arpa.",
44       "dns-servers":
45       [
46         {
47           // There is a GSS-TSIG definition for this server (see
48           // DhcpsDdns/gss-tsig/servers), so it will use
49           // Krb/GSS-TSIG.
50           "ip-address": "192.0.2.1"
51         }
52       ]
53     }
54   ]
55 },
56
57 // The GSS-TSIG hook is loaded and its configuration is specified here.
58 "hooks-libraries": [
59 {
60   "library": "/opt/lib/libddns_gss_tsig.so",
61   "parameters": {
62     // This section governs the GSS-TSIG integration. Each server
63     // mentioned in forward-ddns and/or reverse-ddns needs to have
64     // an entry here to be able to use GSS-TSIG defaults (optional,
65     // if specified they apply to all the GSS-TSIG servers, unless
66     // overwritten on specific server level).
67
68     "server-principal": "DNS/server.example.org@EXAMPLE.ORG",
69     "client-principal": "DHCP/admin.example.org@EXAMPLE.ORG",
70
71     // client-keytab and credentials-cache can both be used to
72     // store client keys. As credentials cache is more flexible,
73     // it is recommended to use it. Typically, using both at the
74     // same time may cause problems.
75     // "client-keytab": "FILE:/etc/dhcp.keytab", // toplevel only
76     "credentials-cache": "FILE:/etc/ccache", // toplevel only
77     "gss-replay-flag": true, // GSS anti replay service
78     "gss-sequence-flag": false, // no GSS sequence service

```

(continues on next page)

(continued from previous page)

```

79     "tkey-lifetime": 3600, // 1 hour
80     "rekey-interval": 2700, // 45 minutes
81     "retry-interval": 120, // 2 minutes
82     "tkey-protocol": "TCP",
83     "fallback": false,
84
85     // The list of GSS-TSIG capable servers
86     "servers": [
87         {
88             // First server (identification is required)
89             "id": "server1",
90             "domain-names": [ ], // if not specified or empty, will
91                                 // match all domains that want to
92                                 // use this IP+port pair
93             "ip-address": "192.0.2.1",
94             "port": 53,
95             "server-principal": "DNS/server1.example.org@EXAMPLE.ORG",
96             "client-principal": "DHCP/admin1.example.org@EXAMPLE.ORG",
97             "gss-replay-flag": false, // no GSS anti replay service
98             "gss-sequence-flag": false, // no GSS sequence service
99             "tkey-lifetime": 7200, // 2 hours
100            "rekey-interval": 5400, // 90 minutes
101            "retry-interval": 240, // 4 minutes
102            "tkey-protocol": "TCP",
103            "fallback": true // if no key is available fallback to the
104                           // standard behavior (vs skip this server)
105        },
106        {
107            // The second server (it has most of the parameters missing
108            // as those are using the defaults specified above)
109            "id": "server2",
110            "ip-address": "192.0.2.2",
111            "port": 5300
112        }
113    ]
114 }
115 ]
116
117
118 // Additional parameters, such as logging, control socket and
119 // others omitted for clarity.
120 }
121
122 }
```

This configuration file contains a number of extra elements.

First, a list of forward and/or reverse domains with related DNS servers identified by their IP+port pairs is defined. If the port is not specified, the default of 53 is assumed. This is similar to basic mode, with no authentication done using TSIG keys, with the exception that static TSIG keys are not referenced by name.

Second, `libddns_gss_tsig.so` must be specified on the `hooks-libraries` list. This hook takes many parameters. The most important one is `servers`, which is a list of GSS-TSIG-capable servers. If there are several servers and they share some characteristics, the values can be specified in the `parameters` scope as defaults. In the example above, the

defaults that apply to all servers, unless otherwise specified on a per-server scope, are defined in lines 63 through 68. The defaults can be skipped if there is only one server defined, or if all servers have different values.

Table 1: List of available parameters

Name	Scope	Type	Default value	Description
client-keytab	global / server	string	empty	the Kerberos client key table
credentials-cache	global / server	string	empty	the Kerberos credentials cache
server-principal	global / server	string	empty	the Kerberos principal name of the DNS server that will receive updates
client-principal	global / server	string	empty	the Kerberos principal name of the Kea D2 service
gss-replay-flag	global / server	true / false	true	require the GSS anti replay service (GSS_C_REPLAY_FLAG)
gss-sequence-flag	global / server	true / false	false	require the GSS sequence service (GSS_C_SEQUENCE_FLAG)
tkey-protocol	global / server	string "TCP" / "UDP"		the protocol used to establish the security context with the DNS servers
tkey-lifetime	global / server	uint32	3600 seconds (1 hour)	the lifetime of GSS-TSIG keys
rekey-interval	global / server	uint32	2700 seconds (45 minutes)	the time interval the keys are checked for rekeying
retry-interval	global / server	uint32	120 seconds (2 minutes)	the time interval to retry to create a key if any error occurred previously
fallback	global / server	true / false	false	the behavior to fallback to non-GSS-TSIG when GSS-TSIG should be used but no GSS-TSIG key is available.
exchange-timeout	global / server	uint32	3000 milliseconds (3 seconds)	the time used to wait for the GSS-TSIG TKEY exchange to finish before it timeouts
user-context	global / server	string	empty	the user-provided data in JSON format (not used by the GSS-TSIG hook)
comment	global / server	string	empty	ignored
22.12. GSS-TSIG	server	string	empty	identifier to a DNS server (required)
domain-names	server	list of strings	empty	the many-to-one relationship between D2 DNS servers and

The global parameters are described below:

- **client-keytab** specifies the Kerberos **client** key table. For instance, `FILE:<filename>` can be used to point to a specific file. This parameter can be specified only once, in the parameters scope, and is the equivalent of setting the `KRB5_CLIENT_KTNAME` environment variable. An empty value is silently ignored.
- **credentials-cache** specifies the Kerberos credentials cache. For instance, `FILE:<filename>` can be used to point to a file or, if using a directory which supports more than one principal, `DIR:<directory-path>`. This parameter can be specified only once, in the parameters scope, and is the equivalent of setting the `KRB5CCNAME` environment variable. An empty value is silently ignored.
- **server-principal** is the Kerberos principal name of the DNS server that receives updates. In other words, this is the DNS server's name in the Kerberos system. This parameter is mandatory, and uses the typical Kerberos notation: `<SERVICE-NAME>/<server-domain-name>@<REALM>`.
- **client-principal** is the Kerberos principal name of the Kea D2 service. It is optional, and uses the typical Kerberos notation: `<SERVICE-NAME>/<server-domain-name>@<REALM>`.
- **gss-replay-flag** determines if the GSS anti replay service is required. It is by default but this can be disabled.
- **gss-sequence-flag** determines if the GSS sequence service is required. It is not by default but is required by the standard so it can be enabled.
- **tkey-protocol** determines which protocol is used to establish the security context with the DNS servers. Currently, the only supported values are TCP (the default) and UDP.
- **tkey-lifetime** determines the lifetime of GSS-TSIG keys in the TKEY protocol. The value must be greater than the **rekey-interval** value. It is expressed in seconds and defaults to 3600 (one hour).
- **rekey-interval** governs the time interval at which the keys for each configured server are checked for rekeying, i.e. when a new key is created to replace the current usable one if its age is greater than the **rekey-interval** value. The value must be smaller than the **tkey-lifetime** value (it is recommended to be set between 50% and 80% of the **tkey-lifetime** value). It is expressed in seconds and defaults to 2700 (45 minutes, or 75% of one hour).
- **retry-interval** governs the time interval at which to retry to create a key if any error occurred previously for any configured server. The value must be smaller than the **rekey-interval** value, and should be at most 1/3 of the difference between **tkey-lifetime** and **rekey-interval**. It is expressed in seconds and defaults to 120 (2 minutes).
- **fallback** governs the behavior when GSS-TSIG should be used (a matching DNS server is configured) but no GSS-TSIG key is available. If set to `false` (the default), this server is skipped; if set to `true`, the DNS server is ignored and the DNS update is sent with the configured DHCP-DDNS protection (e.g. TSIG key), or without any protection when none was configured.
- **exchange-timeout** governs the amount of time to wait for the GSS-TSIG TKEY exchange to finish before the process times out. It is expressed in milliseconds and defaults to 3000 (3 seconds).
- **user-context** is an optional parameter (see [Comments and User Context](#) for a general description of user contexts in Kea).
- **comment** is allowed but currently ignored.
- **servers** specifies the list of DNS servers where GSS-TSIG is enabled.

The server map parameters are described below:

- **id** assigns an identifier to a DNS server. It is used for statistics and commands. It is required, and must be both not empty and unique.
- **domain-names** governs the many-to-one relationship between D2 DNS servers and GSS-TSIG DNS servers: for each domain name on this list, Kea looks for a D2 DNS server for this domain with the specified IP address and port. An empty list (the default) means that all domains match.

- `ip-address` specifies the IP address at which the GSS-TSIG DNS server listens for DDNS and TKEY requests. It is a mandatory parameter.
- `port` specifies the DNS transport port on which the GSS-TSIG DNS server listens for DDNS and TKEY requests. It defaults to 53.
- `server-principal` is the Kerberos principal name of the DNS server that receives updates. The `server-principal` parameter set at the per-server level takes precedence over one set at the global level. It is a mandatory parameter which must be specified at either the global or the server level.
- `client-principal` is the Kerberos principal name of the Kea D2 service for this DNS server. The `client-principal` parameter set at the per-server level takes precedence over one set at the global level. It is an optional parameter.
- `gss-replay-flag` determines if the GSS anti replay service is required. The `gss-replay-flag` parameter set at the per-server level takes precedence over one set at the global level. It is an optional parameter which defaults to true.
- `gss-sequence-flag` determines if the GSS sequence service is required. The `gss-sequence-flag` parameter set at the per-server level takes precedence over one set at the global level. It is an optional parameter which defaults to false.
- `tkey-protocol` determines which protocol is used to establish the security context with the DNS server. The `tkey-protocol` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter.
- `tkey-lifetime` determines the lifetime of GSS-TSIG keys in the TKEY protocol for the DNS server. The `tkey-lifetime` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter.
- `rekey-interval` governs the time interval at which the keys for this particular server are checked for rekeying, i.e. when a new key is created to replace the current usable one if its age is greater than the `rekey-interval` value. The value must be smaller than the `tkey-lifetime` value (it is recommended to be set between 50% and 80% of the `tkey-lifetime` value). The `rekey-interval` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter.
- `retry-interval` governs the time interval at which to retry to create a key if any error occurred previously for this particular server. The value must be smaller than the `rekey-interval` value, and should be at most 1/3 of the difference between `tkey-lifetime` and `rekey-interval`. The `retry-interval` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter.
- `fallback` governs the behavior when GSS-TSIG should be used (a matching DNS server is configured) but no GSS-TSIG key is available. The `fallback` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter..
- `exchange-timeout` governs the amount of time to wait for the GSS-TSIG TKEY exchange to finish before the process times out. The `exchange-timeout` parameter set at the per-server level takes precedence over one set at the global level. The default and supported values for the per-server level parameter are the same as for the global-level parameter.
- `user-context` is an optional parameter (see [Comments and User Context](#) for a general description of user contexts in Kea).
- `comment` is allowed but currently ignored.

Note

Generally it is not recommended to specify both the client keytab (`client-keytab`) and the credentials cache (`credentials-cache`), although this may differ between Kerberos implementations. The client keytab is just for the client key and is typically used to specify the key explicitly in more static manner, while the credentials cache can be used to store multiple credentials and can be dynamically updated by the Kerberos library. As such, the `credentials-cache` is more flexible and thus the recommended alternative.

Also note that only the read access right is needed to use the cache. Fetching credentials and updating the cache requires the write access right.

22.2.4.1 GSS-TSIG Automatic Key Removal

The server periodically deletes keys after they have been expired more than three times the length of the maximum key lifetime (the `tkey-lifetime` parameter). The user has the option to purge keys on demand by using the `gss-tsig-purge-all` command or the `gss-tsig-purge` command.

22.2.4.2 GSS-TSIG Configuration for Deployment

When using Kerberos 5 and BIND 9 as described in *GSS-TSIG Deployment*, the local resolver must point to the BIND 9 named server address. The local Kerberos must also be configured by putting the following text into the `krb5.conf` file:

```
[libdefaults]
    default_realm = EXAMPLE.ORG
    kdc_timesync = 1
    ccache_type = 4
    forwardable = true
    proxiable = true
[realms]
    EXAMPLE.ORG = {
        kdc = kdc.example.org
        admin_server = kdc.example.org
    }
```

With Windows AD, the DNS service is provided by AD, which also provides the Kerberos service. The required text in the `krb5.conf` file becomes:

```
[libdefaults]
    default_realm = <REALM>
    kdc_timesync = 1
    ccache_type = 4
    forwardable = true
    proxiable = true
[realms]
    ${REALM} = {
        kdc = <AD_IP_ADDR>
        admin_server = <AD_IP_ADDR>
    }
```

Even when the GSS-API library can use the secret from the client key table, it is far better for performance to get and cache credentials.

This can be done manually via the command:

```
kinit -k -t /usr/local/share/kea/dhcp.keytab DHCP/admin.example.org
```

or, when using AD:

```
kinit -k -t /usr/local/share/kea/dhcp.keytab DHCP/kea.<domain>
```

The credential cache can be displayed using `klist`.

In production, it is better to rely on a Kerberos Credential Manager as the System Security Services Daemon (`sssd`).

When using BIND 9, the server principal is in the form `"DNS/server.example.org@EXAMPLE.ORG"`; with AD, the format is `"DNS/<server>.<domain>@<REALM>"`.

22.2.5 GSS-TSIG Statistics

The GSS-TSIG hook library introduces new statistics at global and per-DNS-server levels:

- `gss-tsig-key-created` - the number of created GSS-TSIG keys
- `tkey-sent` - the number of sent TKEY exchange initial requests
- `tkey-success` - the number of TKEY exchanges which completed with a success
- `tkey-timeout` - the number of TKEY exchanges which completed on timeout
- `tkey-error` - the number of TKEY exchanges which completed with an error other than a timeout

The relationship between keys and DNS servers is very different between the D2 code and static TSIG keys, and GSS-TSIG keys and DNS servers:

- a static TSIG key can be shared between many DNS servers;
- a GSS-TSIG key is only used by one DNS server inside a dedicated set of keys.

22.2.6 GSS-TSIG Commands

The GSS-TSIG hook library supports some commands, which are described below.

22.2.6.1 The `gss-tsig-get-all` Command

This command lists all the GSS-TSIG servers and keys.

An example command invocation looks like this:

```
{
  "command": "gss-tsig-get-all"
}
```

Here is an example of a response returning one GSS-TSIG server and one key:

```
{
  "result": 0,
  "text": "1 GSS-TSIG servers and 1 keys",
  "arguments": {
    "gss-tsig-servers": [
      {
        "id": "foo",
        "ip-address": "192.1.2.3",
        "port": 53,
```

(continues on next page)

(continued from previous page)

```

        "server-principal": "DNS/foo.com@FOO.COM",
        "key-name-suffix": "foo.com.",
        "tkey-lifetime": 3600,
        "tkey-protocol": "TCP",
        "keys": [
            {
                "name": "1234.sig-foo.com.",
                "inception-date": "2021-09-05 12:23:36.281176",
                "server-id": "foo",
                "expire-date": "2021-09-05 13:23:36.281176",
                "status": "not yet ready",
                "tkey-exchange": true
            }
        ]
    },
    {
        "id": "bar",
        "ip-address": "192.1.2.4",
        "port": 53,
        "server-principal": "DNS/bar.com@FOO.COM",
        "key-name-suffix": "bar.com.",
        "tkey-lifetime": 7200,
        "tkey-protocol": "UDP",
        "keys": [ ]
    }
]
}

```

22.2.6.2 The gss-tsig-get Command

This command retrieves information about the specified GSS-TSIG server.

An example command invocation looks like this:

```

{
    "command": "gss-tsig-get",
    "arguments": {
        "server-id": "foo"
    }
}

```

Here is an example of a response returning information about the server "foo":

```

{
    "result": 0,
    "text": "GSS-TSIG server[foo] found",
    "arguments": {
        "id": "foo",
        "ip-address": "192.1.2.3",
        "port": 53,
        "server-principal": "DNS/foo.com@FOO.COM",
        "key-name-suffix": "foo.com.",

```

(continues on next page)

(continued from previous page)

```

    "tkey-lifetime": 3600,
    "tkey-protocol": "TCP",
    "keys": [
      {
        "name": "1234.sig-foo.com.",
        "server-id": "foo",
        "inception-date": "2021-09-05 12:23:36.281176",
        "expire-date": "2021-09-05 13:23:36.281176",
        "status": "not yet ready",
        "tkey-exchange": true
      }
    ]
  }
}

```

22.2.6.3 The gss-tsig-list Command

This command generates a list of GSS-TSIG server IDs and key names.

An example command invocation looks like this:

```

{
  "command": "gss-tsig-list"
}

```

Here is an example of a response returning two GSS-TSIG servers and three keys:

```

{
  "result": 0,
  "text": "2 GSS-TSIG servers and 3 keys",
  "arguments": {
    "gss-tsig-servers": [
      "foo",
      "bar"
    ],
    "gss-tsig-keys": [
      "1234.example.com.",
      "5678.example.com.",
      "43888.example.org."
    ]
  }
}

```

22.2.6.4 The gss-tsig-key-get Command

This command retrieves information about the specified GSS-TSIG key.

An example command invocation looks like this:

```

{
  "command": "gss-tsig-key-get",
  "arguments": {
    "key-name": "1234.sig-foo.com."
  }
}

```

(continues on next page)

(continued from previous page)

```
}  
}
```

Here is an example of a response returning information about GSS-TSIG key "1234.sig-foo.com.":

```
{  
  "result": 0,  
  "text": "GSS-TSIG key '1234.sig-foo.com.' found",  
  "arguments": {  
    "name": "1234.sig-foo.com.",  
    "server-id": "foo",  
    "inception-date": "2021-09-05 12:23:36.281176",  
    "expire-date": "2021-09-05 13:23:36.281176",  
    "status": "not yet ready",  
    "tkey-exchange": true  
  }  
}
```

22.2.6.5 The gss-tsig-key-expire Command

This command expires the specified GSS-TSIG key.

An example command invocation looks like this:

```
{  
  "command": "gss-tsig-key-expire",  
  "arguments": {  
    "key-name": "1234.sig-foo.com."  
  }  
}
```

Here is an example of a response indicating that GSS-TSIG key "1234.sig-foo.com." has been expired:

```
{  
  "result": 0,  
  "text": "GSS-TSIG key '1234.sig-foo.com.' expired"  
}
```

22.2.6.6 The gss-tsig-key-del Command

This command deletes the specified GSS-TSIG key.

An example command invocation looks like this:

```
{  
  "command": "gss-tsig-key-del",  
  "arguments": {  
    "key-name": "1234.sig-foo.com."  
  }  
}
```

Here is an example of a response indicating that GSS-TSIG key "1234.sig-foo.com." has been deleted:

```
{
  "result": 0,
  "text": "GSS-TSIG key '1234.sig-foo.com.' deleted"
}
```

22.2.6.7 The `gss-tsig-purge-all` Command

This command removes all unusable GSS-TSIG keys.

An example command invocation looks like this:

```
{
  "command": "gss-tsig-purge-all"
}
```

Here is an example of a response indicating that two GSS-TSIG keys have been purged:

```
{
  "result": 0,
  "text": "2 purged GSS-TSIG keys"
}
```

22.2.6.8 The `gss-tsig-purge` Command

This command removes unusable GSS-TSIG keys for the specified server.

An example command invocation looks like this:

```
{
  "command": "gss-tsig-purge",
  "arguments": {
    "server-id": "foo"
  }
}
```

Here is an example of a response indicating that two GSS-TSIG keys for server "foo" have been purged:

```
{
  "result": 0,
  "text": "2 purged keys for GSS-TSIG server[foo]"
}
```

22.2.6.9 The `gss-tsig-rekey-all` Command

This command unconditionally creates new GSS-TSIG keys (rekeys) for all DNS servers.

An example command invocation looks like this:

```
{
  "command": "gss-tsig-rekey-all"
}
```

Here is an example of a response indicating that a rekey was performed:


```
{
  "result": 0,
  "text": "rekeyed"
}
```

This command is useful when, for instance, the DHCP-DDNS server is reconnected to the network.

22.2.6.10 The gss-tsig-rekey Command

This command unconditionally creates new GSS-TSIG keys (rekeys) for a specified DNS server.

An example command invocation looks like this:

```
{
  "command": "gss-tsig-rekey",
  "arguments": {
    "server-id": "foo"
  }
}
```

Here is an example of a response indicating that a rekey was performed:

```
{
  "result": 0,
  "text": "GSS-TSIG server[foo] rekeyed"
}
```

This command is typically used when a DNS server has been rebooted, so that existing GSS-TSIG keys shared with this server can no longer be used.

22.3 RADIUS

22.3.1 RADIUS Overview

This hook library allows Kea to interact with two types of RADIUS services: access and accounting. The most common DHCP and RADIUS integration is done on the DHCP relay-agent level: DHCP clients send DHCP packets to DHCP relays; those relays contact the RADIUS server and either send the packet to the DHCP server or drop it, depending on the response. It does require DHCP relay hardware to support RADIUS communication, though, and even if the relay has the necessary support, it is often not flexible enough to send and receive additional RADIUS attributes. There is a more appealing alternative: to extend the DHCP server to talk to RADIUS directly. That is the goal of this library.

Note

This library can only be loaded by the *kea-dhcp4* or the *kea-dhcp6* process.

The major feature of this hook library is the ability to use RADIUS authorization. When a DHCP packet is received, the Kea server sends an Access-Request to the RADIUS server and waits for a response. The server then sends back either an Access-Accept with specific client attributes, or an Access-Reject. There are two cases supported here: first, the Access-Accept includes a Framed-IP-Address attribute (for DHCPv4) or a Framed-IPv6-Address attribute/Delegated-IPv6-Prefix (for DHCPv6), which are interpreted by Kea as instructions to assign the specified IPv4 or IPv6 address/IPv6 prefix. This effectively means RADIUS can act as an address-reservation database.

The second supported case is the ability to assign clients to specific pools based on a RADIUS response. In this case, the RADIUS server sends back an Access-Accept with a Framed-Pool attribute. For both DHCPv4 and DHCPv6, Kea

interprets this attribute as a client class. With the addition of the ability to limit access to pools to specific classes (see *Configuring Pools With Class Information*), RADIUS can be used to force the client to be assigned a dynamic address or prefix from a specific pool. Furthermore, the same mechanism can be used to control what kind of options the client gets if there are DHCP options specified for a particular class.

22.3.2 RADIUS Hook Library Configuration

The RADIUS hook is a library that must be loaded by either *kea-dhcp4* or *kea-dhcp6* servers. Unlike some other available hook libraries, this one takes many parameters. For example, this configuration can be used:

```
{
  "Dhcp4": {

    // Your regular DHCPv4 configuration parameters go here.

    "hooks-libraries": [
      {
        // Note that the RADIUS access service requires host-cache for
        // proper operation, so this library is loaded as well.
        "library": "/usr/local/lib/kea/hooks/libdhcp_host_cache.so"
      },
      {
        "library": "/usr/local/lib/kea/hooks/libdhcp_radius.so",
        "parameters": {

          // Specify where the dictionary is located.
          "dictionary": "/etc/kea/radius/dictionary",

          // Specify which address to use to communicate with RADIUS servers
          "bindaddr": "*"

          // More RADIUS parameters go here.
        }
      }
    ]
  }
}
```

RADIUS is a complicated environment. As such, it is not feasible to provide a default configuration that works for everyone. However, we do have an example that showcases some of the more common features; please see `doc/examples/kea4/hooks-radius.json` in the Kea sources.

The RADIUS hook library supports the following global configuration flags:

- `bindaddr` (default `"*"`) - specifies the address to be used by the hook library in communication with the RADIUS servers. The `"*"` special value tells the kernel to choose the address at hook library load time.
- `canonical-mac-address` (default `false`) - specifies whether MAC addresses in attributes follow the canonical RADIUS format (lowercase pairs of hexadecimal digits separated by `-`).
- `client-id-pop0` (default `false`) - is used with *libdhcp_flex_id.so*; it removes the leading zero (or pair of zeroes in DHCPv6) type in the client ID (DUID in DHCPv6). See `client-id-printable` for any value implications when used in conjunction with it.
- `client-id-printable` (default `false`) - checks whether the `client-id/duid` content is printable and uses it as is instead of in hexadecimal. It implies that `client-id-pop0` and `extract-duid` as 0 and 255 are not printable.

- **deadtime** (default 0) - is a mechanism that helps in sorting the servers so those that have proved responsive so far are contacted first, and the servers that have proved unresponsive are left to try later. The deadtime value specifies the number of seconds after which a server is considered unresponsive. 0 disables the mechanism.
- **dictionary** (default "/etc/kea/radius/dictionary") - is the attribute and value dictionary; note that it is a critical parameter. A dictionary is provided by Kea and is set by default.
- **extract-duid** (default true) - extracts the embedded DUID from an RFC 4361-compliant DHCPv4 client ID. See **client-id-printable** for any value implications when used in conjunction with it.
- **identifier-type4** (default "client-id") - specifies the identifier type to build the User-Name attribute for DHCPv4; it should be the same as the host identifier. When [libdhcp_flex_id.so](#) is used, then **replace-client-id** must be set to true and **client-id** must be used with **client-id-pop0** enabled.
- **identifier-type6** (default "duid") - specifies the identifier type to build the User-Name attribute for DHCPv6; it should be the same as the host identifier. When [libdhcp_flex_id.so](#) is used, then **replace-client-id** must be set to true and **duid** must be used with **client-id-pop0** enabled.
- **nas-ports** (default []), specifies the NAS port to use in place of a subnet ID (default behavior). It is an array of maps, each map having two elements at most: the mandatory NAS port value, and optionally, a selector consisting of either a subnet ID, a subnet prefix, or a shared-network name. If the selector is applied to the packet, the NAS port is used instead of the subnet ID. When the subnet ID is 0 or missing, the specified NAS port acts as a default. The substitution happens for all packets that did not match a selector.
- **realm** (default "") - is the default realm.
- **reselect-subnet-address** (default false) - enables subnet reselection according to the value of the Framed-IP-Address or, respectively, the Framed-IPv6-Address attribute from the RADIUS access response. With this flag enabled, if the IP address is not in the range of the currently selected subnet, but is in the range of another subnet that is selectable with regards to other criteria, the latter subnet is selected and used further in the lease assignment process.
- **reselect-subnet-pool** (default false) - enables subnet reselection according to the value of the Framed-Pool attribute from the RADIUS access response. With this flag enabled, if the currently selected subnet is not guarded by the client class represented by the attribute value, but there is another selectable subnet that is guarded by it, the latter subnet is selected and used further in the lease assignment process. This reselection is attempted first, and if successful, it prevents the function of **reselect-subnet-address** from coming into effect.
- **retries** (default 3) - is the number of retries before trying the next server.
- **session-history** (default "") - is the name of the file providing persistent storage for accounting session history.
- **thread-pool-size** (default 0) indicates the number of threads that are used for sending RADIUS requests and processing RADIUS responses for both access and accounting services before passing them to the core thread pool. A value of 0 instructs the RADIUS hook library to use the same number of threads used for core DHCP processing. This value is only relevant if the Kea core is configured as multi-threaded; a single-threaded Kea core results in single-threaded RADIUS processing.
- **timeout** (default 10) - is the number of seconds during which a response is awaited.

Two services are supported:

- **access** - the authorization service.
- **accounting** - the accounting service.

At the service level, three sections can be configured:

- Servers that define RADIUS services that the library is expected to contact. Each server may have the following items specified:

- **name** - specifies the IP address of the server. A domain name may be used and will be resolved at hook library load time.
- **port** - specifies the UDP port of the server. By default, it is 1812 for access and 1813 for accounting.
- **secret** - authenticates messages.

When no server is specified, the service is disabled.

- Attributes which define additional information that the Kea server sends to a RADIUS server. The parameter must be identified either by a name or type. Its value can be specified in one of three possible ways: **data** (which defines a plain text value), **raw** (which defines the value in hex), or **expr** (which defines an expression that is evaluated for each incoming packet independently).
 - **name** - is the name of the attribute.
 - **type** - is the type of the attribute. Either the type or the name must be provided, and the attribute must be defined in the dictionary.
 - **data** - is the first of three ways to specify the attribute content. It specifies the textual representation of the attribute content.
 - **raw** - is the second of three ways to specify the attribute content. It specifies the content in hexadecimal.
 - **expr** - is the last of the three ways to specify the attribute content. It specifies an evaluation expression on the DHCP query packet.

Attributes are supported only for the access service.

- The **peer-updates** boolean flag (default **true**) controls whether lease updates coming from an active High-Availability (HA) partner should result in an accounting request. This may be desirable to remove duplicates if HA partners are configured to send requests to the same RADIUS server. The flag is only supported by the accounting service. The lease synchronization process at the startup of an HA node does not trigger a RADIUS accounting request, regardless of the value of this flag.
- The **max-pending-requests** positive integer (default **0**) limits the number of pending RADIUS requests. It is supported only by the access service. The value **0** means no limit; **64** is a recommended setting.

For example, to specify a single access server available on localhost that uses "1234" as a secret, and tell Kea to send three additional attributes (User-Password, Connect-Info, and Configuration-Token), the following snippet could be used:

```
{
  "parameters": {

    // Other RADIUS parameters here

    "access": {

      // This starts the list of access servers.
      "servers": [
        {
          // These are parameters for the first (and only) access server
          "name": "127.0.0.1",
          "port": 1812,
          "secret": "1234"
        }
      ]
      // Additional access servers could be specified here.
    ],

    // This defines a list of additional attributes Kea will send to each
```

```

// access server in Access-Request.
"attributes": [
  {
    // This attribute is identified by name (must be present in the
    // dictionary) and has static value (i.e. the same value will be
    // sent to every server for every packet).
    "name": "User-Password",
    "data": "mysecretpassword"
  },
  {
    // It is also possible to specify an attribute using its type,
    // rather than a name. 77 is Connect-Info. The value is specified
    // using hex. Again, this is a static value. It will be sent the
    // same for every packet and to every server.
    "type": 77,
    "raw": "65666a6a71"
  },
  {
    // This example shows how an expression can be used to send dynamic value.
    // The expression (see Using Expressions in Classification) may take any
    // value from the incoming packet or even its metadata e.g. the
    // interface it was received over from.
    "name": "Configuration-Token",
    "expr": "hexstring(pkt4.mac,':')"
  }
] // End of attributes
}, // End of access

// Accounting parameters.
"accounting": {
  // This starts the list of accounting servers.
  "servers": [
    {
      // These are parameters for the first (and only) accounting server
      "name": "127.0.0.1",
      "port": 1813,
      "secret": "1234"
    }
  ] // Additional accounting servers could be specified here.
}
}
}

```

Customization is sometimes required for certain attributes by devices belonging to various vendors. This is a great way to leverage the expression evaluation mechanism. For example, MAC addresses which might be used as a convenience value for the User-Password attribute are most likely to appear in colon-hexadecimal notation (de:ad:be:ef:ca:fe), but they might need to be expressed in hyphen-hexadecimal notation (de-ad-be-ef-ca-fe). Here's how to specify that:

```

{
  "parameters": {
    "access": {
      "attributes": [

```

(continues on next page)

(continued from previous page)

```

    {
      "name": "User-Password",
      "expr": "hexstring(pkt4.mac, '-')"
    }
  ]
}

```

And here's how to specify period-separated hexadecimal notation (`dead.beef.cafe`), preferred by Cisco devices:

```

{
  "parameters": {
    "access": {
      "attributes": [
        {
          "name": "User-Password",
          "expr": "substring(hexstring(pkt4.mac, ''), 0, 4) + '.' +
↳ substring(hexstring(pkt4.mac, ''), 4, 4) + '.' + substring(hexstring(pkt4.mac, ''), 8,
↳ 4)"
        }
      ]
    }
  }
}

```

For the access service of *libdhcp_radius.so* to operate properly, *libdhcp_host_cache.so* must also be loaded. The reason for this is somewhat complex. In a typical deployment, DHCP clients send their packets via DHCP relay, which inserts certain Relay Agent Information options, such as `circuit-id` or `remote-id`. The values of those options are then used by the Kea DHCP server to formulate the necessary attributes in the Access-Request message sent to the RADIUS server. However, once the DHCP client gets its address, it then renews by sending packets directly to the DHCP server. The relays cannot insert their RAI options at that point, and the DHCP server cannot send the Access-Request queries to the RADIUS server by using just the information from incoming packets. Kea needs to keep the information received during the initial Discover/Offer exchanges and use it again later when sending accounting messages.

This mechanism is implemented based on user context in host reservations. (See *Comments and User Context* and *User Contexts in Hooks* for details.) The host-cache mechanism allows the information retrieved by RADIUS to be stored and used later instead of doing another authorization round-trip to the RADIUS server. In other words, the host-cache mechanism is mandatory for the access service. The Kea configuration is rejected if *libdhcp_host_cache.so* is not loaded alongside *libdhcp_radius.so*.

Note

Currently the RADIUS hook library is incompatible with the `early-global-reservations-lookup` global parameter, i.e. setting the parameter to `true` raises an error when the hook library is loaded.

Note

Currently the RADIUS hook library is incompatible with the multi-subnet shared networks that have host reservations other than global. Loading the RADIUS hook library in a Kea DHCP server that has this configuration raises

an error.

Note

Currently the `reselect-subnet-address` and `reselect-subnet-pool` parameters are incompatible with the hub-and-spoke configuration of the *libdhcp_ha.so*. This hook library uses an originally selected subnet for choosing an HA relationship to process a packet. The subnet reselection may interfere with this choice. See the *Hub and Spoke Configuration* for details.

22.3.3 RADIUS Server Setup Example

The RADIUS hook library requires at least one RADIUS server to function. One popular open source implementation is FreeRADIUS; here's how to set it up to enable basic functionality in Kea.

1. Install FreeRADIUS through the package manager or from the tarballs available on [the freeradius.org download page](<https://freeradius.org/releases/>).
2. Establish the FreeRADIUS configuration directory. It's commonly `/etc/freeradius`, but it may be `/etc/raddb`.
3. Generate certificates. Go to `/etc/freeradius/certs`. Run `./bootstrap` or `make`. Wait until finished. It should take a few seconds.
4. Check that the server is able to start. Running with the `-X` flag is a good way to display potential errors. Run `radiusd -X` or `freeradius -X`, whichever is available. It should display `Ready to process requests` on the final line.
5. If the Kea DHCP server and the RADIUS server are on different machines, edit `/etc/freeradius/clients.conf` with the proper address under `ipaddr`. This file is also where the secret is set; it needs to match the one set in the hook library's configuration.
6. If RADIUS is used to authorize DHCP clients, each DHCP client needs to have an entry in the `authorize` file, which can be commonly found at:
 - `/etc/raddb/mods-config/files/authorize`
 - `/etc/freeradius/3.0/mods-config/files/authorize`
 - `/etc/freeradius/users` (for RADIUS 2.x series)

The passwords for entries must match the passwords in the configuration of the RADIUS hook library under the `User-Password` attribute. Each entry can have zero or more attributes.

In the following example, there are six entries with the password set to the client ID, which needs to be dynamically set in the hook library's configuration. Here's how the entries might look:

```
01:00:0c:01:02:03:04  Cleartext-password := "00:0c:01:02:03:04"

01:00:0c:01:02:03:05  Cleartext-password := "00:0c:01:02:03:05"
    Framed-IP-Address = "192.0.2.5"

01:00:0c:01:02:03:06  Cleartext-password := "00:0c:01:02:03:06"
    Framed-IP-Address = "192.0.2.6",
    Framed-Pool = "classical"

00:03:00:01:00:0c:01:02:03:07  Cleartext-password := "00:0c:01:02:03:07"
```

(continues on next page)

(continued from previous page)

```
00:03:00:01:00:0c:01:02:03:08    Cleartext-password := "00:0c:01:02:03:08"
    Framed-IPv6-Address = "2001:db8::8"

00:03:00:01:00:0c:01:02:03:09    Cleartext-password := "00:0c:01:02:03:09"
    Framed-IPv6-Address = "2001:db8::9",
    Framed-Pool = "classroom"
```

7. Accounting does not need to be modified to work with Kea, but customizations are possible in the accounting file, which can commonly be found at:

- /etc/radius-config/mods-config/files/accounting
- /etc/freeradius/3.0/mods-config/files/accounting

22.3.4 RADIUS Workflows for Lease Allocation

The following diagrams show a high-level view of how RADIUS assists with the lease allocation process in *kea-dhcp4* and *kea-dhcp6*.

Somewhat tangential to lease allocation, and not shown in the diagrams above, is the `command_processed` callout, which sends Accounting-Request messages when a lease command is received.



22.3.4.1 Parked-Packet Limit

Refer to *Parked-Packet Limit* for a basic introduction to packet parking.

The RADIUS hook library uses this mechanism. To allow for asynchronous communication between Kea and the RADIUS server and concurrent processing of DHCP packets by the Kea server, the DHCP request is parked, before the access request is sent on the subnet select callout. When the access response becomes available to the Kea DHCP server, the request is unparked, and the server continues processing it.

22.3.5 Differences Between RADIUS Hook Libraries Prior To Kea 2.4.0 and As Of 2.6.0

The RADIUS hook library in 2.4.0 and prior versions relied on the FreeRADIUS client library to function. Starting with 2.6.0 and onward, the RADIUS hook library is standalone, with its own RADIUS client implementation and its own RADIUS dictionary. There are differences:



Feature	Old	New
Support for Attribute Data Types	string, ipaddr, ipv4prefix, integer, integer64, date, ifid, ipv6addr, ipv6prefix, tlv, abinary, byte, ether, short, signed, octets	string (can simulate any other unsupported data type too), ipaddr, integer, date (interpreted as integer), ipv6addr, ipv6prefix
Names of Standard Attributes	Taken from the FreeRADIUS dictionary.	Taken from the Kea RADIUS dictionary and the IANA registry. There is an aliasing mechanism built into the library that ensures backward compatibility, e.g. <code>Password</code> translates to the standard name of the attribute <code>User-Password</code> .
Resolution of RADIUS Server Domain Names	At run time.	At hook library load time.
Automatic Deduction of Source Address for Reaching RADIUS Servers (configured with <code>bindaddr: "*"</code>)	At run time.	At hook library load time.
RADIUS Server Limit per Service	8	Unlimited
Support for Including Dictionaries Inside Dictionaries	Yes	No
Support for Vendor Attributes	Yes	No
Attribute Names and Attribute Values	Case-insensitive	Case-sensitive
Integer Values	Do not require an attribute definition.	Must have an associated attribute definition in the dictionary.
Reply-Message Presence in the Kea Logs	Only as part of the aggregated list of attributes in <code>RADIUS_AUTHENTICATION_ACCEPT</code> , <code>RADIUS_ACCESS_CACHE_INSERT</code> , <code>RADIUS_ACCESS_CACHE_GET</code> log messages.	Also has a dedicated <code>RADIUS_REPLY_MESSAGE_ATTRIBUTE</code> message per each Reply-Message attribute logged after a valid RADIUS reply is received.
Behavior of Multiple Attributes of the Same Type (except Reply-Message)	Experimentally, only the first attribute on the wire from an Access-Accept message is considered.	Experimentally, only the last attribute on the wire from an Access-Accept message is considered.

MONITORING KEA WITH STORK

Most administrators want to be able to monitor any Kea services that are running. Kea offers so many pieces of information - configuration files, API, statistics, logs, open database content, and more - that it may sometimes be overwhelming to keep up. ISC's Stork project is intended to address this problem for both Kea and BIND 9. Stork is useful in a variety of ways:

- Stork can be used as a dashboard. It provides insight into what exactly is happening on the servers. In particular, it allows users to: see up-to-date details regarding pool utilization in subnets and shared networks; monitor the state of the HA pair (and provide extra insight in case of failover and recovery events); list, filter, and search for specific host reservations; and more. Only a single Stork server needs to be deployed, and one Stork agent on each machine to be monitored.
- The Stork agent can integrate Kea with Prometheus and Grafana. Once the Stork agent is active on the server, it serves as a Prometheus exporter. Users who have deployed Prometheus in their networks can visualize statistics as time series using Grafana.
- Stork can act as both a dashboard and an integrator for Prometheus/Grafana. Once Stork is linked to where Grafana is deployed on the network, users can inspect the current status and visit a customized link to Grafana to see how a given property behaves over time.

Stork is available as source code, but also as native deb and RPM packages, which makes it easy to install on most popular systems. For more details, please see the [Stork ARM](#) or the [Stork project page](#). The ARM has a nice collection of screenshots that is frequently updated, to give users an idea of what is currently available. Stork is in the midst of full development with monthly releases, so please check back frequently.

Note

Stork uses Kea API calls from the `libdhcp_host_cmds.so` and `libdhcp_subnet_cmds.so` hooks to make changes to Kea's configuration. Such usage is incompatible with the *Kea Configuration Backend* (CB). Stork does not currently implement support for the `libdhcp_cb_cmds.so` needed to make changes to CB configuration. If using the Kea CB, Stork can only be safely used in a "read-only" fashion. Support for the CB is planned for a future release of Stork.

23.1 Kea Statistics in Grafana

The ISC Stork project provides an agent that can be deployed alongside Kea. It exposes Kea statistics in a format that is accepted by Prometheus. One of the major benefits of Prometheus is that it turns repeated one-time observations into time series, which lets users monitor how certain behaviors change over time. It is easy to use other tools to visualize data available in Prometheus; the most common approach is to use Grafana to provide visual dashboards. The Stork project provides dashboard definitions for Kea that can be imported into Grafana very easily.

Learn more about Prometheus and Grafana on their websites: *Prometheus* <<https://prometheus.io/>> and *Grafana* <<https://grafana.com/>>.

API REFERENCE

Kea currently supports 201 commands in *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6* daemons and *cb_cmds*, *class_cmds*, *gss_tsig*, *high_availability*, *host_cache*, *host_cmds*, *lease_cmds*, *lease_query*, *perfmon*, *stat_cmds*, *subnet_cmds* hook libraries.

Commands supported by *kea-ctrl-agent* daemon: *build-report*, *config-get*, *config-hash-get*, *config-reload*, *config-set*, *config-test*, *config-write*, *list-commands*, *shutdown*, *status-get*, *version-get*.

Commands supported by *kea-dhcp-ddns* daemon: *build-report*, *config-get*, *config-hash-get*, *config-reload*, *config-set*, *config-test*, *config-write*, *gss-tsig-get*, *gss-tsig-get-all*, *gss-tsig-key-del*, *gss-tsig-key-expire*, *gss-tsig-key-get*, *gss-tsig-list*, *gss-tsig-purge*, *gss-tsig-purge-all*, *gss-tsig-rekey*, *gss-tsig-rekey-all*, *list-commands*, *shutdown*, *statistic-get*, *statistic-get-all*, *statistic-reset*, *statistic-reset-all*, *status-get*, *version-get*.

Commands supported by *kea-dhcp4* daemon: *build-report*, *cache-clear*, *cache-flush*, *cache-get*, *cache-get-by-id*, *cache-insert*, *cache-load*, *cache-remove*, *cache-size*, *cache-write*, *class-add*, *class-del*, *class-get*, *class-list*, *class-update*, *config-backend-pull*, *config-get*, *config-hash-get*, *config-reload*, *config-set*, *config-test*, *config-write*, *dhcp-disable*, *dhcp-enable*, *extended-info4-upgrade*, *ha-continue*, *ha-heartbeat*, *ha-maintenance-cancel*, *ha-maintenance-notify*, *ha-maintenance-start*, *ha-reset*, *ha-scopes*, *ha-sync*, *ha-sync-complete-notify*, *lease4-add*, *lease4-del*, *lease4-get*, *lease4-get-all*, *lease4-get-by-client-id*, *lease4-get-by-hostname*, *lease4-get-by-hw-address*, *lease4-get-page*, *lease4-resend-ddns*, *lease4-update*, *lease4-wipe*, *lease4-write*, *leases-reclaim*, *list-commands*, *network4-add*, *network4-del*, *network4-get*, *network4-list*, *network4-subnet-add*, *network4-subnet-del*, *perfmon-control*, *perfmon-get-all-durations*, *remote-class4-del*, *remote-class4-get*, *remote-class4-get-all*, *remote-class4-set*, *remote-global-parameter4-del*, *remote-global-parameter4-get*, *remote-global-parameter4-get-all*, *remote-global-parameter4-set*, *remote-network4-del*, *remote-network4-get*, *remote-network4-list*, *remote-network4-set*, *remote-option-def4-del*, *remote-option-def4-get*, *remote-option-def4-get-all*, *remote-option-def4-set*, *remote-option4-global-del*, *remote-option4-global-get*, *remote-option4-global-get-all*, *remote-option4-global-set*, *remote-option4-network-del*, *remote-option4-network-set*, *remote-option4-pool-del*, *remote-option4-pool-set*, *remote-option4-subnet-del*, *remote-option4-subnet-set*, *remote-server4-del*, *remote-server4-get*, *remote-server4-get-all*, *remote-server4-set*, *remote-subnet4-del-by-id*, *remote-subnet4-del-by-prefix*, *remote-subnet4-get-by-id*, *remote-subnet4-get-by-prefix*, *remote-subnet4-list*, *remote-subnet4-set*, *reservation-add*, *reservation-del*, *reservation-get*, *reservation-get-all*, *reservation-get-by-address*, *reservation-get-by-hostname*, *reservation-get-by-id*, *reservation-get-page*, *reservation-update*, *server-tag-get*, *shutdown*, *stat-lease4-get*, *statistic-get*, *statistic-get-all*, *statistic-remove*, *statistic-remove-all*, *statistic-reset*, *statistic-reset-all*, *statistic-sample-age-set*, *statistic-sample-age-set-all*, *statistic-sample-count-set*, *statistic-sample-count-set-all*, *status-get*, *subnet4-add*, *subnet4-del*, *subnet4-delta-add*, *subnet4-delta-del*, *subnet4-get*, *subnet4-list*, *subnet4-select-test*, *subnet4-update*, *subnet4o6-select-test*, *version-get*.

Commands supported by *kea-dhcp6* daemon: *build-report*, *cache-clear*, *cache-flush*, *cache-get*, *cache-get-by-id*, *cache-insert*, *cache-load*, *cache-remove*, *cache-size*, *cache-write*, *class-add*, *class-del*, *class-get*, *class-list*, *class-update*, *config-backend-pull*, *config-get*, *config-hash-get*, *config-reload*, *config-set*, *config-test*, *config-write*, *dhcp-disable*, *dhcp-enable*, *extended-info6-upgrade*, *ha-continue*, *ha-heartbeat*, *ha-maintenance-cancel*, *ha-maintenance-notify*, *ha-maintenance-start*, *ha-reset*, *ha-scopes*, *ha-sync*, *ha-sync-complete-notify*, *lease6-add*, *lease6-bulk-apply*, *lease6-del*, *lease6-get*, *lease6-get-all*, *lease6-get-by-duid*, *lease6-get-by-hostname*, *lease6-get-page*, *lease6-resend-ddns*, *lease6-update*, *lease6-wipe*, *lease6-write*, *leases-reclaim*, *list-commands*, *network6-add*, *network6-del*, *network6-get*, *network6-list*, *network6-subnet-add*, *network6-subnet-del*, *perfmon-control*,

perfmmon-get-all-durations, remote-class6-del, remote-class6-get, remote-class6-get-all, remote-class6-set, remote-global-parameter6-del, remote-global-parameter6-get, remote-global-parameter6-get-all, remote-global-parameter6-set, remote-network6-del, remote-network6-get, remote-network6-list, remote-network6-set, remote-option-def6-del, remote-option-def6-get, remote-option-def6-get-all, remote-option-def6-set, remote-option6-global-del, remote-option6-global-get, remote-option6-global-get-all, remote-option6-global-set, remote-option6-network-del, remote-option6-network-set, remote-option6-pd-pool-del, remote-option6-pd-pool-set, remote-option6-pool-del, remote-option6-pool-set, remote-option6-subnet-del, remote-option6-subnet-set, remote-server6-del, remote-server6-get, remote-server6-get-all, remote-server6-set, remote-subnet6-del-by-id, remote-subnet6-del-by-prefix, remote-subnet6-get-by-id, remote-subnet6-get-by-prefix, remote-subnet6-list, remote-subnet6-set, reservation-add, reservation-del, reservation-get, reservation-get-all, reservation-get-by-address, reservation-get-by-hostname, reservation-get-by-id, reservation-get-page, reservation-update, server-tag-get, shutdown, stat-lease6-get, statistic-get, statistic-get-all, statistic-remove, statistic-remove-all, statistic-reset, statistic-reset-all, statistic-sample-age-set, statistic-sample-age-set-all, statistic-sample-count-set, statistic-sample-count-set-all, status-get, subnet6-add, subnet6-del, subnet6-delta-add, subnet6-delta-del, subnet6-get, subnet6-list, subnet6-select-test, subnet6-update, version-get.

Commands supported by *cb_cmds* hook library: *remote-class4-del, remote-class4-get, remote-class4-get-all, remote-class4-set, remote-class6-del, remote-class6-get, remote-class6-get-all, remote-class6-set, remote-global-parameter4-del, remote-global-parameter4-get, remote-global-parameter4-get-all, remote-global-parameter4-set, remote-global-parameter6-del, remote-global-parameter6-get, remote-global-parameter6-get-all, remote-global-parameter6-set, remote-network4-del, remote-network4-get, remote-network4-list, remote-network4-set, remote-network6-del, remote-network6-get, remote-network6-list, remote-network6-set, remote-option-def4-del, remote-option-def4-get, remote-option-def4-get-all, remote-option-def4-set, remote-option-def6-del, remote-option-def6-get, remote-option-def6-get-all, remote-option-def6-set, remote-option4-global-del, remote-option4-global-get, remote-option4-global-get-all, remote-option4-global-set, remote-option4-network-del, remote-option4-network-set, remote-option4-pool-del, remote-option4-pool-set, remote-option4-subnet-del, remote-option4-subnet-set, remote-option6-global-del, remote-option6-global-get, remote-option6-global-get-all, remote-option6-global-set, remote-option6-network-del, remote-option6-network-set, remote-option6-pd-pool-del, remote-option6-pd-pool-set, remote-option6-pool-del, remote-option6-pool-set, remote-option6-subnet-del, remote-option6-subnet-set, remote-server4-del, remote-server4-get, remote-server4-get-all, remote-server4-set, remote-server6-del, remote-server6-get, remote-server6-get-all, remote-server6-set, remote-subnet4-del-by-id, remote-subnet4-del-by-prefix, remote-subnet4-get-by-id, remote-subnet4-get-by-prefix, remote-subnet4-list, remote-subnet4-set, remote-subnet6-del-by-id, remote-subnet6-del-by-prefix, remote-subnet6-get-by-id, remote-subnet6-get-by-prefix, remote-subnet6-list, remote-subnet6-set.*

Commands supported by *class_cmds* hook library: *class-add, class-del, class-get, class-list, class-update.*

Commands supported by *gss_tsig* hook library: *gss-tsig-get, gss-tsig-get-all, gss-tsig-key-del, gss-tsig-key-expire, gss-tsig-key-get, gss-tsig-list, gss-tsig-purge, gss-tsig-purge-all, gss-tsig-rekey, gss-tsig-rekey-all.*

Commands supported by *high_availability* hook library: *ha-continue, ha-heartbeat, ha-maintenance-cancel, ha-maintenance-notify, ha-maintenance-start, ha-reset, ha-scopes, ha-sync, ha-sync-complete-notify.*

Commands supported by *host_cache* hook library: *cache-clear, cache-flush, cache-get, cache-get-by-id, cache-insert, cache-load, cache-remove, cache-size, cache-write.*

Commands supported by *host_cmds* hook library: *reservation-add, reservation-del, reservation-get, reservation-get-all, reservation-get-by-address, reservation-get-by-hostname, reservation-get-by-id, reservation-get-page, reservation-update.*

Commands supported by *lease_cmds* hook library: *lease4-add, lease4-del, lease4-get, lease4-get-all, lease4-get-by-client-id, lease4-get-by-hostname, lease4-get-by-hw-address, lease4-get-page, lease4-resend-ddns, lease4-update, lease4-wipe, lease4-write, lease6-add, lease6-bulk-apply, lease6-del, lease6-get, lease6-get-all, lease6-get-by-duid, lease6-get-by-hostname, lease6-get-page, lease6-resend-ddns, lease6-update, lease6-wipe, lease6-write.*

Commands supported by *lease_query* hook library: *extended-info4-upgrade, extended-info6-upgrade.*

Commands supported by *perfmmon* hook library: *perfmmon-control, perfmmon-get-all-durations.*

Commands supported by *stat_cmds* hook library: *stat-lease4-get, stat-lease6-get.*

Commands supported by *subnet_cmds* hook library: *network4-add*, *network4-del*, *network4-get*, *network4-list*, *network4-subnet-add*, *network4-subnet-del*, *network6-add*, *network6-del*, *network6-get*, *network6-list*, *network6-subnet-add*, *network6-subnet-del*, *subnet4-add*, *subnet4-del*, *subnet4-delta-add*, *subnet4-delta-del*, *subnet4-get*, *subnet4-list*, *subnet4-update*, *subnet6-add*, *subnet6-del*, *subnet6-delta-add*, *subnet6-delta-del*, *subnet6-get*, *subnet6-list*, *subnet6-update*.

24.1 build-report

This command returns the list of compilation options that this particular binary was built with.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *build-report command*

Command syntax:

```
{
  "command": "build-report"
}
```

Response syntax:

```
{
  "result": 0,
  "text": <string with build details>
}
```

24.2 cache-clear

This command removes all cached host reservations.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-clear command*

Command syntax:

```
{
  "command": "cache-clear"
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.3 cache-flush

This command removes up to the given number or all cached host reservations.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-flush command*

Command syntax:

```
{
  "command": "cache-flush",
  "arguments": 5
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.4 cache-get

This command returns the full content of the host cache.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *cache-get command*

Command syntax:

```
{
  "command": "cache-get"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "123 entries returned.",
  "arguments": <list of host reservations>
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.5 cache-get-by-id

This command returns entries matching the given identifier from the host cache.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (*host_cache* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *cache-get-by-id command*

Command syntax:

```
{
  "command": "cache-get-by-id",
  "arguments": {
    "hw-address": "01:02:03:04:05:06"
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "2 entries returned.",
  "arguments": <list of host reservations>
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.6 cache-insert

This command inserts a host into the cache.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-insert command*

Command syntax:

```
{
  "command": "cache-insert",
  "arguments": {
    "hw-address": "01:02:03:04:05:06",
    "subnet-id4": 4,
    "subnet-id6": 0,
    "ip-address": "192.0.2.100",
    "hostname": "somehost.example.org",
    "client-classes4": [ ],
    "client-classes6": [ ],
    "option-data4": [ ],
    "option-data6": [ ],
    "next-server": "192.0.0.2",
    "server-hostname": "server-hostname.example.org",
    "boot-file-name": "bootfile.efi",
    "host-id": 0
  }
}
```

Another example that adds IPv6 host reservation to cache **is**:

```
{
  "command": "cache-insert",
  "arguments": {
    "hw-address": "01:02:03:04:05:06",
    "subnet-id4": 0,
    "subnet-id6": 6,
    "ip-addresses": [ "2001:db8::cafe:babe" ],
    "prefixes": [ "2001:db8:dead:beef::/64" ],
    "hostname": "",
    "client-classes4": [ ],
    "client-classes6": [ ],
    "option-data4": [ ],
    "option-data6": [ ],
    "next-server": "0.0.0.0",
    "server-hostname": ""
  }
}
```

(continues on next page)

(continued from previous page)

```

    "boot-file-name": "",
    "host-id": 0
  }
}

```

Response syntax:

```

{
  "result": <integer>,
  "text": "<string>"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.7 cache-load

This command allows the contents of a file on disk to be loaded into an in-memory cache.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-load command*

Command syntax:

```

{
  "command": "cache-load",
  "arguments": "/usr/local/share/kea/kea-host-cache.json"
}

```

Response syntax:

```

{
  "result": <integer>,
  "text": "<string>"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)

- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.8 cache-remove

This command removes entries from the host cache. It takes parameters similar to the `reservation-get` command.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-remove command*

Command syntax:

```
{
  "command": "cache-remove",
  "arguments": {
    "ip-address": "192.0.2.1",
    "subnet-id": 123
  }
}
```

Another example that removes the IPv6 host identifier by DUID **and** specific subnet-id is:

```
{
  "command": "cache-remove",
  "arguments": {
    "duid": "00:01:ab:cd:f0:a1:c2:d3:e4",
    "subnet-id": 123
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.9 cache-size

This command returns the number of entries in the host cache.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (*host_cache* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *cache-size command*

Command syntax:

```
{
  "command": "cache-size"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "123 entries.",
  "arguments": { "size": 123 }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.10 cache-write

This command instructs Kea to write its host cache content to disk.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*host_cache* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *cache-write command*

Command syntax:

```
{
  "command": "cache-write",
  "arguments": "/path/to/the/file.json"
}
```

The command takes one mandatory argument that specifies the filename path of a file to be written.

Response syntax:

```
{  
  "result": <integer>,  
  "text": "<string>"  
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.11 class-add

This command adds a new class to the existing server configuration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.5.0 (*class_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *class-add command*

Command syntax:

```
{  
  "command": "class-add",  
  "arguments": {  
    "client-classes": [ {  
      "name": <name of the class>,  
      "test": <test expression to be evaluated on incoming packets>,  
      "option-data": [ <option values here> ],  
      "option-def": [ <option definitions here> ],  
      "next-server": <ipv4 address>,  
      "server-hostname": <string>,  
      "boot-file-name": <name of the boot file>  
    } ]  
  }  
}
```

The *next-server*, *server-hostname*, and *boot-file-name* are DHCPv4-specific. Only one client class can be added with a single command.

Response syntax:

```
{  
  "result": 0,  
  "text": "Class '<class-name>' added."  
}
```

The command is successful (result 0), unless the class name is a duplicate or another error occurs (result 1).

24.12 class-del

This command removes a client class from the server configuration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.5.0 (*class_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *class-del command*

Command syntax:

```
{
  "command": "class-del",
  "arguments": {
    "name": <name of the class>
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Class '<class-name>' deleted."
}
```

The command returns a result of 3 (empty) if the client class does not exist. If the client class exists, the returned result is 0 if the deletion was successful; the result is 1 if the deletion is unsuccessful.

24.13 class-get

This command returns detailed information about an existing client class.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.5.0 (*class_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *class-get command*

Command syntax:

```
{
  "command": "class-get",
  "arguments": {
    "name": <name of the class>
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Class '<class-name>' definition returned",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "client-classes": [
      {
        "name": <name of the class>,
        "only-in-additional-list": <only in additional list boolean value>,
        "test": <test expression to be evaluated on incoming packets>,
        "option-data": [ <option values here> ],
        "option-def": [ <option definitions here> ],
        "next-server": <ipv4 address>,
        "server-hostname": <string>,
        "boot-file-name": <name of the boot file>
      }
    ]
  }
}

```

The returned information depends on the DHCP server type, i.e. some parameters are specific to the DHCPv4 server. Also, some parameters may not be returned if they are not set for the client class. If a class with the specified name does not exist, a result of 3 (empty) is returned. If the client class is found, the result of 0 is returned. If there is an error while processing the command, the result of 1 is returned.

24.14 class-list

This command retrieves a list of all client classes from the server configuration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.5.0 (*class_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *class-list command*

Command syntax:

```

{
  "command": "class-list"
}

```

This command includes no arguments.

Response syntax:

```

{
  "result": 0,
  "text": "'<number of>' classes found",
  "arguments": {
    "client-classes": [
      {
        "name": <first class name>
      },
      {
        "name": <second class name>
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

The returned list of classes merely contains their names. In order to retrieve full information about one of these classes, use *The class-get Command*. The returned result is 3 (empty) if no classes are found. If the command is processed successfully and the list of client classes is not empty, the result of 0 is returned. If there is an error while processing the command, the result of 1 is returned.

24.15 class-update

This command updates an existing client class in the server configuration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.5.0 (*class_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *class-update command*

Command syntax:

```
{
  "command": "class-update",
  "arguments": {
    "client-classes": [ {
      "name": <name of the class>,
      "test": <test expression to be evaluated on incoming packets>,
      "option-data": [ <option values here> ],
      "option-def": [ <option definitions here> ],
      "next-server": <ipv4 address>,
      "server-hostname": <string>,
      "boot-file-name": <name of the boot file>
    } ]
  }
}
```

The *next-server*, *server-hostname*, and *boot-file-name* are DHCPv4-specific. Only one client class can be updated with a single command.

Response syntax:

```
{
  "result": 0,
  "text": "Class '<class-name>' updated."
}
```

The command returns the result of 3 (empty) if the client class does not exist. If the client class exists, the returned result is 0 if the update was successful, or 1 if the update is unsuccessful.

24.16 config-backend-pull

This command forces an immediate update of the server using Config Backends. This command does not take any parameters.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.1 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *config-backend-pull command*

Command syntax:

```
{  
  "command": "config-backend-pull"  
}
```

Response syntax:

```
{  
  "result": 0,  
  "text": "On demand configuration update successful."  
}
```

When no Config Backends are configured this command returns empty (3); If an error occurs error (1) is returned with the error details; otherwise success (0) is returned.

24.17 config-get

This command retrieves the current configuration used by the server. The configuration is essentially the same as the contents of the configuration file, but includes additional changes made by other commands and due to parameters' inheritance.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *config-get command*

Command syntax:

```
{  
  "command": "config-get"  
}
```

This command takes no parameters.

Response syntax:

```
{  
  "result": <integer>,  
  "arguments": {  
    <Dhcp4, Dhcp6, or Control-agent object>: <JSON configuration here>  
  }  
}
```

Starting with Kea 2.4.0, the successful response contains an SHA256 digest of the configuration that was just retrieved. It might be used to determine if a configuration has been modified or not, possibly with the use of *config-hash-get*.

24.18 config-hash-get

This command retrieves the hash of the current configuration used by the server. The configuration is essentially the same as the contents of the configuration file, but includes additional changes made by other commands and due to parameters' inheritance. Currently it is SHA256, but the algorithm may change in the future.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 2.4.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *config-hash-get command*

Command syntax:

```
{
  "command": "config-hash-get"
}
```

This command takes no parameters.

Response syntax:

```
{
  "result": <integer>,
  "arguments": {
    "hash": <SHA256 hash of the config in hexa>
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.19 config-reload

This command instructs Kea to reload the configuration file that was used previously.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *config-reload command*

Command syntax:

```
{
  "command": "config-reload"
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.20 config-set

This command instructs the server to replace its current configuration with the new configuration supplied in the command's arguments.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *config-set command*

Command syntax:

```
{
  "command": "config-set",
  "arguments": {
    "'<server>': {
      }
  }
}
```

In the example below, '<server>' is the configuration element name for a given server such as "Dhcp4" or "Dhcp6".

Response syntax:

```
{
  "arguments": {
    "hash": "8B5F5822E93178B65CE658304C37EA511BDE7D29F792AA5E88012FC741F2BE32"
  },
  "result": 0,
  "text": "Configuration successful."
}

or

{ "result": 1, "text": "unsupported parameter: BOGUS (<string>:16:26)" }
```

Starting with Kea 2.4.0, the successful response from a DHCPv4, DHCPv6, or DHCP-DDNS daemons also contain a SHA-256 digest of the newly set configuration. The digest can be used to easily determine if a configuration has been modified or not, possibly with the use of `config-hash-get`.

24.21 config-test

This command instructs the server to check whether the new configuration supplied in the command's arguments can be loaded.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *config-test command*

Command syntax:

```
{
  "command": "config-test",
  "arguments": {
    "'<server>': {
      }
  }
}
```

In the example below, '`<server>`' is the configuration element name for a given server such as "Dhcp4" or "Dhcp6".

Response syntax:

```
{ "result": 0, "text": "Configuration seems sane." }

or

{ "result": 1, "text": "unsupported parameter: BOGUS (<string>:16:26)" }
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.22 config-write

This command instructs the Kea server to write its current configuration to a file on disk.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *config-write command*

Command syntax:

```
{
  "command": "config-write",
  "arguments": {
    "filename": "config-modified-2017-03-15.json"
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.23 dhcp-disable

This command disables the DHCP service.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *dhcp-disable command*

Command syntax:

```
{
  "command": "dhcp-disable",
  "arguments": {
    "max-period": 20,
    "origin": "ha-partner",
    "origin-id": 2002
  }
}
```

Response syntax:

```
{
  "result": <integer>,

```

(continues on next page)

(continued from previous page)

```
{
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.24 dhcp-enable

This command enables the DHCP service.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *dhcp-enable command*

Command syntax:

```
{
  "command": "dhcp-enable",
  "arguments": {
    "origin": "ha-partner",
    "origin-id": 2002
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.25 extended-info4-upgrade

This command sanitizes extended info of all IPv4 leases and fills relay and remote ID columns in the SQL lease database.

Supported by: *kea-dhcp4*

Availability: 2.3.8 (*lease_query* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *extended-info4-upgrade command*

Command syntax:

```
{  
  "command": "extended-info4-upgrade"  
}
```

Response syntax:

```
{  
  "arguments": {  
  },  
  "result": 0,  
  "text": "Upgraded 1000 leases"  
}
```

This command should be used when some old IPv4 leases are present in the lease database using a SQL backend.

24.26 extended-info6-upgrade

This command sanitizes extended info of all IPv6 leases and rebuilds when they are enabled relay and remote ID tables in the SQL lease database.

Supported by: *kea-dhcp6*

Availability: 2.5.0 (*lease_query* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *extended-info6-upgrade command*

Command syntax:

```
{  
  "command": "extended-info6-upgrade"  
}
```

Response syntax:

```
{  
  "arguments": {  
  },  
  "result": 0,  
  "text": "Upgraded 1000 leases"  
}
```

This command should be used when some old IPv6 leases are present in the lease database using a SQL backend to make them visible to the .

24.27 gss-tsig-get

This command retrieves information about the specified GSS-TSIG server.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-get command*

Command syntax:

```
{
  "command": "gss-tsig-get",
  "arguments": {
    "server-id": "foo"
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "GSS-TSIG server[foo] found",
  "arguments": {
    "id": "foo",
    "ip-address": "192.1.2.3",
    "port": 53,
    "server-principal": "DNS/foo.com@FOO.COM",
    "key-name-suffix": "foo.com.",
    "tkey-lifetime": 3600,
    "tkey-protocol": "TCP",
    "keys": [
      {
        "name": "1234.sig-foo.com.",
        "server-id": "foo",
        "inception-date": "2021-09-05 12:23:36.281176",
        "expire-date": "2021-09-05 13:23:36.281176",
        "status": "not yet ready",
        "tkey-exchange": true
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.28 gss-tsig-get-all

This command lists GSS-TSIG servers and keys.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-get-all command*

Command syntax:

```
{  
  "command": "gss-tsig-get-all"  
}
```

Response syntax:

```
{  
  "result": 0,  
  "text": "1 GSS-TSIG servers and 1 keys",  
  "arguments": {  
    "gss-tsig-servers": [  
      {  
        "id": "foo",  
        "ip-address": "192.1.2.3",  
        "port": 53,  
        "server-principal": "DNS/foo.com@FOO.COM",  
        "key-name-suffix": "foo.com.",  
        "tkey-lifetime": 3600,  
        "tkey-protocol": "TCP",  
        "keys": [  
          {  
            "name": "1234.sig-foo.com.",  
            "inception-date": "2021-09-05 12:23:36.281176",  
            "server-id": "foo",  
            "expire-date": "2021-09-05 13:23:36.281176",  
            "status": "not yet ready",  
            "tkey-exchange": true  
          }  
        ]  
      },  
      {  
        "id": "bar",  
        "ip-address": "192.1.2.4",  
        "port": 53,  
        "server-principal": "DNS/bar.com@FOO.COM",  
        "key-name-suffix": "bar.com.",  
        "tkey-lifetime": 7200,  
        "tkey-protocol": "UDP",  
        "keys": [ ]  
      }  
    ]  
  }  
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.29 gss-tsig-key-del

This command deletes the specified GSS-TSIG key.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-key-del command*

Command syntax:

```
{
  "command": "gss-tsig-key-del",
  "arguments": {
    "key-name": "1234.sig-foo.com."
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "GSS-TSIG key '1234.sig-foo.com.' deleted"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.30 gss-tsig-key-expire

This command expires the specified GSS-TSIG key.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-key-expire command*

Command syntax:

```
{
  "command": "gss-tsig-key-expire",
  "arguments": {
    "key-name": "1234.sig-foo.com."
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "GSS-TSIG key '1234.sig-foo.com.' expired"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.31 gss-tsig-key-get

This command retrieves information about the specified GSS-TSIG key.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-key-get command*

Command syntax:

```
{
  "command": "gss-tsig-key-get",
  "arguments": {
    "key-name": "1234.sig-foo.com."
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "GSS-TSIG key '1234.sig-foo.com.' found",
  "arguments": {
    "name": "1234.sig-foo.com.",
    "server-id": "foo",
    "inception-date": "2021-09-05 12:23:36.281176",
    "expire-date": "2021-09-05 13:23:36.281176",
    "status": "not yet ready",
    "tkey-exchange": true
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.32 gss-tsig-list

This command lists GSS-TSIG server IDs and key names.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-list command*

Command syntax:

```
{
  "command": "gss-tsig-list"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "2 GSS-TSIG servers and 3 keys",
  "arguments": {
    "gss-tsig-servers": [
      "foo",
      "bar"
    ],
    "gss-tsig-keys": [
      "1234.example.com.",

```

(continues on next page)

(continued from previous page)

```
        "5678.example.com.",  
        "43888.example.org."  
    ]  
}  
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.33 gss-tsig-purge

This command removes not usable GSS-TSIG keys for the specified server.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-purge command*

Command syntax:

```
{  
    "command": "gss-tsig-purge",  
    "arguments": {  
        "server-id": "foo"  
    }  
}
```

Response syntax:

```
{  
    "result": 0,  
    "text": "2 purged keys for GSS-TSIG server[foo]"  
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.34 gss-tsig-purge-all

This command removes not usable GSS-TSIG keys.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-purge-all command*

Command syntax:

```
{
  "command": "gss-tsig-purge-all"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "2 purged GSS-TSIG keys"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.35 gss-tsig-rekey

The command unconditionally creates new GSS-TSIG keys for (rekeys) a specified DNS server.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-rekey command*

Command syntax:

```
{
  "command": "gss-tsig-rekey",
  "arguments": {
    "server-id": "foo"
  }
}
```

Response syntax:


```
{
  "result": 0,
  "text": "GSS-TSIG server[foo] rekeyed"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.36 gss-tsig-rekey-all

This command unconditionally creates new GSS-TSIG keys (rekeys) for all DNS servers.

Supported by: *kea-dhcp-ddns*

Availability: 2.0.0 (*gss_tsig* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *gss-tsig-rekey-all command*

Command syntax:

```
{
  "command": "gss-tsig-rekey-all"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "rekeyed"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.37 ha-continue

This command resumes the operation of a paused HA state machine.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *ha-continue command*

Command syntax:

```
{
  "command": "ha-continue",
  "arguments": {
    "server-name": "server1"
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.38 ha-heartbeat

This command is sent internally by a Kea partner when operating in High Availability (HA) mode or by the system administrator to verify the state of the servers with regards to the High Availability. It retrieves the server's HA state and clock value.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *ha-heartbeat command*

Command syntax:

```
{
  "command": "ha-heartbeat",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "server-name": "server1"
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "HA peer status returned.",
  "arguments": {
    "state": <server state>,
    "date-time": <server notion of time>,
    "scopes": [ <first scope>, <second scope>, ... ],
    "unsent-update-count": <total number of lease allocations in partner-down state>
  }
}

```

The response includes a server state (see [Server States](#)), current clock value, served scopes and the counter indicating how many leases the server has allocated without sending lease updates to its partner. The partner uses this counter to determine if it should synchronize its lease database.

24.39 ha-maintenance-cancel

This command is sent to instruct a server in the partner-in-maintenance state to transition back to the previous state, effectively canceling the maintenance.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.4 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [ha-maintenance-cancel command](#)

Command syntax:

```

{
  "command": "ha-maintenance-cancel"
}

```

This command takes no arguments.

Response syntax:

```

{
  "result": <integer>,
  "text": "<string>"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)

- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.40 ha-maintenance-notify

This command is sent by the server receiving the `ha-maintenance-start` to its partner to cause the partner to transition to the in-maintenance state or to revert it from the in-maintenance state as a result of receiving the `ha-maintenance-cancel` command.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.4 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *ha-maintenance-notify command*

Command syntax:

```
{
  "command": "ha-maintenance-notify",
  "arguments": {
    "cancel": true,
    "state": "ready"
  }
}
```

This command includes a boolean argument which, if false, indicates that the server should transition to the in-maintenance state. If the cancel argument is set to true it instructs the server to revert from the in-maintenance state to its previous state. The state argument informs the recipient about the state of the sending server. This command is not meant to be used by the administrator. It is merely used for internal communication between the HA partners.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

The response may include a special error code of 1001 to indicate that the partner refused to enter the maintenance state.

24.41 ha-maintenance-start

This command is sent to instruct one of the servers to transition to the partner-in-maintenance state in which it will be responding to all DHCP queries. The server receiving this command sends the `ha-maintenance-notify` to its partner to cause the partner to transition to the in-maintenance state.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.4 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *ha-maintenance-start command*

Command syntax:

```
{
  "command": "ha-maintenance-start"
}
```

This command takes no arguments.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.42 ha-reset

This command resets the HA state machine of the server by transitioning it to the waiting state.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.9.4 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [ha-reset command](#)

Command syntax:

```
{
  "command": "ha-reset",
  "arguments": {
    "server-name": "server1"
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.43 ha-scopes

This command modifies the scope that the server is responsible for serving when operating in High Availability (HA) mode.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *ha-scopes command*

Command syntax:

```
{
  "command": "ha-scopes",
  "arguments": {
    "scopes": [ "HA_server1", "HA_server2" ],
    "server-name": "server1"
  }
}
```

In the example below, the arguments configure the server to handle traffic from both the HA_server1 and HA_server2 scopes.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.44 ha-sync

This command instructs the server running in HA mode to synchronize its local lease database with the selected peer.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.4.0 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [ha-sync command](#)

Command syntax:

```
{
  "command": "ha-sync",
  "arguments": {
    "server-name": "server1",
    "max-period": 60
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.45 ha-sync-complete-notify

A server notifies its partner with this command that it has finished the lease database synchronization. If the partner is in the partner-down state it temporarily stops allocating new leases until it transitions to a normal operation state (e.g. load-balancing). If the partner observes a failing heartbeat it can resume allocating new leases in the partner-down state.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.9.11 (*high_availability* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [ha-sync-complete-notify command](#)

Command syntax:

```
{
  "command": "ha-sync-complete-notify",
  "arguments": {
    "server-name": "server1",
    "origin": 2001,
    "max-period": 60
  }
}
```

This command takes no arguments.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.46 lease4-add

This command administratively adds a new IPv4 lease.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-add command*

Command syntax:

```
{
  "command": "lease4-add",
  "arguments": {
    "ip-address": "192.0.2.202",
    "hw-address": "1a:1b:1c:1d:1e:1f"
  }
}
```

Note that Kea 1.4 requires an additional argument, subnet-ID, which is optional as of Kea 1.5. A number of other, more-detailed, optional arguments are also supported.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

If the returned result is equal to 4, it indicates that the lease could not be created because it was in conflict with the server's state or its notion of the configuration. The High Availability hook library can handle such a result differently than a general error. A general error of 1 can indicate issues with processing the command, database availability etc.

24.47 lease4-del

This command deletes a lease from the lease database.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-del command*

Command syntax:

```
{
  "command": "lease4-del",
  "arguments": {
    "ip-address": "192.0.2.202"
  }
}
```

The lease to be deleted can be specified either by IP address or by identifier-type and identifier value. The currently supported identifiers are "hw-address" and "client-id".

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.48 lease4-get

This command queries the lease database and retrieves existing leases.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get command*

Command syntax:

```
{
  "command": "lease4-get",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "ip-address": "192.0.2.1"
  }
}

```

Response syntax:

```

{
  "arguments": {
    "client-id": "42:42:42:42:42:42:42:42",
    "cltt": 12345678,
    "fqdn-fwd": false,
    "fqdn-rev": true,
    "hostname": "myhost.example.com.",
    "hw-address": "08:08:08:08:08:08",
    "ip-address": "192.0.2.1",
    "state": 0,
    "subnet-id": 44,
    "valid-lft": 3600
  },
  "result": 0,
  "text": "IPv4 lease found."
}

```

lease4-get returns a result that indicates the outcome of the operation and lease details, if found. It has one of the following values: 0 (success), 1 (error), or 3 (empty). Result 3 is returned if no leases are found with specified IP address.

24.49 lease4-get-all

This command retrieves all IPv4 leases or all leases for the specified set of subnets.

Supported by: *kea-dhcp4*

Availability: 1.4.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get-all command*

Command syntax:

```

{
  "command": "lease4-get-all",
  "arguments": {
    "subnets": [ 1, 2, 3, ... ]
  }
}

```

The lease4-get-all command may result in very large responses. Please consider using lease4-get-page to get them in chunks. The subnets parameter is optional. If not specified, leases from all subnets are returned.

Response syntax:

```

{
  "arguments": {

```

(continues on next page)

(continued from previous page)

```

"leases": [
  {
    "client-id": "01:00:0c:01:02:03:04",
    "cltt": 1600432232,
    "fqdn-fwd": false,
    "fqdn-rev": false,
    "hostname": "",
    "hw-address": "00:0c:01:02:03:04",
    "ip-address": "192.168.1.150",
    "state": 0,
    "subnet-id": 1,
    "valid-lft": 4000
  },
  {
    "client-id": "01:00:0c:01:02:03:05",
    "cltt": 1600432234,
    "fqdn-fwd": false,
    "fqdn-rev": false,
    "hostname": "",
    "hw-address": "00:0c:01:02:03:05",
    "ip-address": "192.168.1.151",
    "state": 0,
    "subnet-id": 1,
    "valid-lft": 4000
  }
]
},
"result": 0,
"text": "2 IPv4 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameter.

24.50 lease4-get-by-client-id

This command retrieves all IPv4 leases with the specified client id.

Supported by: *kea-dhcp4*

Availability: 1.7.1 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get-by-client-id command*

Command syntax:

```

{
  "command": "lease4-get-by-client-id",
  "arguments": {
    "client-id": "01:00:0c:01:02:03:04"
  }
}

```

Response syntax:

```
{
  "arguments": {
    "leases": [
      {
        "client-id": "01:00:0c:01:02:03:04",
        "cltt": 1600432232,
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "",
        "hw-address": "00:0c:01:02:03:04",
        "ip-address": "192.168.1.150",
        "state": 0,
        "subnet-id": 1,
        "valid-lft": 4000
      }
    ]
  },
  "result": 0,
  "text": "1 IPv4 lease(s) found."
}
```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found.

24.51 lease4-get-by-hostname

This command retrieves all IPv4 leases with the specified hostname.

Supported by: *kea-dhcp4*

Availability: 1.7.1 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get-by-hostname command*

Command syntax:

```
{
  "command": "lease4-get-by-hostname",
  "arguments": {
    "hostname": "myhost.example.com."
  }
}
```

Response syntax:

```
{
  "arguments": {
    "leases": [
      {
        "client-id": "01:00:0c:01:02:03:04",
        "cltt": 1600432232,
        "fqdn-fwd": false,
```

(continues on next page)

(continued from previous page)

```

        "fqdn-rev": false,
        "hostname": "myhost.example.com.",
        "hw-address": "00:0c:01:02:03:04",
        "ip-address": "192.168.1.150",
        "state": 0,
        "subnet-id": 1,
        "valid-lft": 4000
    }
],
"result": 0,
"text": "1 IPv4 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found.

24.52 lease4-get-by-hw-address

This command retrieves all IPv4 leases with the specified hardware address.

Supported by: *kea-dhcp4*

Availability: 1.7.1 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get-by-hw-address command*

Command syntax:

```

{
  "command": "lease4-get-by-hw-address",
  "arguments": {
    "hw-address": "00:0c:01:02:03:04"
  }
}

```

Response syntax:

```

{
  "arguments": {
    "leases": [
      {
        "client-id": "01:00:0c:01:02:03:04",
        "cltt": 1600432232,
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "myhost.example.com.",
        "hw-address": "00:0c:01:02:03:04",
        "ip-address": "192.168.1.150",
        "state": 0,
        "subnet-id": 1,
        "valid-lft": 4000
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ]
},
"result": 0,
"text": "1 IPv4 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found.

24.53 lease4-get-page

This command retrieves all IPv4 leases by page.

Supported by: *kea-dhcp4*

Availability: 1.5.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease4-get-page command*

Command syntax:

```

{
  "command": "lease4-get-page",
  "arguments": {
    "limit": <integer>,
    "from": <IPv4 address or 'start'>
  }
}

```

The from address and the page size limit are mandatory.

Response syntax:

```

{
  "arguments": {
    "leases": [
      {
        "client-id": "01:00:0c:01:02:03:04",
        "cltt": 1600432232,
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "",
        "hw-address": "00:0c:01:02:03:04",
        "ip-address": "192.168.1.150",
        "state": 0,
        "subnet-id": 1,
        "valid-lft": 4000
      },
      {
        "client-id": "01:00:0c:01:02:03:05",
        "cltt": 1600432234,
        "fqdn-fwd": false,

```

(continues on next page)

(continued from previous page)

```

        "fqdn-rev": false,
        "hostname": "",
        "hw-address": "00:0c:01:02:03:05",
        "ip-address": "192.168.1.151",
        "state": 0,
        "subnet-id": 1,
        "valid-lft": 4000
    }
  ],
  "result": 0,
  "text": "2 IPv4 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameters.

24.54 lease4-resend-ddns

This command resends a request to kea-dhcp-ddns to update DNS for an existing lease.

Supported by: *kea-dhcp4*

Availability: 1.7.6 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-resend-ddns command*

Command syntax:

```

{
  "command": "lease4-resend-ddns",
  "arguments": {
    "ip-address": "192.0.2.1"
  }
}

```

Response syntax:

```

{
  "arguments": {
  },
  "result": 0,
  "text": "NCR generated for: 192.0.2.1, hostname: example.com."
}

```

lease4-resend-ddns returns a result that indicates the outcome of the operation and lease details, if found. It has one of the following values: 0 (success), 1 (error), or 3 (empty).

24.55 lease4-update

This command updates existing leases.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-update command*

Command syntax:

```
{
  "command": "lease4-update",
  "arguments": {
    "ip-address": "192.0.2.1",
    "hostname": "newhostname.example.org",
    "hw-address": "1a:1b:1c:1d:1e:1f",
    "subnet-id": 44,
    "force-create": true
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

If the returned result is equal to 4, it indicates that the lease could not be updated because it was in conflict with the server's state or its notion of the configuration. The High Availability hook library can handle such a result differently than a general error. A general error of 1 can indicate issues with processing the command, database availability etc.

24.56 lease4-wipe

This command removes all leases associated with a given subnet.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-wipe command*

Command syntax:

```
{
  "command": "lease4-wipe",
  "arguments": {
    "subnet-id": 44
  }
}
```

Response syntax:


```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.57 lease4-write

This command writes the IPv4 memfile lease database into a CSV file.

Supported by: *kea-dhcp4*

Availability: 2.3.1 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease4-write command*

Command syntax:

```
{
  "command": "lease4-write",
  "arguments": {
    "filename": "a_file.csv"
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.58 lease6-add

This command administratively creates a new lease.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-add command*

Command syntax:

```
{
  "command": "lease6-add",
  "arguments": {
    "subnet-id": 66,
    "ip-address": "2001:db8::3",
    "duid": "1a:1b:1c:1d:1e:1f:20:21:22:23:24",
    "iaid": 1234
  }
}
```

lease6-add can be also used to add leases for IPv6 prefixes.

Response syntax:

```
{ "result": 0, "text": "Lease added." }
or
{ "result": 1, "text": "missing parameter 'ip-address' (<string>:3:19)" }
```

If the returned result is equal to 4, it indicates that the lease could not be created because it was in conflict with the server's state or its notion of the configuration. The High Availability hook library can handle such a result differently than a general error. A general error of 1 can indicate issues with processing the command, database availability etc.

24.59 lease6-bulk-apply

This command creates, updates, or deletes multiple IPv6 leases in a single transaction. It communicates lease changes between HA peers, but may be used in all cases where it is desirable to apply multiple lease updates in a single transaction.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-bulk-apply command*

Command syntax:

```
{
  "command": "lease6-bulk-apply",
  "arguments": {
    "deleted-leases": [
      {
        "ip-address": "2001:db8:abcd::",

```

(continues on next page)

(continued from previous page)

```

        "type": "IA_PD",
        ...
    },
    {
        "ip-address": "2001:db8:abcd::234",
        "type": "IA_NA",
        ...
    }
],
"leases": [
    {
        "subnet-id": 66,
        "ip-address": "2001:db8:cafe::",
        "type": "IA_PD",
        ...
    },
    {
        "subnet-id": 66,
        "ip-address": "2001:db8:abcd::333",
        "type": "IA_NA",
        ...
    }
]
}

```

If any of the leases is malformed, all changes are rolled back. If the leases are well-formed but the operation fails for one or more leases, these leases are listed in the response; however, the changes are preserved for all leases for which the operation was successful. The "deleted-leases" and "leases" are optional parameters, but one of them must be specified.

Response syntax:

```

{
    "result": 0,
    "text": "IPv6 leases bulk apply completed.",
    "arguments": {
        "failed-deleted-leases": [
            {
                "ip-address": "2001:db8:abcd::",
                "type": "IA_PD",
                "result": <control result>,
                "error-message": <error message>
            }
        ],
        "failed-leases": [
            {
                "ip-address": "2001:db8:cafe::",
                "type": "IA_PD",
                "result": <control result>,
                "error-message": <error message>
            }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

The "failed-deleted-leases" holds the list of leases which failed to delete; this includes leases which were not found in the database. The "failed-leases" includes the list of leases which failed to create or update. For each lease for which there was an error during processing, insertion into the database, etc., the result is set to 1. If an error occurs due to a conflict between the lease and the server's configuration or state, the result of 4 is returned instead of 1. For each lease which was not deleted because the server did not find it in the database, the result of 3 is returned.

24.60 lease6-del

This command deletes a lease from the lease database.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-del command*

Command syntax:

```
{
  "command": "lease6-del",
  "arguments": {
    "ip-address": "2001:db8::3"
  }
}
```

lease6-del returns a result that indicates the outcome of the operation. It has one of the following values: 0 (success), 1 (error), or 3 (empty).

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.61 lease6-get

This command queries the lease database and retrieves existing leases.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease6-get command*

Command syntax:

```
{
  "command": "lease6-get",
  "arguments": {
    "ip-address": "2001:db8:1234:ab::",
    "type": "IA_PD"
  }
}
```

lease6-get returns a result that indicates the outcome of the operation and lease details, if found.

Response syntax:

```
{
  "arguments": {
    "leases": [
      {
        "cltt": 1600439560,
        "duid": "00:01:00:01:26:f7:81:88:00:0c:01:02:03:04",
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "foobar.example.org",
        "hw-address": "00:0c:01:02:03:04",
        "iaid": 1,
        "ip-address": "2001:db8:1::",
        "preferred-lft": 3000,
        "state": 0,
        "subnet-id": 1,
        "type": "IA_NA",
        "valid-lft": 4000
      }
    ]
  },
  "result": 0,
  "text": "1 IPv6 lease(s) found."
}
```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameter.

24.62 lease6-get-all

This command retrieves all IPv6 leases or all leases for the specified set of subnets.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease6-get-all command*

Command syntax:

```
{
  "command": "lease6-get-all",
  "arguments": {
    "subnets": [ 1, 2, 3, 4 ]
  }
}
```

The `lease6-get-all` command may result in very large responses. Please consider using `lease6-get-page` to get them in chunks. the `subnets` parameter is optional. If not specified, leases from all subnets are returned.

Response syntax:

```
{
  "arguments": {
    "leases": [
      {
        "cltt": 12345678,
        "duid": "42:42:42:42:42:42:42:42",
        "fqdn-fwd": false,
        "fqdn-rev": true,
        "hostname": "myhost.example.com.",
        "hw-address": "08:08:08:08:08:08",
        "iaid": 1,
        "ip-address": "2001:db8:2::1",
        "preferred-lft": 500,
        "state": 0,
        "subnet-id": 44,
        "type": "IA_NA",
        "valid-lft": 3600
      },
      {
        "cltt": 12345678,
        "duid": "21:21:21:21:21:21:21:21",
        "fqdn-fwd": false,
        "fqdn-rev": true,
        "hostname": "",
        "iaid": 1,
        "ip-address": "2001:db8:0:0:2::",
        "preferred-lft": 500,
        "prefix-len": 80,
        "state": 0,
        "subnet-id": 44,

```

(continues on next page)

(continued from previous page)

```

        "type": "IA_PD",
        "valid-lft": 3600
    }
],
},
"result": 0,
"text": "2 IPv6 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameter.

24.63 lease6-get-by-duid

This command retrieves all IPv6 leases with the specified DUID.

Supported by: *kea-dhcp6*

Availability: 1.7.1 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease6-get-by-duid command*

Command syntax:

```

{
  "command": "lease6-get-by-duid",
  "arguments": {
    "duid": "1a:1b:1c:1d:1e:1f:20:21:22:23:24"
  }
}

```

Response syntax:

```

{
  "arguments": {
    "leases": [
      {
        "cltt": 1600439560,
        "duid": "00:01:00:01:26:f7:81:88:00:0c:01:02:03:04",
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "foobar.example.org",
        "hw-address": "00:0c:01:02:03:04",
        "iaid": 1,
        "ip-address": "2001:db8:1::",
        "preferred-lft": 3000,
        "state": 0,
        "subnet-id": 1,
        "type": "IA_NA",
        "valid-lft": 4000
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

},
"result": 0,
"text": "1 IPv6 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameter.

24.64 lease6-get-by-hostname

This command retrieves all IPv6 leases with the specified hostname.

Supported by: *kea-dhcp6*

Availability: 1.7.1 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease6-get-by-hostname command*

Command syntax:

```

{
  "command": "lease6-get-by-hostname",
  "arguments": {
    "hostname": "myhost.example.com."
  }
}

```

Response syntax:

```

{
  "arguments": {
    "leases": [
      {
        "cltt": 1600439560,
        "duid": "00:01:00:01:26:f7:81:88:00:0c:01:02:03:04",
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "foobar.example.org",
        "hw-address": "00:0c:01:02:03:04",
        "iaid": 1,
        "ip-address": "2001:db8:1::",
        "preferred-lft": 3000,
        "state": 0,
        "subnet-id": 1,
        "type": "IA_NA",
        "valid-lft": 4000
      }
    ]
  },
  "result": 0,
  "text": "1 IPv6 lease(s) found."
}

```


Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found with specified parameter.

24.65 lease6-get-page

This command retrieves all IPv6 leases by page.

Supported by: *kea-dhcp6*

Availability: 1.5.0 (*lease_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *lease6-get-page command*

Command syntax:

```
{
  "command": "lease6-get-page",
  "arguments": {
    "limit": <integer>,
    "from": <IPv6 address or 'start'>
  }
}
```

The from address and the page size limit are mandatory.

Response syntax:

```
{
  "arguments": {
    "leases": [
      {
        "cltt": 1600439560,
        "duid": "00:01:00:01:26:f7:81:88:00:0c:01:02:03:04",
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "foo.example.org",
        "hw-address": "00:0c:01:02:03:04",
        "iaid": 1,
        "ip-address": "2001:db8:1::1",
        "preferred-lft": 3000,
        "state": 0,
        "subnet-id": 1,
        "type": "IA_NA",
        "valid-lft": 4000
      },
      {
        "cltt": 1600439570,
        "duid": "00:01:00:01:26:f7:81:88:00:0c:01:02:03:05",
        "fqdn-fwd": false,
        "fqdn-rev": false,
        "hostname": "bar.example.org",
        "hw-address": "00:0c:01:02:03:05",
        "iaid": 1,
        "ip-address": "2001:db8:1::2",

```

(continues on next page)

(continued from previous page)

```

        "preferred-lft": 3000,
        "state": 0,
        "subnet-id": 1,
        "type": "IA_NA",
        "valid-lft": 4000
    }
]
},
"result": 0,
"text": "1 IPv6 lease(s) found."
}

```

Result 0 is returned when at least one lease is found, 1 when parameters are malformed or missing, 3 is returned if no leases are found.

24.66 lease6-resend-ddns

This command resends a request to kea-dhcp-ddns to update DNS for an existing lease.

Supported by: *kea-dhcp6*

Availability: 1.7.6 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-resend-ddns command*

Command syntax:

```

{
  "command": "lease6-resend-ddns",
  "arguments": {
    "ip-address": "2001:db8::1"
  }
}

```

Response syntax:

```

{
  "arguments": {
  },
  "result": 0,
  "text": "NCR generated for: 2001:db8::1, hostname: example.com."
}

```

lease6-resend-ddns returns a result that indicates the outcome of the operation and lease details, if found. It has one of the following values: 0 (success), 1 (error), or 3 (empty).

24.67 lease6-update

This command updates existing leases.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-update command*

Command syntax:

```
{
  "command": "lease6-update",
  "arguments": {
    "ip-address": "2001:db8::1",
    "duid": "88:88:88:88:88:88:88:88",
    "iaid": 7654321,
    "hostname": "newhostname.example.org",
    "subnet-id": 66,
    "force-create": false
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

If the returned result is equal to 4, it indicates that the lease could not be updated because it was in conflict with the server's state or its notion of the configuration. The High Availability hook library can handle such a result differently than a general error. A general error of 1 can indicate issues with processing the command, database availability etc.

24.68 lease6-wipe

This command removes all leases associated with a given subnet.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-wipe command*

Command syntax:

```
{
  "command": "lease6-wipe",
  "arguments": {
    "subnet-id": 66
  }
}
```

Note: not all backends support this command.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.69 lease6-write

This command writes the IPv6 memfile lease database into a CSV file.

Supported by: *kea-dhcp6*

Availability: 2.3.1 (*lease_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *lease6-write command*

Command syntax:

```
{
  "command": "lease6-write",
  "arguments": {
    "filename": "a_file.csv"
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.70 leases-reclaim

This command instructs the server to reclaim all expired leases immediately.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [leases-reclaim command](#)

Command syntax:

```
{
  "command": "leases-reclaim",
  "arguments": {
    "remove": true
  }
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.71 list-commands

This command retrieves a list of all commands supported by the server.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see [list-commands command](#)

Command syntax:

```
{
  "command": "list-commands",
  "arguments": { }
}
```

The server responds with a list of all supported commands.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.72 network4-add

This command adds a new shared network.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network4-add command*

Command syntax:

```
{
  "command": "network4-add",
  "arguments": {
    "shared-networks": [ {
      "name": "floor13",
      "subnet4": [
        {
          "id": 100,
          "pools": [ { "pool": "192.0.2.2-192.0.2.99" } ],
          "subnet": "192.0.2.0/24",
          "option-data": [
            {
              "name": "routers",
              "data": "192.0.2.1"
            }
          ]
        }
      ],
    },
    {
      "id": 101,
      "pools": [ { "pool": "192.0.3.2-192.0.3.99" } ],
      "subnet": "192.0.3.0/24",
      "option-data": [
        {
          "name": "routers",
          "data": "192.0.3.1"
        }
      ]
    }
  ]
} ]
}
```

Response syntax:

```
{
  "arguments": {
    "shared-networks": [ { "name": "floor13" } ]
  },
  "result": 0,
  "text": "A new IPv4 shared network 'floor13' added"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.73 network4-del

This command deletes existing shared networks.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network4-del command*

Command syntax:

```
{
  "command": "network4-del",
  "arguments": {
    "name": "floor13"
  }
}
```

Response syntax:

```
{
  "command": "network4-del",
  "arguments": {
    "shared-networks": [
      {
        "name": "floor13"
      }
    ]
  },
  "result": 0,
  "text": "IPv4 shared network 'floor13' deleted"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.74 network4-get

This command retrieves detailed information about shared networks, including subnets that are currently part of a given network.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *network4-get command*

Command syntax:

```
{
  "command": "network4-get",
  "arguments": {
    "name": "floor13"
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Info about IPv4 shared network 'floor13' returned",
  "arguments": {
    "shared-networks": [
      {
        "match-client-id": true,
        "name": "floor13",
        "option-data": [ ],
        "rebind-timer": 90,
        "relay": {
          "ip-address": "0.0.0.0"
        },
        "renew-timer": 60,
        "subnet4": [
          {
            "subnet": "192.0.2.0/24",
            "id": 5,
            <many other subnet specific details here>
          },
          {

```

(continues on next page)

(continued from previous page)

```

        "subnet": "192.0.3.0/31",
        "id": 6,
        <many other subnet specific details here>
    },
    ],
    "valid-lifetime": 120
}
]
}

```

Note that the actual response contains many additional fields that are omitted here for clarity.

24.75 network4-list

This command retrieves the full list of currently configured shared networks.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *network4-list command*

Command syntax:

```

{
  "command": "network4-list"
}

```

Response syntax:

```

{
  "arguments": {
    "shared-networks": [
      { "name": "floor1" },
      { "name": "office" }
    ]
  },
  "result": 0,
  "text": "2 IPv4 network(s) found"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.76 network4-subnet-add

This command adds existing subnets to existing shared networks.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network4-subnet-add command*

Command syntax:

```
{
  "command": "network4-subnet-add",
  "arguments": {
    "name": "floor13",
    "id": 5
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 subnet 10.0.0.0/8 (id 5) is now part of shared network 'floor1'"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.77 network4-subnet-del

This command removes a subnet that is part of an existing shared network and demotes it to a plain, stand-alone subnet.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network4-subnet-del command*

Command syntax:

```
{
  "command": "network4-subnet-del",
  "arguments": {
    "name": "floor13",
```

(continues on next page)

(continued from previous page)

```

    "id": 5
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "IPv4 subnet 10.0.0.0/8 (id 5) is now removed from shared network 'floor13'"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.78 network6-add

This command adds a new shared network.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network6-add command*

Command syntax:

```

{
  "command": "network6-add",
  "arguments": {
    "shared-networks": [ {
      "name": "floor13",
      "subnet6": [
        {
          "id": 100,
          "pools": [ { "pool": "2003:db8:1::1-2003:db8:1::ff" } ],
          "subnet": "2003:db8:1::/64",
          "option-data": [
            {
              "name": "dns-servers",
              "data": "2005:db8:1::1"
            }
          ]
        }
      ],
    },
    {
      "id": 101,

```

(continues on next page)

(continued from previous page)

```

        "pools": [ { "pool": "2003:db8:2::1-2003:db8:2::ff" } ],
        "subnet": "2003:db8:2::/64",
        "option-data": [
            {
                "name": "dns-servers",
                "data": "2006:db8:1::1"
            }
        ]
    } ]
}

```

Response syntax:

```

{
    "arguments": {
        "shared-networks": [ { "name": "floor13" } ]
    },
    "result": 0,
    "text": "A new IPv6 shared network 'floor13' added"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.79 network6-del

This command deletes existing shared networks.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network6-del command*

Command syntax:

```

{
    "command": "network6-del",
    "arguments": {
        "name": "floor13"
    }
}

```

Response syntax:

```
{
  "command": "network6-del",
  "arguments": {
    "shared-networks": [
      {
        "name": "floor13"
      }
    ]
  },
  "result": 0,
  "text": "IPv6 shared network 'floor13' deleted"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.80 network6-get

This command retrieves detailed information about shared networks, including subnets that are currently part of a given network.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *network6-get command*

Command syntax:

```
{
  "command": "network4-get",
  "arguments": {
    "name": "floor13"
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Info about IPv6 shared network 'floor13' returned",
  "arguments": {
    "shared-networks": [
      {
```

(continues on next page)

(continued from previous page)

```

    "match-client-id": true,
    "name": "floor13",
    "option-data": [ ],
    "rebind-timer": 90,
    "relay": {
        "ip-address": "::"
    },
    "renew-timer": 60,
    "subnet6": [
        {
            "subnet": "2003:db8:1::/64",
            "id": 5,
            <many other subnet specific details here>
        },
        {
            "subnet": "2003:db8:2::/71",
            "id": 6,
            <many other subnet specific details here>
        }
    ],
    "valid-lifetime": 120
}
]
}

```

Note that the actual response contains many additional fields that are omitted here for clarity.

24.81 network6-list

This command retrieves the full list of currently configured shared networks.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *network6-list command*

Command syntax:

```

{
    "command": "network6-list"
}

```

Response syntax:

```

{
    "arguments": {
        "shared-networks": [
            { "name": "floor1" },
            { "name": "office" }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```

},
"result": 0,
"text": "2 IPv6 network(s) found"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.82 network6-subnet-add

This command adds existing subnets to existing shared networks.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network6-subnet-add command*

Command syntax:

```

{
  "command": "network6-subnet-add",
  "arguments": {
    "name": "floor13",
    "id": 5
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "IPv6 subnet 2003:db8::/64 (id 5) is now part of shared network 'floor1'"
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.83 network6-subnet-del

This command removes a subnet that is part of an existing shared network and demotes it to a plain, stand-alone subnet.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *network6-subnet-del command*

Command syntax:

```
{
  "command": "network6-subnet-del",
  "arguments": {
    "name": "floor13",
    "id": 5
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet 2003:db8::/64 (id 5) is now removed from shared network 'floor13"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.84 perfmon-control

This command enables/disables active monitoring and statistics reporting.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 2.7.0 (*perfmon* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *perfmon-control command*

Command syntax:

```
{
  "command": "perfmon-control",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "enable-monitoring": true,
    "stats-mgr-reporting": false
  }
}

```

Response syntax:

```

{
  "arguments": {
    "enable-monitoring": true,
    "stats-mgr-reporting": false
  },
  "result": 0,
  "text": "perfmon-control success."
}

```

Result 0 is returned if command succeeds along with the resultant values of both flags. Result is 1 when parameters are malformed or missing.

24.85 perfmon-get-all-durations

This command fetches all monitored duration data currently held by Perfmon.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 2.7.0 (*perfmon* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *perfmon-get-all-durations command*

Command syntax:

```

{
  "command": "perfmon-get-all-durations",
  "arguments": {
    "result-set-format": true
  }
}

```

Response syntax:

If *result-set-format* is false, the output will be as shown below:

```

{
  "arguments": {
    "durations": [{
      "duration-key": {
        "query-type": "DHCPDISCOVER",
        "response-type": "DHCPOFFER",
        "start-event": "buffer_read",
        "stop-event": "mt_queued",
        "subnet-id": 0
      },
      "max-duration-usecs": 118,

```

(continues on next page)

(continued from previous page)

```

        "mean-duration-usecs": 47,
        "min-duration-usecs": 31,
        "occurrences": 501,
        "start-time": "2024-06-12 17:52:06.814884",
        "total-duration-usecs": 23951
    },
    ...
],
"result-set-format": false,
"interval-width-secs": 5,
"timestamp": "2024-06-12 17:52:22.397233"
},
"result": 0,
"text": "perfmon-get-all-durations: n found"
}

```

If `result-set-format` is true, the output will be as shown below:

```

{
  "arguments": {
    "durations-result-set": {
      "columns": [
        "query-type",
        "response-type",
        "start-event",
        "end-event",
        "subnet-id",
        "interval-start",
        "occurences",
        "min-duration-usecs",
        "max-duration-usecs",
        "total-duration-usecs",
        "mean-duration-usecs"
      ],
      "rows": [
        [
          "DHCPDISCOVER",
          "DHCPOFFER",
          "buffer_read",
          "mt_queued",
          0,
          "2024-06-12 17:52:06.814884",
          501,
          31,
          118,
          23951,
          47
        ],
        ...
      ]
    },
    "result-set-format": true,
    "interval-width-secs": 5,
    "timestamp": "2024-06-12 17:52:22.397233"
  }
}

```

(continues on next page)

(continued from previous page)

```

},
"result": 0,
"text": "perfmon-get-all-durations: n found"
}

```

Result 0 is returned if command succeeds along with monitored duration data. Result is 1 if command is invalid or command processing encounters an error.

24.86 remote-class4-del

This command deletes a DHCPv4 client class from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-class4-del command*

Command syntax:

```

{
  "command": "remote-class4-del",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}

```

This command includes a list with exactly one name of the client class to be deleted. The *server-tags* parameter must not be specified for this command.

Response syntax:

```

{
  "result": 0,
  "text": "1 DHCPv4 client class(es) deleted.",
  "arguments": {
    "count": 1
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported

- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.87 remote-class4-get

This command fetches a selected DHCPv4 client class by name from the specified database.

Supported by: *kea-dhcp4*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-class4-get command*

Command syntax:

```
{
  "command": "remote-class4-get",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the client class to be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 client class found.",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the client class information>
      }
    ],
    "count": 1
  }
}
```

The metadata is included in the returned shared network definition and provides the database-specific information associated with the returned object.

24.88 remote-class4-get-all

This command fetches all DHCPv4 client classes for specified servers from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-class4-get-all command*

Command syntax:

```
{
  "command": "remote-class4-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The `server-tags` list is required for this command, and must not be empty.

Response syntax:

```
{
  "result": 0,
  "text": "2 DHCPv4 client class(es) found.",
  "arguments": {
    "client-classes": [
      {
        <first client class specification>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
      },
      {
        <second client class specification>,
        "metadata": {
          "server-tags": [ <first server tag>, ... ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains a client class name and the metadata, which provides database-specific information associated with the client class.

24.89 remote-class4-set

This command creates or replaces a DHCPv4 client class in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-class4-set command*

Command syntax:

```
{
  "command": "remote-class4-set",
  "arguments": {
    "client-class": [
      {
        <client class specification>,
        "follow-class-name": <existing class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The provided list must contain exactly one client class specification. It may contain an optional parameter "follow-class-name" which can specify an existing class name to indicate that the class from the command is placed right after this existing class in the hierarchy. This parameter can be omitted or set to "null" to indicate that the new client class should be appended at the end of the hierarchy or an updated class should remain at the current position. The **server-tags** list is mandatory and must contain one or more server tags as strings to explicitly associate the client class with one or more user-defined servers. It may include the special server tag "all" to associate the class with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 shared network successfully set.",
  "arguments": {
    "client-classes": [
      {
        "name": <set client class name>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported

- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.90 remote-class6-del

This command deletes a DHCPv6 client class from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-class6-del command*

Command syntax:

```
{
  "command": "remote-class6-del",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the client class to be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 client class(es) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.91 remote-class6-get

This command fetches a selected DHCPv6 client class by name from the specified database.

Supported by: *kea-dhcp6*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-class6-get command*

Command syntax:

```
{
  "command": "remote-class6-get",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the client class to be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 client class found.",
  "arguments": {
    "client-classes": [
      {
        "name": <client class name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the client class information>
      }
    ],
    "count": 1
  }
}
```

The metadata is included in the returned shared network definition and provides the database-specific information associated with the returned object.

24.92 remote-class6-get-all

This command fetches all DHCPv6 client classes for specified servers from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-class6-get-all command*

Command syntax:

```
{
  "command": "remote-class6-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The `server-tags` list is required for this command, and must not be empty.

Response syntax:

```
{
  "result": 0,
  "text": "2 DHCPv6 client class(es) found.",
  "arguments": {
    "client-classes": [
      {
        <first client class specification>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
      },
      {
        <second client class specification>,
        "metadata": {
          "server-tags": [ <first server tag>, ... ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains a client class name and the metadata, which provides database-specific information associated with the client class.

24.93 remote-class6-set

This command creates or replaces a DHCPv6 client class in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.9.10 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-class6-set command*

Command syntax:

```
{
  "command": "remote-class6-set",
  "arguments": {
    "client-class": [
      {
        <client class specification>,
        "follow-class-name": <existing class name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The provided list must contain exactly one client class specification. It may contain an optional parameter "follow-class-name" which can specify an existing class name to indicate that the class from the command is placed right after this existing class in the hierarchy. This parameter can be omitted or set to "null" to indicate that the new client class should be appended at the end of the hierarchy or an updated class should remain at the current position. The **server-tags** list is mandatory and must contain one or more server tags as strings to explicitly associate the client class with one or more user-defined servers. It may include the special server tag "all" to associate the class with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 shared network successfully set.",
  "arguments": {
    "client-classes": [
      {
        "name": <set client class name>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported

- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.94 remote-global-parameter4-del

This command deletes a global DHCPv4 parameter from the configuration database. The server uses the value specified in the configuration file, or a default value if the parameter is not specified, after deleting the parameter from the database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter4-del command*

Command syntax:

```
{
  "command": "remote-global-parameter4-del",
  "arguments": {
    "parameters": [ <parameter name as string> ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries the list including exactly one name of the parameter to be deleted. If deleting a map parameter, the `map-name.parameter-name` format must be used. The `server-tags` list is mandatory and it must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 global parameter(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.95 remote-global-parameter4-get

This command fetches the selected global parameter for the server from the specified database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter4-get command*

Command syntax:

```
{
  "command": "remote-global-parameter4-get",
  "arguments": {
    "parameters": [ <parameter name as string> ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries a list including exactly one name of the parameter to be fetched. If retrieving a map parameter, the `map-name.parameter-name` format must be used. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed; it fetches the global parameter value shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 global parameter found.",
  "arguments": {
    "parameters": {
      <parameter name>: <parameter value>,
      "metadata": {
        "server-tags": [ <server tag> ]
      }
    },
    "count": 1
  }
}
```

The returned response contains a map with a global parameter name:value pair. The value may be a JSON string, integer, real, boolean or a map containing only one of these types. The metadata is included and provides database-specific information associated with the returned object. If the "all" server tag is specified, the command attempts to fetch the global parameter value associated with all servers. If the explicit server tag is specified, the command fetches the value associated with the given server. If the server-specific value does not exist, the `remote-global-parameter4-get` command fetches the value associated with all servers.

24.96 remote-global-parameter4-get-all

This command fetches all global parameters for the server from the specified database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter4-get-all command*

Command syntax:

```
{
  "command": "remote-global-parameter4-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed; it fetches the global parameters shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 global parameters found.",
  "arguments": {
    "parameters": [
      {
        <first parameter name>: <first parameter value>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      },
      {
        <second parameter name>: <second parameter value>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains a global parameter name:value pair. The value may be a JSON string, integer, real, boolean or a map containing only one of these types. The metadata is appended to each parameter and provides database-specific information associated with the returned objects. If the server tag "all" is included in the command, the response contains the global parameters shared among all servers. It excludes server-specific global parameters. If an explicit server tag is included in the command, the response contains all global

parameters directly associated with the given server, and the global parameters associated with all servers when server-specific values are not present.

24.97 remote-global-parameter4-set

This command creates or updates one or more global parameters in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter4-set command*

Command syntax:

```
{
  "command": "remote-global-parameter4-set",
  "arguments": {
    "parameters": {
      <first parameter name>: <first parameter value>,
      <second parameter name>: <second parameter value>
    },
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries multiple global parameters with their values (including maps with scalar parameters). Care should be taken when specifying more than one parameter; in some cases, only a subset of the parameters may be successfully stored in the database and other parameters may fail to be stored. In such cases the `remote-global-parameter4-get-all` command may be useful to verify the contents of the database after the update. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of null, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified parameters with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 global parameter(s) successfully set.",
  "arguments": {
    "parameters": {
      <first parameter name>: <first parameter value>,
      <second parameter name>: <second parameter value>
    },
    "count": 2
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.98 remote-global-parameter6-del

This command deletes a global DHCPv6 parameter from the configuration database. The server uses the value specified in the configuration file, or a default value if the parameter is not specified in the configuration file, after deleting the parameter from the database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter6-del command*

Command syntax:

```
{
  "command": "remote-global-parameter6-del",
  "arguments": {
    "parameters": [ <parameter name as string> ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries the list including exactly one name of the parameter to be deleted. If deleting a map parameter, the `map-name.parameter-name` format must be used. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 global parameter(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.99 remote-global-parameter6-get

This command fetches the selected global parameter for the server from the specified database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter6-get command*

Command syntax:

```
{
  "command": "remote-global-parameter6-get",
  "arguments": {
    "parameters": [ <parameter name as string> ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries a list including exactly one name of the parameter to be fetched. If retrieving a map parameter, the `map-name.parameter-name` format must be used. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed; it fetches the global parameter value shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 global parameter found.",
  "arguments": {
    "parameters": {
      <parameter name>: <parameter value>,
      "metadata": {
        "server-tags": [ <server tag> ]
      }
    },
    "count": 1
  }
}
```

The returned response contains a map with a global parameter name:value pair. The value may be a JSON string, integer, real, boolean or a map containing only one of these types. The metadata is included and provides database-specific information associated with the returned object. If the "all" server tag is specified, the command attempts to fetch the global parameter value associated with all servers. If the explicit server tag is specified, the command fetches the value associated with the given server. If the server-specific value does not exist, the `remote-global-parameter6-get` fetches the value associated with all servers.

24.100 remote-global-parameter6-get-all

This command fetches all global parameters for the server from the specified database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter6-get-all command*

Command syntax:

```
{
  "command": "remote-global-parameter6-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed; it fetches the global parameters shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 global parameters found.",
  "arguments": {
    "parameters": [
      {
        <first parameter name>: <first parameter value>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      },
      {
        <second parameter name>: <second parameter value>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains a global parameter name:value pair. The value may be a JSON string, integer, real, boolean or a map containing only one of these types. The metadata is appended to each parameter and provides database-specific information associated with the returned objects. If the server tag "all" is included in the command, the response contains the global parameters shared among all servers. It excludes server-specific global parameters. If an explicit server tag is included in the command, the response contains all global

parameters directly associated with the given server, and the global parameters associated with all servers when server-specific values are not present.

24.101 remote-global-parameter6-set

This command creates or updates one or more global parameters in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-global-parameter6-set command*

Command syntax:

```
{
  "command": "remote-global-parameter6-set",
  "arguments": {
    "parameters": {
      <first parameter name>: <first parameter value>,
      <second parameter name>: <second parameter value>
    },
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command carries multiple global parameters with their values (including maps with scalar parameters). Care should be taken when specifying more than one parameter; in some cases, only a subset of the parameters may be successfully stored in the database and other parameters may fail to be stored. In such cases the `remote-global-parameter6-get-all` command may be useful to verify the contents of the database after the update. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of null, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified parameters with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 global parameter(s) successfully set.",
  "arguments": {
    "parameters": {
      <first parameter name>: <first parameter value>,
      <second parameter name>: <second parameter value>
    },
    "count": 2
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.102 remote-network4-del

This command deletes an IPv4 shared network from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-network4-del command*

Command syntax:

```
{
  "command": "remote-network4-del",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "subnets-action": <'keep' | 'delete'>,
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the shared network to be deleted. The `subnets-action` parameter denotes whether the subnets in this shared network should be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 IPv4 shared network(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)

- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.103 remote-network4-get

This command fetches the selected IPv4 shared network for the server from the specified database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-network4-get command*

Command syntax:

```
{
  "command": "remote-network4-get",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "subnets-include": <'full' | 'no'>,
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the shared network to be returned. The `subnets-include` optional parameter allows for specifying whether the subnets belonging to the shared network should also be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 shared network found.",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the shared network information, potentially including
↪ subnets>
      }
    ],
    "count": 1
  }
}
```

If the subnets are returned with the shared network, they are carried in the `subnet4` list within the shared network definition. The metadata is included in the returned shared network definition and provides the database-specific information associated with the returned object.

24.104 remote-network4-list

This command fetches a list of all IPv4 shared networks from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-network4-list command*

Command syntax:

```
{
  "command": "remote-network4-list",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The `server-tags` list is required for this command, and must not be empty. It may either contain one or multiple server tags as strings, or a single null value.

Response syntax:

```
{
  "result": 0,
  "text": "2 IPv4 shared network(s) found.",
  "arguments": {
    "shared-networks": [
      {
        "name": <first shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
      },
      {
        "name": <second shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, ... ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains the list of maps. Each map contains the shared network name and the metadata, which provides database-specific information associated with the shared network. The returned list does not contain full

definitions of the shared networks; use `remote-network4-get` to fetch the full information about the selected shared networks. If the command includes explicit server tags as strings (including the special server tag "all"), the list contains all shared networks which are associated with any of the specified tags. A network is returned even if it is associated with multiple servers and only one of the specified tags matches. If the command includes the `null` value in the `server-tags` list, the response contains all shared networks which are assigned to no servers (unassigned).

24.105 remote-network4-set

This command creates or replaces an IPv4 shared network in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-network4-set command*

Command syntax:

```
{
  "command": "remote-network4-set",
  "arguments": {
    "shared-networks": [
      {
        <shared network specification excluding subnets list>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The provided list must contain exactly one shared network specification, and must not contain subnets (the "subnet4" parameter). The subnets are added to the shared network using the `remote-subnet4-set` command. The `server-tags` list is mandatory and must contain one or more server tags as strings to explicitly associate the shared network with one or more user-defined servers. It may include the special server tag "all" to associate the network with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 shared network successfully set."
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.106 remote-network6-del

This command deletes an IPv6 shared network from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-network6-del command*

Command syntax:

```
{
  "command": "remote-network6-del",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "subnets-action": <'keep' | 'delete'>,
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the shared network to be deleted. The `subnets-action` parameter indicates whether the subnets in this shared network should be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 IPv6 shared network(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.107 remote-network6-get

This command fetches the selected IPv6 shared network for the server from the specified database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-network6-get command*

Command syntax:

```
{
  "command": "remote-network6-get",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "subnets-include": <'full' | 'no'>,
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one name of the shared network to be returned. The `subnets-include` optional parameter allows for specifying whether the subnets belonging to the shared network should also be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 shared network found.",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the shared network information, potentially including
↳ subnets>
      }
    ],
    "count": 1
  }
}
```

If the subnets are returned with the shared network, they are carried in the `subnet6` list within the shared network definition. The metadata is included in the returned shared network definition and provides the database-specific information associated with the returned object.

24.108 remote-network6-list

This command fetches a list of all IPv6 shared networks from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-network6-list command*

Command syntax:

```
{
  "command": "remote-network6-list",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The `server-tags` list is required for this command, and must not be empty. It may either contain one or multiple server tags as strings, or a single null value.

Response syntax:

```
{
  "result": 0,
  "text": "2 IPv6 shared network(s) found.",
  "arguments": {
    "shared-networks": [
      {
        "name": <first shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
      },
      {
        "name": <second shared network name>,
        "metadata": {
          "server-tags": [ <first server tag>, ... ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains the list of maps. Each map contains the shared network name and the metadata, which provides database-specific information associated with the shared network. The returned list does not contain full definitions of the shared networks; use `remote-network6-get` to fetch the full information about the selected shared networks. If the command includes explicit server tags as strings (including the special server tag "all"), the list contains all shared networks which are associated with any of the specified tags. A network is returned even if it is associated

with multiple servers and only one of the specified tags matches. If the command includes the `null` value in the `server-tags` list, the response contains all shared networks which are assigned to no servers (unassigned).

24.109 remote-network6-set

This command creates or replaces an IPv6 shared network in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-network6-set command*

Command syntax:

```
{
  "command": "remote-network6-set",
  "arguments": {
    "shared-networks": [
      {
        <shared network specification excluding subnets list>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}
```

The provided list must contain exactly one shared network specification, and must not contain subnets (the "subnet6" parameter). The subnets are added to the shared network using the `remote-subnet6-set` command. The `server-tags` list is mandatory and must contain one or more server tags as strings to explicitly associate the shared network with one or more user-defined servers. It may include the special server tag "all" to associate the network with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 shared network successfully set."
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.110 remote-option-def4-del

This command deletes a DHCPv4 option definition from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def4-del command*

Command syntax:

```
{
  "command": "remote-option-def4-del",
  "arguments": {
    "option-defs": [ {
      "code": <option code>,
      "space": <option space>
    } ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command includes a list with exactly one option definition specification, comprising an option name and code. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 option definition(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.111 remote-option-def4-get

This command fetches a DHCPv4 option definition from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def4-get command*

Command syntax:

```
{
  "command": "remote-option-def4-get",
  "arguments": {
    "option-defs": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The desired option definition is identified by the pair of option code/space values. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed, to fetch the option definition instance shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option definition found.",
  "arguments": {
    "option-defs": [
      {
        <option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 1
  }
}
```

The metadata is included and provides database-specific information associated with the returned object. If the "all" server tag is specified, the command attempts to fetch the option definition associated with all servers. If the explicit server tag is specified, the command fetches the option definition associated with the given server. If the server-specific option definition does not exist, the `remote-option-def4-get` command fetches the option definition associated with all servers.

24.112 remote-option-def4-get-all

This command fetches all DHCPv4 option definitions from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def4-get-all command*

Command syntax:

```
{
  "command": "remote-option-def4-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed, to fetch the option definitions shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "2 DHCPv4 option definition(s) found.",
  "arguments": {
    "option-defs": [
      {
        <first option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      },
      {
        <second option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains an option definition specification and the metadata, including database-specific information associated with the returned objects. If the server tag "all" is included in the command, the response contains the option definitions shared among all servers. It excludes server-specific option definitions. If an explicit server tag is included in the command, the response contains all option definitions directly

associated with the given server, and the option definitions associated with all servers when server-specific option definitions are not present.

24.113 remote-option-def4-set

This command creates or replaces a DHCPv4 option definition in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def4-set command*

Command syntax:

```
{
  "command": "remote-option-def4-set",
  "arguments": {
    "option-defs": [
      {
        <option definition specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The provided list must contain exactly one option definition specification. The **server-tags** list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of **null**, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified option definition with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option definition set."
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.114 remote-option-def6-del

This command deletes a DHCPv6 option definition from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def6-del command*

Command syntax:

```
{
  "command": "remote-option-def6-del",
  "arguments": {
    "option-defs": [ {
      "code": <option code>,
      "space": <option space>
    } ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command includes a list with exactly one option definition specification, comprising an option name and code. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 option definition(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.115 remote-option-def6-get

This command fetches a DHCPv6 option definition from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def6-get command*

Command syntax:

```
{
  "command": "remote-option-def6-get",
  "arguments": {
    "option-defs": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The desired option definition is identified by the pair of option code/space values. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed, to fetch the option definition instance shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 option definition found.",
  "arguments": {
    "option-defs": [
      {
        <option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 1
  }
}
```

The metadata is included and provides database-specific information associated with the returned object. If the "all" server tag is specified, the command fetches the option definition associated with all servers. If the explicit server tag is specified, the command fetches the option definition associated with the given server. If the server-specific option definition does not exist, the `remote-option-def6-get` command fetches the option definition associated with all servers.

24.116 remote-option-def6-get-all

This command fetches all DHCPv6 option definitions from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def6-get-all command*

Command syntax:

```
{
  "command": "remote-option-def6-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed, to fetch the option definitions shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "2 DHCPv6 option definition(s) found.",
  "arguments": {
    "option-defs": [
      {
        <first option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      },
      {
        <second option definition>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains an option definition specification and the metadata, including database-specific information associated with the returned objects. If the server tag "all" is included in the command, the response contains the option definitions shared among all servers. It excludes server-specific option definitions. If an explicit server tag is included in the command, the response contains all option definitions directly

associated with the given server, and the option definitions associated with all servers when server-specific option definitions are not present.

24.117 remote-option-def6-set

This command creates or replaces a DHCPv6 option definition in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option-def6-set command*

Command syntax:

```
{
  "command": "remote-option-def6-set",
  "arguments": {
    "option-defs": [
      {
        <option definition specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The provided list must contain exactly one option definition specification. The **server-tags** list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of **null**, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified option definition with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 option definition set."
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.118 remote-option4-global-del

This command deletes a DHCPv4 global option from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-global-del command*

Command syntax:

```
{
  "command": "remote-option4-global-del",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

This command includes a list with exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or multiple server tags will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.119 remote-option4-global-get

This command fetches a global DHCPv4 option for the server from the specified database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-global-get command*

Command syntax:

```
{
  "command": "remote-option4-global-get",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The option is identified by the pair of option code/space values. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed, to fetch the global option instance shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option is found.",
  "arguments": {
    "options": [
      {
        <option information>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ]
  }
}
```

The metadata is included and provides database specific information associated with the returned object. If the "all" server tag is specified, the command fetches the global option associated with all servers. If the explicit server tag is specified, the command fetches the global option associated with the given server. If the server specific option does not exist, it fetches the option associated with all servers.

24.120 remote-option4-global-get-all

This command fetches all DHCPv4 global options for the server from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-global-get-all command*

Command syntax:

```
{
  "command": "remote-option4-global-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed, to fetch the global options shared by all servers.

Response syntax:

```
{
  "result": 0,
  "text": "2 DHCPv4 option(s) found.",
  "arguments": {
    "options": [
      {
        <first option specification>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      },
      {
        <second option specification>,
        "metadata": {
          "server-tags": [ <server tag> ]
        }
      }
    ],
    "count": 2
  }
}
```

The returned response contains a list of maps. Each map contains a global option specification and the metadata, including database-specific information associated with the returned object. If the server tag "all" is included in the command, the response contains the global options shared among all servers. It excludes server-specific global options. If an explicit server tag is included in the command, the response contains all global options directly associated with the given server, and the options associated with all servers when server-specific options are not present.

24.121 remote-option4-global-set

This command creates or replaces a DHCPv4 global option in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-global-set command*

Command syntax:

```
{
  "command": "remote-option4-global-set",
  "arguments": {
    "options": [
      {
        <global option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}
```

The provided list must contain exactly one option specification. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified option with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.122 remote-option4-network-del

This command deletes a DHCPv4 option from a shared network from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-network-del command*

Command syntax:

```
{
  "command": "remote-option4-network-del",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly one name of the shared network and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.123 remote-option4-network-set

This command creates or replaces a DHCPv4 option in a shared network in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-network-set command*

Command syntax:

```
{
  "command": "remote-option4-network-set",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "options": [
      {
        <shared network option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

The provided lists must contain exactly one name of the shared network and one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported

- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.124 remote-option4-pool-del

This command deletes a DHCPv4 option from an address pool from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-pool-del command*

Command syntax:

```
{
  "command": "remote-option4-pool-del",
  "arguments": {
    "pools": [
      {
        "pool": <pool range or prefix>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly one address pool specification and exactly one option specification comprising an option space name and code. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.125 remote-option4-pool-set

This command creates or replaces a DHCPv4 option in an address pool in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-pool-set command*

Command syntax:

```
{
  "command": "remote-option4-pool-set",
  "arguments": {
    "pools": [
      {
        "pool": <pool range or prefix>
      }
    ],
    "options": [
      {
        <address pool option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly address pool specification and exactly one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.126 remote-option4-subnet-del

This command deletes a DHCPv4 option from a subnet from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-subnet-del command*

Command syntax:

```
{
  "command": "remote-option4-subnet-del",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly one ID of the subnet and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.127 remote-option4-subnet-set

This command creates or replaces a DHCPv4 option in a subnet in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option4-subnet-set command*

Command syntax:

```
{
  "command": "remote-option4-subnet-set",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "options": [
      {
        <subnet option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

The provided lists must contain exactly one ID of the subnet and one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,

```

(continues on next page)

(continued from previous page)

```

        "space": <option space>
    }
]
}
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.128 remote-option6-global-del

This command deletes a DHCPv6 global option from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-global-del command*

Command syntax:

```

{
    "command": "remote-option6-global-del",
    "arguments": {
        "options": [
            {
                "code": <option code>,
                "space": <option space>
            }
        ],
        "remote": {
            <specification of the database to connect to>
        },
        "server-tags": [ <single server tag as string> ]
    }
}

```

This command includes a list with exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or multiple server tags will result in an error.

Response syntax:

```

{
    "result": 0,
    "text": "1 DHCPv6 option(s) deleted.",

```

(continues on next page)

(continued from previous page)

```

    "arguments": {
        "count": 1
    }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.129 remote-option6-global-get

This command fetches a global DHCPv6 option for the server from the specified database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-global-get command*

Command syntax:

```

{
    "command": "remote-option6-global-get",
    "arguments": {
        "options": [
            {
                "code": <option code>,
                "space": <option space>
            }
        ],
        "remote": {
            <specification of the database to connect to>
        },
        "server-tags": [ <single server tag as string> ]
    }
}

```

The option is identified by the pair of option code/space values. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed, to fetch the global option instance shared by all servers.

Response syntax:

```

{
    "result": 0,
    "text": "DHCPv6 option is found.",

```

(continues on next page)

(continued from previous page)

```

    "arguments": {
      "options": [
        {
          <option information>,
          "metadata": {
            "server-tags": [ <server tag> ]
          }
        }
      ]
    }
  }
}

```

The metadata is included and provides database-specific information associated with the returned object. If the "all" server tag is specified, the command attempts to fetch the global option associated with all servers. If the explicit server tag is specified, the command will fetch the global option associated with the given server. If the server-specific option does not exist, it fetches the option associated with all servers.

24.130 remote-option6-global-get-all

This command fetches all DHCPv6 global options for the server from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-global-get-all command*

Command syntax:

```

{
  "command": "remote-option6-global-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <single server tag as string> ]
  }
}

```

The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The special server tag "all" is allowed, to fetch the global options shared by all servers.

Response syntax:

```

{
  "result": 0,
  "text": "2 DHCPv6 option(s) found.",
  "arguments": {
    "options": [
      {
        <first option specification>,
        "metadata": {

```

(continues on next page)

(continued from previous page)

```

        "server-tags": [ <server tag> ]
    },
    {
        <second option specification>,
        "metadata": {
            "server-tags": [ <server tag> ]
        }
    }
],
"count": 2
}

```

The returned response contains a list of maps. Each map contains a global option specification and the metadata, including database-specific information associated with the returned object. If the server tag "all" is included in the command, the response contains the global options shared between all servers. It excludes server-specific global options. If an explicit server tag is included in the command, the response contains all global options directly associated with the given server, and the options associated with all servers when server-specific options are not present.

24.131 remote-option6-global-set

This command creates or replaces a DHCPv6 global option in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-global-set command*

Command syntax:

```

{
    "command": "remote-option6-global-set",
    "arguments": {
        "options": [
            {
                <global option specification>
            }
        ],
        "remote": {
            <specification of the database to connect to>
        },
        "server-tags": [ <single server tag as string> ]
    }
}

```

The provided list must contain exactly one option specification. The `server-tags` list is mandatory and must contain exactly one server tag. Specifying an empty list, a value of `null`, or multiple server tags will result in an error. The server tag "all" is allowed; it associates the specified option with all servers.

Response syntax:


```
{
  "result": 0,
  "text": "DHCPv6 option set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.132 remote-option6-network-del

This command deletes a DHCPv6 option from a shared network from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-network-del command*

Command syntax:

```
{
  "command": "remote-option6-network-del",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

This command includes two lists with exactly one name of the shared network and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of `null`, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.133 remote-option6-network-set

This command creates or replaces a DHCPv6 option in a shared network in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-network-set command*

Command syntax:

```
{
  "command": "remote-option6-network-set",
  "arguments": {
    "shared-networks": [
      {
        "name": <shared network name>
      }
    ],
    "options": [
      {
        <shared network option specification>
      }
    ],
    "remote": {
```

(continues on next page)

(continued from previous page)

```

    <specification of the database to connect to>
  }
}

```

The provided lists must contain exactly one name of the shared network and one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```

{
  "result": 0,
  "text": "DHCPv6 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.134 remote-option6-pd-pool-del

This command deletes a DHCPv6 option from a prefix delegation pool from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-pd-pool-del command*

Command syntax:

```

{
  "command": "remote-option6-pd-pool-del",
  "arguments": {
    "pd-pools": [
      {
        "prefix": <pool prefix (address part)>,
        "prefix-len": <pool prefix (length part)>
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ],
  "options": [
    {
      "code": <option code>,
      "space": <option space>
    }
  ],
  "remote": {
    <specification of the database to connect to>
  }
}

```

This command includes two lists with exactly one prefix delegation pool specification and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of `null`, or a server tag will result in an error.

Response syntax:

```

{
  "result": 0,
  "text": "1 DHCPv6 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.135 remote-option6-pd-pool-set

This command creates or replaces a DHCPv6 option in a prefix delegation pool in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-pd-pool-set command*

Command syntax:

```

{
  "command": "remote-option6-pd-pool-set",

```

(continues on next page)

(continued from previous page)

```

"arguments": {
  "pd-pools": [
    {
      "prefix": <pool prefix (address part)>,
      "prefix-len": <pool prefix (length part)>
    }
  ],
  "options": [
    {
      <prefix delegation pool option specification>
    }
  ],
  "remote": {
    <specification of the database to connect to>
  }
}

```

This command includes two lists with exactly one prefix delegation pool specification and exactly one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```

{
  "result": 0,
  "text": "DHCPv6 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.136 remote-option6-pool-del

This command deletes a DHCPv6 option from an address pool from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see [remote-option6-pool-del command](#)

Command syntax:

```
{
  "command": "remote-option6-pool-del",
  "arguments": {
    "pools": [
      {
        "pool": <pool range or prefix>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly one address pool specification and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.137 remote-option6-pool-set

This command creates or replaces a DHCPv6 option in an address pool in the configuration database.

Supported by: [kea-dhcp6](#)

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-pool-set command*

Command syntax:

```
{
  "command": "remote-option6-pool-set",
  "arguments": {
    "pools": [
      {
        "pool": <pool range or prefix>
      }
    ],
    "options": [
      {
        <address pool option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly address pool specification and exactly one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.138 remote-option6-subnet-del

This command deletes a DHCPv6 option from a subnet from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-subnet-del command*

Command syntax:

```
{
  "command": "remote-option6-subnet-del",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes two lists with exactly one ID of the subnet and exactly one option specification, comprising an option name and code. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 option(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.139 remote-option6-subnet-set

This command creates or replaces a DHCPv6 option in a subnet in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-option6-subnet-set command*

Command syntax:

```
{
  "command": "remote-option6-subnet-set",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "options": [
      {
        <subnet option specification>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

The provided lists must contain exactly one ID of the subnet and one option specification. Specifying an empty list, a value of null, or a server tag will result in an error.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 option successfully set.",
  "arguments": {
    "options": [
      {
        "code": <option code>,
        "space": <option space>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported

- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.140 remote-server4-del

This command deletes information about a DHCPv4 server from the configuration database. Any configuration explicitly associated with the deleted server is automatically disassociated. In addition, configuration elements not shareable with other servers (e.g. global DHCP parameters) are deleted. Shareable configuration elements (e.g. subnets, shared networks) are not deleted as they may be used by other servers.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-server4-del* command

Command syntax:

```
{
  "command": "remote-server4-del",
  "arguments": {
    "servers": [
      {
        "server-tag": <server name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command carries the list including exactly one map with the tag of the server to be deleted.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv4 server(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.141 remote-server4-get

This command fetches information about the DHCPv4 server, such as the server tag and description.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-server4-get command*

Command syntax:

```
{
  "command": "remote-server4-get",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command carries the list including exactly one map with the tag of the server to be fetched.

Response syntax:

```
{
  "result": 0,
  "text": "DHCP server 'server tag' found.",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ],
    "count": 1
  }
}
```

The server tag is the unique identifier of the server, used to associate the configuration elements in the database with the particular server instance. The returned server description is specified by the user when setting the server information.

24.142 remote-server4-get-all

This command fetches information about all DHCPv4 servers specified by the user.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-server4-get-all command*

Command syntax:

```
{
  "command": "remote-server4-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command contains no arguments besides the optional `remote`.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv4 servers found.",
  "arguments": {
    "servers": [
      {
        "server-tag": <first server tag>,
        "description": <first server description>
      },
      {
        "server-tag": <second server tag>,
        "description": <second server description>
      }
    ],
    "count": 2
  }
}
```

The returned response contain a list of maps. Each map contains a server tag uniquely identifying a server, and the user-defined description of the server. The Kea Configuration Backend uses the keyword `all` to associate parts of the configuration with all servers. Internally, it creates the logical server `all` for this purpose. However, this logical server is not returned as a result of the `remote-server4-get-all` command; only the user-defined servers are returned.

24.143 remote-server4-set

This command creates or replaces information about the DHCPv4 server in the database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-server4-set command*

Command syntax:

```
{
  "command": "remote-server4-set",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}

```

The provided list must contain exactly one server specification. The `server-tag` must be unique across all servers within the configuration database. The `description` is the arbitrary text describing the server, its location within the network, etc.

Response syntax:

```

{
  "result": 0,
  "text": "DHCPv4 server successfully set.",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ]
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.144 remote-server6-del

This command deletes information about a DHCPv6 server from the configuration database. Any configuration explicitly associated with the deleted server is automatically disassociated. In addition, configuration elements not shareable with other servers (e.g. global DHCP parameters) are deleted. Shareable configuration elements (e.g. subnets, shared networks) are not deleted as they may be used by other servers.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-server6-del command*

Command syntax:

```
{
  "command": "remote-server6-del",
  "arguments": {
    "servers": [
      {
        "server-tag": <server name>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command carries the list including exactly one map with the tag of the server to be deleted.

Response syntax:

```
{
  "result": 0,
  "text": "1 DHCPv6 server(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.145 remote-server6-get

This command fetches information about the DHCPv6 server, such as the server tag and description.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-server6-get command*

Command syntax:

```
{
  "command": "remote-server6-get",
```

(continues on next page)

(continued from previous page)

```

"arguments": {
  "servers": [
    {
      "server-tag": <server tag>
    }
  ],
  "remote": {
    <specification of the database to connect to>
  }
}

```

This command carries the list including exactly one map with the tag of the server to be fetched.

Response syntax:

```

{
  "result": 0,
  "text": "DHCP server 'server tag' found.",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ],
    "count": 1
  }
}

```

The server tag is the unique identifier of the server, used to associate the configuration elements in the database with the particular server instance. The returned server description is specified by the user when setting the server information.

24.146 remote-server6-get-all

This command fetches information about all DHCPv6 servers specified by the user.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-server6-get-all command*

Command syntax:

```

{
  "command": "remote-server6-get-all",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    }
  }
}

```

This command contains no arguments besides the optional `remote`.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 servers found.",
  "arguments": {
    "servers": [
      {
        "server-tag": <first server tag>,
        "description": <first server description>
      },
      {
        "server-tag": <second server tag>,
        "description": <second server description>
      }
    ],
    "count": 2
  }
}
```

The returned response contain a list of maps. Each map contains a server tag uniquely identifying a server, and the user-defined description of the server. The Kea Configuration Backend uses the keyword `all` to associate parts of the configuration with all servers. Internally, it creates the logical server `all` for this purpose. However, this logical server is not returned as a result of the `remote-server6-get-all` command; only the user-defined servers are returned.

24.147 remote-server6-set

This command creates or replaces information about the DHCPv6 server in the database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-server6-set command*

Command syntax:

```
{
  "command": "remote-server6-set",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```


The provided list must contain exactly one server specification. The `server-tag` must be unique across all servers within the configuration database. The `description` is the arbitrary text describing the server, its location within the network, etc.

Response syntax:

```
{
  "result": 0,
  "text": "DHCPv6 server successfully set.",
  "arguments": {
    "servers": [
      {
        "server-tag": <server tag>,
        "description": <server description>
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.148 remote-subnet4-del-by-id

This command deletes an IPv4 subnet by ID from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-del-by-id command*

Command syntax:

```
{
  "command": "remote-subnet4-del-by-id",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one ID of the subnet to be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 IPv4 subnet(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.149 remote-subnet4-del-by-prefix

This command deletes an IPv4 subnet by prefix from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-del-by-prefix command*

Command syntax:

```
{
  "command": "remote-subnet4-del-by-prefix",
  "arguments": {
    "subnets": [
      {
        "subnet": <subnet prefix>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one prefix of the subnet to be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 IPv4 subnet(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.150 remote-subnet4-get-by-id

This command fetches the selected IPv4 subnet by ID from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-get-by-id command*

Command syntax:

```
{
  "command": "remote-subnet4-get-by-id",
  "arguments": {
    "subnets": [ {
      "id": <subnet identifier>
    } ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one ID of the subnet to be returned. The *server-tags* parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 subnet found.",
  "arguments": {
    "subnets": [ {
      "id": <subnet identifier>,

```

(continues on next page)

(continued from previous page)

```

    "subnet": <subnet prefix>,
    "shared-network-name": <shared network name or null>,
    "metadata": {
        "server-tags": [ <first server tag>, <second server tag>, ... ]
    },
    <the rest of the subnet specification here>
} ],
"count": 1
}
}

```

If the shared network name is null, it means that the returned subnet does not belong to any shared network (a global subnet). The metadata is included in the returned subnet definition and provides database-specific information associated with the returned object.

24.151 remote-subnet4-get-by-prefix

This command fetches the selected IPv4 subnet by prefix from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-get-by-prefix command*

Command syntax:

```

{
  "command": "remote-subnet4-get-by-prefix",
  "arguments": {
    "subnets": [ {
      "subnet": <subnet prefix>
    } ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}

```

This command includes a list with exactly one prefix of the subnet to be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```

{
  "result": 0,
  "text": "IPv4 subnet found.",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>,
        "subnet": <subnet prefix>,
        "shared-network-name": <shared network name or null>,

```

(continues on next page)

(continued from previous page)

```

        "metadata": {
            "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the subnet specification here>
    }
],
"count": 1
}

```

If the shared network name is null, it means that the returned subnet does not belong to any shared network (global subnet). The metadata is included in the returned subnet definition and provides database-specific information associated with the returned object.

24.152 remote-subnet4-list

This command fetches a list of all IPv4 subnets from the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-list command*

Command syntax:

```

{
  "command": "remote-subnet4-list",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}

```

The `server-tags` list is required for this command, and must not be empty. It may either contain one or multiple server tags as strings, or a single null value.

Response syntax:

```

{
  "result": 0,
  "text": "2 IPv4 subnets found.",
  "arguments": {
    "subnets": [
      {
        "id": <first subnet identifier>,
        "subnet": <first subnet prefix>,
        "shared-network-name": <shared network name or null>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    },
    {
      "id": <second subnet identifier>,
      "subnet": <second subnet prefix>,
      "shared-network-name": <shared network name or null>,
      "metadata": {
        "server-tags": [ <first server tag>, ... ]
      }
    }
  ],
  "count": 2
}

```

The returned response contains a list of maps. Each map contains a subnet identifier, prefix, and shared network name to which the subnet belongs. If the subnet does not belong to a shared network, the name is null. The metadata includes database-specific information associated with the subnets. The returned list does not contain full subnet definitions; use `remote-subnet4-get` to fetch the full information about the selected subnets. If the command includes explicit server tags as strings (including the special server tag "all"), the list contains all subnets which are associated with any of the specified tags. A subnet is returned even if it is associated with multiple servers and only one of the specified tags matches. If the command includes the null value in the `server-tags` list, the response contains all subnets which are assigned to no servers (unassigned).

24.153 remote-subnet4-set

This command creates or replaces an IPv4 subnet in the configuration database.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet4-set command*

Command syntax:

```

{
  "command": "remote-subnet4-set",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>,
        "subnet": <subnet prefix>,
        "shared-network-name": <shared network name or null>,
        <the rest of the subnet specification here>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}

```

The provided list must contain exactly one subnet specification. The `shared-network-name` parameter is required for these commands; it associates the subnet with the shared network by its name. If the subnet must not belong to any shared network (a global subnet), the null value must be specified for the shared network name. The `server-tags` list is mandatory and must contain one or more server tags as strings to explicitly associate the subnet with one or more user-defined servers. The `remote-subnet4-set` command may include the special server tag "all" to associate the subnet with all servers.

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 subnet successfully set.",
  "arguments": {
    "id": <subnet identifier>,
    "subnet": <subnet prefix>
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.154 remote-subnet6-del-by-id

This command deletes an IPv6 subnet by ID from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-del-by-id command*

Command syntax:

```
{
  "command": "remote-subnet6-del-by-id",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one ID of the subnet to be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "1 IPv6 subnet(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.155 remote-subnet6-del-by-prefix

This command deletes an IPv6 subnet by prefix from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-del-by-prefix command*

Command syntax:

```
{
  "command": "remote-subnet6-del-by-prefix",
  "arguments": {
    "subnets": [
      {
        "subnet": <subnet prefix>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one prefix of the subnet to be deleted. The `server-tags` parameter must not be specified for this command.

Response syntax:


```
{
  "result": 0,
  "text": "1 IPv6 subnet(s) deleted.",
  "arguments": {
    "count": 1
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.156 remote-subnet6-get-by-id

This command fetches the selected IPv6 subnet by ID from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-get-by-id command*

Command syntax:

```
{
  "command": "remote-subnet6-get-by-id",
  "arguments": {
    "subnets": [
      {
        "id": <subnet identifier>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}
```

This command includes a list with exactly one ID of the subnet to be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet found.",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "subnets": [
      {
        "id": <subnet identifier>,
        "subnet": <subnet prefix>,
        "shared-network-name": <shared network name or null>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the subnet specification here>
      }
    ],
    "count": 1
  }
}

```

If the shared network name is null, it means that the returned subnet does not belong to any shared network (a global subnet). The metadata is included in the returned subnet definition and provides database-specific information associated with the returned object.

24.157 remote-subnet6-get-by-prefix

This command fetches the selected IPv6 subnet by prefix from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-get-by-prefix command*

Command syntax:

```

{
  "command": "remote-subnet6-get-by-prefix",
  "arguments": {
    "subnets": [
      {
        "subnet": <subnet prefix>
      }
    ],
    "remote": {
      <specification of the database to connect to>
    }
  }
}

```

This command includes a list with exactly one prefix of the subnet to be returned. The `server-tags` parameter must not be specified for this command.

Response syntax:

```

{
  "result": 0,
  "text": "IPv6 subnet found.",

```

(continues on next page)

(continued from previous page)

```

    "arguments": {
      "subnets": [ {
        "id": <subnet identifier>,
        "subnet": <subnet prefix>,
        "shared-network-name": <shared network name or null>,
        "metadata": {
          "server-tags": [ <first server tag>, <second server tag>, ... ]
        },
        <the rest of the subnet specification here>
      } ],
      "count": 1
    }
  }
}

```

If the shared network name is null, it means that the returned subnet does not belong to any shared network (global subnet). The metadata is included in the returned subnet definition and provides database-specific information associated with the returned object.

24.158 remote-subnet6-list

This command fetches a list of all IPv6 subnets from the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-list command*

Command syntax:

```

{
  "command": "remote-subnet6-list",
  "arguments": {
    "remote": {
      <specification of the database to connect to>
    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}

```

The **server-tags** list is required for this command, and must not be empty. It may either contain one or multiple server tags as strings, or a single null value.

Response syntax:

```

{
  "result": 0,
  "text": "2 IPv6 subnets found.",
  "arguments": {
    "subnets": [
      {
        "id": <first subnet identifier>,
        "subnet": <first subnet prefix>,

```

(continues on next page)

(continued from previous page)

```

        "shared-network-name": <shared network name or null>,
        "metadata": {
            "server-tags": [ <first server tag>, <second server tag>, ... ]
        }
    },
    {
        "id": <second subnet identifier>,
        "subnet": <second subnet prefix>,
        "shared-network-name": <shared network name or null>,
        "metadata": {
            "server-tags": [ <first server tag>, ... ]
        }
    }
],
"count": 2
}

```

The returned response contains a list of maps. Each map contains a subnet identifier, prefix, and shared network name to which the subnet belongs. If the subnet does not belong to a shared network, the name is null. The metadata includes database-specific information associated with the subnets. The returned list does not contain full subnet definitions; use `remote-subnet6-get` to fetch the full information about the selected subnets. If the command includes explicit server tags as strings (including the special server tag "all"), the list contains all subnets which are associated with any of the specified tags. A subnet is returned even if it is associated with multiple servers and only one of the specified tags matches. If the command includes the null value in the `server-tags` list, the response contains all subnets which are assigned to no servers (unassigned).

24.159 remote-subnet6-set

This command creates or replaces an IPv6 subnet in the configuration database.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*cb_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *remote-subnet6-set command*

Command syntax:

```

{
    "command": "remote-subnet6-set",
    "arguments": {
        "subnets": [
            {
                "id": <subnet identifier>,
                "subnet": <subnet prefix>,
                "shared-network-name": <shared network name or null>,
                <the rest of the subnet specification here>
            }
        ],
        "remote": {
            <specification of the database to connect to>
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

    },
    "server-tags": [ <first server tag>, <second server tag>, ... ]
  }
}

```

The provided list must contain exactly one subnet specification. The `shared-network-name` parameter is required for these commands; it associates the subnet with the shared network by its name. If the subnet must not belong to any shared network (a global subnet), the `null` value must be specified for the shared network name. The `server-tags` list is mandatory and must contain one or more server tags as strings to explicitly associate the subnet with one or more user-defined servers. The `remote-subnet6-set` command may include the special server tag "all" to associate the subnet with all servers.

Response syntax:

```

{
  "result": 0,
  "text": "IPv6 subnet successfully set.",
  "arguments": {
    "id": <subnet identifier>,
    "subnet": <subnet prefix>
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.160 reservation-add

This command adds a new host reservation. The reservation may include IPv4 addresses, IPv6 addresses, IPv6 prefixes, various identifiers, a class the client will be assigned to, DHCPv4 and DHCPv6 options, and more.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (*host_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *reservation-add command*

Command syntax:

```

{
  "command": "reservation-add",
  "arguments": {
    "reservation": {
      "boot-file-name": <string>,
      "client-id": <string>,

```

(continues on next page)

(continued from previous page)

```

    "circuit-id": <string>,
    "duid": <string>,
    "flex-id": <string>,
    "ip-address": <string (IPv4 address)>,
    "ip-addresses": [ <comma-separated strings> ],
    "hw-address": <string>,
    "hostname": <string>,
    "next-server": <string (IPv4 address)>,
    "option-data": [ <comma-separated structures defining options> ],
    "prefixes": [ <comma-separated IPv6 prefixes> ],
    "client-classes": [ <comma-separated strings> ],
    "server-hostname": <string>,
    "subnet-id": <integer>,
    "user-context": <any valid JSON>
  },
  "operation-target": <string (memory, database, all, default)>
}

```

Note that boot-file-name, circuit-id, client-id, ip-address, next-server, and server-hostname are IPv4-specific. ip-addresses, and prefixes are IPv6-specific. Operation-target is optional; default is 'alternate'.

Response syntax:

```

{
  "result": <integer>,
  "text": <string>
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.161 reservation-del

This command deletes an existing host reservation.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (*host_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *reservation-del command*

Command syntax:

```
{
  "command": "reservation-del",
  "arguments": {
    "subnet-id": <integer>,
    "ip-address": <string>,
    "identifier-type": <one of 'hw-address', 'duid', 'circuit-id', 'client-id' and
↪ 'flex-id'>,
    "identifier": <string>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

The host reservation can be identified by either the (subnet-id, ip-address) pair or a triplet of (subnet-id, identifier-type, identifier). Operation-target is optional; default is 'alternate'.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.162 reservation-get

This command retrieves an existing host reservation.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get command*

Command syntax:

```
{
  "command": "reservation-get",
  "arguments": {
    "subnet-id": <integer>,
    "identifier-type": <one of 'hw-address', 'duid', 'circuit-id', 'client-id' and
↪ 'flex-id'>,
    "identifier": <string>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

The host reservation can be identified by either the (subnet-id, ip-address) pair or a triplet of (subnet-id, identifier-type, identifier). Operation-target is optional; default is 'all'.

Response syntax:

```
{
  "result": <integer>,
  "text": <string>,
  "arguments": {
    "boot-file-name": <string>,
    "comment": <string>,
    "client-id": <string>,
    "circuit-id": <string>,
    "duid": <string>,
    "flex-id": <string>,
    "ip-address": <string (IPv4 address)>,
    "ip-addresses": [ <comma-separated strings> ],
    "hw-address": <string>,
    "hostname": <string>,
    "next-server": <string (IPv4 address)>,
    "option-data": [ <comma-separated structures defining options> ],
    "prefixes": [ <comma-separated IPv6 prefixes> ],
    "client-classes": [ <comma-separated strings> ],
    "server-hostname": <string>,
    "subnet-id": <integer>,
    "user-context": <any valid JSON>
  }
}
```

The arguments object appears only if a host is found. Many fields in the arguments object appear only if a specific field is set.

24.163 reservation-get-all

This command retrieves all host reservations for a specified subnet.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get-all command*

Command syntax:

```
{
  "command": "reservation-get-all",
  "arguments": {
    "subnet-id": <integer>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

Operation-target is optional; default is 'all'.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

The reservation-get-all command may result in very large responses.

24.164 reservation-get-by-address

This command retrieves all host reservations for given ip-address and optionally a specified subnet.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 2.4.0 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get-by-address command*

Command syntax:

```
{
  "command": "reservation-get-by-address",
  "arguments": {
    "ip-address": <string>,
    "subnet-id": <integer>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

The host reservations can be identified by a pair of 'ip-address' and 'subnet-id'. 'subnet-id' is optional. 'operation-target' is optional; default is 'all'.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

The reservation-get-by-address may return many reservations for the same ip-address. This command may be useful in case ip-reservations-unique configuration flag is set to false.

24.165 reservation-get-by-hostname

This command retrieves all host reservations for a specified hostname and optionally a specified subnet.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.1 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get-by-hostname command*

Command syntax:

```
{
  "command": "reservation-get-by-hostname",
  "arguments": {
    "hostname": <hostname>,
    "subnet-id": <integer>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

Operation-target is optional; default is 'all'.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

The reservation-get-by-hostname command may result in large responses.

24.166 reservation-get-by-id

This command retrieves all host reservations for a specified identifier (type and value).

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.9.0 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get-by-id command*

Command syntax:

```
{
  "command": "reservation-get-by-id",
  "arguments": {
    "identifier-type": <one of 'hw-address', 'duid', 'circuit-id', 'client-id' and
    ↪ 'flex-id'>,
    "identifier": <string>,
    "operation-target": <string (memory, database, all, default)>
  }
}
```

Operation-target is optional; default is 'all'.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

The reservation-get-by-id command may result in large responses.

24.167 reservation-get-page

This command retrieves all host reservations or host reservations for a specified subnet by page.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (*host_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *reservation-get-page command*

Command syntax:

```
{
  "command": "reservation-get-page",
  "arguments": {
    "subnet-id": <integer>,
    "limit": <integer>,
    "source-index": <integer>,
    "from": <integer>
  }
}
```

The page size limit is mandatory. The subnet-id is optional since version 1.9.0. The source-index and from host-id are optional and default to 0. Values to use to load the next page are returned in responses in a next map.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.168 reservation-update

This command updates an existing host reservation. The reservation has to include host identifiers and a subnet identifier and may include IPv4 addresses, IPv6 addresses, IPv6 prefixes, various identifiers, a class the client will be assigned to, DHCPv4 and DHCPv6 options, and more.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 2.3.7 (*host_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *reservation-update command*

Command syntax:

```
{
  "command": "reservation-update",
  "arguments": {
    "reservation": {
      "boot-file-name": <string>,
      "client-id": <string>,
      "circuit-id": <string>,
      "duid": <string>,
      "flex-id": <string>,
      "ip-address": <string (IPv4 address)>,
      "ip-addresses": [ <comma-separated strings> ],
      "hw-address": <string>,
      "hostname": <string>,
      "next-server": <string (IPv4 address)>,
      "option-data": [ <comma-separated structures defining options> ],
      "prefixes": [ <comma-separated IPv6 prefixes> ],
      "client-classes": [ <comma-separated strings> ],
      "server-hostname": <string>,
      "subnet-id": <integer>,
      "user-context": <any valid JSON>
    },
    "operation-target": <string (memory, database, all, default)>
  }
}
```

Note that boot-file-name, circuit-id, client-id, ip-address, next-server, and server-hostname are IPv4-specific. ip-addresses, and prefixes are IPv6-specific. Operation-target is optional; default is 'alternate'.

Response syntax:

```
{
  "result": <integer>,
  "text": <string>
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.169 server-tag-get

This command returns the server tag used by the server. Server tag is essential configuration parameter in the Config Backend configuration. This parameter is configured in the local config file. This command does not take any parameters.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *server-tag-get command*

Command syntax:

```
{
  "command": "server-tag-get"
}
```

Response syntax:

```
{
  "result": 0,
  "arguments": {
    "server-tag": "office1"
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.170 shutdown

This command instructs the server to initiate its shutdown procedure.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *shutdown command*

Command syntax:

```
{
  "command": "shutdown",
  "arguments": {
    "exit-value": 123
  }
}
```

The server responds with a confirmation that the shutdown procedure has been initiated.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.171 stat-lease4-get

This command fetches lease statistics for a range of known IPv4 subnets.

Supported by: *kea-dhcp4*

Availability: 1.4.0 (*stat_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *stat-lease4-get command*

Command syntax:

```
{
  "command": "stat-lease4-get"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "stat-lease4-get: 2 rows found",
  "arguments": {
    "result-set": {
      "columns": [ "subnet-id", "total-addresses", "cumulative-assigned-addresses",
↪ "assigned-addresses", "declined-addresses" ],
      "rows": [
        [ 10, 256, 200, 111, 0 ],
        [ 20, 4098, 5000, 2034, 4 ]
      ],
      "timestamp": "2018-05-04 15:03:37.000000"
    }
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error

- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.172 stat-lease6-get

This command fetches lease statistics for a range of known IPv6 subnets.

Supported by: *kea-dhcp6*

Availability: 1.4.0 (*stat_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *stat-lease6-get command*

Command syntax:

```
{
  "command": "stat-lease6-get",
  "arguments": {
    "subnet-id" : 10
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "stat-lease6-get: 2 rows found",
  "arguments": {
    "result-set": {
      "columns": [ "subnet-id", "total-nas", "cumulative-assigned-nas", "assigned-nas",
↪ "declined-addresses", "total-pds", "cumulative-assigned-pds", "assigned-pds" ],
      "rows": [
        [ 10, 4096, 3000, 2400, 3, 0, 0 ],
        [ 20, 0, 0, 0, 1048, 500, 233 ],
        [ 30, 256, 300, 60, 0, 1048, 15, 15 ]
      ],
      "timestamp": "2018-05-04 15:03:37.000000"
    }
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.173 statistic-get

This command retrieves a single statistic. It takes a single string parameter called *name* that specifies the statistic name.

Supported by: *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *statistic-get command*

Command syntax:

```
{
  "command": "statistic-get",
  "arguments": {
    "name": "pkt4-received"
  }
}
```

The server responds with the details of the requested statistic, with a result of 0 indicating success, and the specified statistic as the value of the "arguments" parameter.

Response syntax:

```
{
  "result": 0,
  "arguments": {
    "pkt4-received": [ [ "first_value", "2019-07-30 10:11:19.498739" ], [ "second_
↪value", "2019-07-30 10:11:19.498662" ] ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.174 statistic-get-all

This command retrieves all recorded statistics.

Supported by: *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *statistic-get-all command*

Command syntax:


```
{
  "command": "statistic-get-all",
  "arguments": { }
}
```

The server responds with the details of all recorded statistics, with a result of 0 indicating that it iterated over all statistics (even when the total number of statistics is zero).

Response syntax:

```
{
  "result": 0,
  "arguments": {
    "cumulative-assigned-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836166"
      ]
    ],
    "declined-addresses": [
      [
        0,
        "2023-06-13 20:42:46.836156"
      ]
    ],
    "pkt4-ack-received": [
      [
        0,
        "2023-06-13 20:42:46.616247"
      ]
    ],
    "pkt4-ack-sent": [
      [
        0,
        "2023-06-13 20:42:46.616290"
      ]
    ],
    "pkt4-decline-received": [
      [
        0,
        "2023-06-13 20:42:46.616296"
      ]
    ],
    "pkt4-discover-received": [
      [
        0,
        "2023-06-13 20:42:46.616303"
      ]
    ],
    "pkt4-inform-received": [
      [
        0,
        "2023-06-13 20:42:46.616308"
      ]
    ]
  ]
}
```

(continues on next page)

(continued from previous page)

```
    ],
    "pkt4-nak-received": [
        [
            0,
            "2023-06-13 20:42:46.616312"
        ]
    ],
    "pkt4-nak-sent": [
        [
            0,
            "2023-06-13 20:42:46.616314"
        ]
    ],
    "pkt4-offer-received": [
        [
            0,
            "2023-06-13 20:42:46.616318"
        ]
    ],
    "pkt4-offer-sent": [
        [
            0,
            "2023-06-13 20:42:46.616323"
        ]
    ],
    "pkt4-parse-failed": [
        [
            0,
            "2023-06-13 20:42:46.616326"
        ]
    ],
    "pkt4-receive-drop": [
        [
            0,
            "2023-06-13 20:42:46.616330"
        ]
    ],
    "pkt4-received": [
        [
            0,
            "2023-06-13 20:42:46.616335"
        ]
    ],
    "pkt4-release-received": [
        [
            0,
            "2023-06-13 20:42:46.616339"
        ]
    ],
    "pkt4-request-received": [
        [
```

(continues on next page)

(continued from previous page)

```
        0,
        "2023-06-13 20:42:46.616343"
    ],
    "pkt4-sent": [
        [
            0,
            "2023-06-13 20:42:46.616348"
        ]
    ],
    "pkt4-unknown-received": [
        [
            0,
            "2023-06-13 20:42:46.616354"
        ]
    ],
    "reclaimed-declined-addresses": [
        [
            0,
            "2023-06-13 20:42:46.836159"
        ]
    ],
    "reclaimed-leases": [
        [
            0,
            "2023-06-13 20:42:46.836163"
        ]
    ],
    "subnet[1].assigned-addresses": [
        [
            0,
            "2023-06-13 20:42:46.836173"
        ]
    ],
    "subnet[1].cumulative-assigned-addresses": [
        [
            0,
            "2023-06-13 20:42:46.836098"
        ]
    ],
    "subnet[1].declined-addresses": [
        [
            0,
            "2023-06-13 20:42:46.836178"
        ]
    ],
    "subnet[1].pool[0].assigned-addresses": [
        [
            0,
            "2023-06-13 20:42:46.836205"
        ]
    ],
    ],
```

(continues on next page)

(continued from previous page)

```

"subnet[1].pool[0].cumulative-assigned-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836137"
  ]
],
"subnet[1].pool[0].declined-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836213"
  ]
],
"subnet[1].pool[0].reclaimed-declined-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836225"
  ]
],
"subnet[1].pool[0].reclaimed-leases": [
  [
    0,
    "2023-06-13 20:42:46.836236"
  ]
],
"subnet[1].pool[0].total-addresses": [
  [
    11010049,
    "2023-06-13 20:42:46.836128"
  ]
],
"subnet[1].reclaimed-declined-addresses": [
  [
    0,
    "2023-06-13 20:42:46.836186"
  ]
],
"subnet[1].reclaimed-leases": [
  [
    0,
    "2023-06-13 20:42:46.836194"
  ]
],
"subnet[1].total-addresses": [
  [
    11010049,
    "2023-06-13 20:42:46.836083"
  ]
],
"subnet[1].v4-lease-reuses": [
  [
    0,
    "2023-06-13 20:42:46.836105"
  ]
]

```

(continues on next page)

(continued from previous page)

```
    ],
    "subnet[1].v4-reservation-conflicts": [
        [
            0,
            "2023-06-13 20:42:46.836111"
        ]
    ],
    "v4-allocation-fail": [
        [
            0,
            "2023-06-13 20:42:46.616358"
        ]
    ],
    "v4-allocation-fail-classes": [
        [
            0,
            "2023-06-13 20:42:46.616363"
        ]
    ],
    "v4-allocation-fail-no-pools": [
        [
            0,
            "2023-06-13 20:42:46.616368"
        ]
    ],
    "v4-allocation-fail-shared-network": [
        [
            0,
            "2023-06-13 20:42:46.616372"
        ]
    ],
    "v4-allocation-fail-subnet": [
        [
            0,
            "2023-06-13 20:42:46.616376"
        ]
    ],
    "v4-lease-reuses": [
        [
            0,
            "2023-06-13 20:42:46.616410"
        ]
    ],
    "v4-reservation-conflicts": [
        [
            0,
            "2023-06-13 20:42:46.616412"
        ]
    ]
}
}
```

(continues on next page)

(continued from previous page)

```

or
{
  "result": 0,
  "arguments": {
    "cumulative-assigned-nas": [
      [
        0,
        "2023-06-13 21:28:57.196757"
      ]
    ],
    "cumulative-assigned-pds": [
      [
        0,
        "2023-06-13 21:28:57.196758"
      ]
    ],
    "declined-addresses": [
      [
        0,
        "2023-06-13 21:28:57.196754"
      ]
    ],
    "pkt6-advertise-received": [
      [
        0,
        "2023-06-13 21:28:57.177731"
      ]
    ],
    "pkt6-advertise-sent": [
      [
        0,
        "2023-06-13 21:28:57.177739"
      ]
    ],
    "pkt6-decline-received": [
      [
        0,
        "2023-06-13 21:28:57.177739"
      ]
    ],
    "pkt6-dhcpv4-query-received": [
      [
        0,
        "2023-06-13 21:28:57.177740"
      ]
    ],
    "pkt6-dhcpv4-response-received": [
      [
        0,
        "2023-06-13 21:28:57.177740"
      ]
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```
    ],
    "pkt6-dhcpv4-response-sent": [
        [
            0,
            "2023-06-13 21:28:57.177741"
        ]
    ],
    "pkt6-infrequest-received": [
        [
            0,
            "2023-06-13 21:28:57.177742"
        ]
    ],
    "pkt6-parse-failed": [
        [
            0,
            "2023-06-13 21:28:57.177742"
        ]
    ],
    "pkt6-rebind-received": [
        [
            0,
            "2023-06-13 21:28:57.177743"
        ]
    ],
    "pkt6-receive-drop": [
        [
            0,
            "2023-06-13 21:28:57.177743"
        ]
    ],
    "pkt6-received": [
        [
            0,
            "2023-06-13 21:28:57.177744"
        ]
    ],
    "pkt6-release-received": [
        [
            0,
            "2023-06-13 21:28:57.177744"
        ]
    ],
    "pkt6-renew-received": [
        [
            0,
            "2023-06-13 21:28:57.177745"
        ]
    ],
    "pkt6-reply-received": [
        [
```

(continues on next page)

(continued from previous page)

```

        0,
        "2023-06-13 21:28:57.177745"
    ]
],
"pkt6-reply-sent": [
    [
        0,
        "2023-06-13 21:28:57.177746"
    ]
],
"pkt6-request-received": [
    [
        0,
        "2023-06-13 21:28:57.177747"
    ]
],
"pkt6-sent": [
    [
        0,
        "2023-06-13 21:28:57.177747"
    ]
],
"pkt6-solicit-received": [
    [
        0,
        "2023-06-13 21:28:57.177748"
    ]
],
"pkt6-unknown-received": [
    [
        0,
        "2023-06-13 21:28:57.177748"
    ]
],
"reclaimed-declined-addresses": [
    [
        0,
        "2023-06-13 21:28:57.196755"
    ]
],
"reclaimed-leases": [
    [
        0,
        "2023-06-13 21:28:57.196756"
    ]
],
"subnet[1].assigned-nas": [
    [
        0,
        "2023-06-13 21:28:57.196760"
    ]
],
],

```

(continues on next page)

(continued from previous page)

```

"subnet[1].assigned-pds": [
    [
        0,
        "2023-06-13 21:28:57.196761"
    ]
],
"subnet[1].cumulative-assigned-nas": [
    [
        0,
        "2023-06-13 21:28:57.196727"
    ]
],
"subnet[1].cumulative-assigned-pds": [
    [
        0,
        "2023-06-13 21:28:57.196729"
    ]
],
"subnet[1].declined-addresses": [
    [
        0,
        "2023-06-13 21:28:57.196763"
    ]
],
"subnet[1].pd-pool[0].assigned-pds": [
    [
        0,
        "2023-06-13 21:28:57.196785"
    ]
],
"subnet[1].pd-pool[0].cumulative-assigned-pds": [
    [
        0,
        "2023-06-13 21:28:57.196744"
    ]
],
"subnet[1].pd-pool[0].reclaimed-leases": [
    [
        0,
        "2023-06-13 21:28:57.196789"
    ]
],
"subnet[1].pd-pool[0].total-pds": [
    [
        256,
        "2023-06-13 21:28:57.196741"
    ]
],
"subnet[1].pool[0].assigned-nas": [
    [
        0,
        "2023-06-13 21:28:57.196773"
    ]
]

```

(continues on next page)

(continued from previous page)

```

    ]
  ],
  "subnet[1].pool[0].cumulative-assigned-nas": [
    [
      0,
      "2023-06-13 21:28:57.196739"
    ]
  ],
  "subnet[1].pool[0].declined-addresses": [
    [
      0,
      "2023-06-13 21:28:57.196775"
    ]
  ],
  "subnet[1].pool[0].reclaimed-declined-addresses": [
    [
      0,
      "2023-06-13 21:28:57.196779"
    ]
  ],
  "subnet[1].pool[0].reclaimed-leases": [
    [
      0,
      "2023-06-13 21:28:57.196783"
    ]
  ],
  "subnet[1].pool[0].total-nas": [
    [
      281474976710656,
      "2023-06-13 21:28:57.196736"
    ]
  ],
  "subnet[1].reclaimed-declined-addresses": [
    [
      0,
      "2023-06-13 21:28:57.196766"
    ]
  ],
  "subnet[1].reclaimed-leases": [
    [
      0,
      "2023-06-13 21:28:57.196770"
    ]
  ],
  "subnet[1].total-nas": [
    [
      281474976710656,
      "2023-06-13 21:28:57.196720"
    ]
  ],
  "subnet[1].total-pds": [
    [

```

(continues on next page)

(continued from previous page)

```

        256,
        "2023-06-13 21:28:57.196724"
    ],
    "subnet[1].v6-ia-na-lease-reuses": [
        [
            0,
            "2023-06-13 21:28:57.196731"
        ]
    ],
    "subnet[1].v6-ia-pd-lease-reuses": [
        [
            0,
            "2023-06-13 21:28:57.196733"
        ]
    ],
    "v6-allocation-fail": [
        [
            0,
            "2023-06-13 21:28:57.177749"
        ]
    ],
    "v6-allocation-fail-classes": [
        [
            0,
            "2023-06-13 21:28:57.177755"
        ]
    ],
    "v6-allocation-fail-no-pools": [
        [
            0,
            "2023-06-13 21:28:57.177756"
        ]
    ],
    "v6-allocation-fail-shared-network": [
        [
            0,
            "2023-06-13 21:28:57.177756"
        ]
    ],
    "v6-allocation-fail-subnet": [
        [
            0,
            "2023-06-13 21:28:57.177757"
        ]
    ],
    "v6-ia-na-lease-reuses": [
        [
            0,
            "2023-06-13 21:28:57.177757"
        ]
    ],

```

(continues on next page)

(continued from previous page)

```

        "v6-ia-pd-lease-reuses": [
            [
                0,
                "2023-06-13 21:28:57.177758"
            ]
        ]
    }
}

or

{
    "result": 0,
    "arguments": {
        "ncr-error": [
            [
                0,
                "2023-06-13 21:42:54.627751"
            ]
        ],
        "ncr-invalid": [
            [
                0,
                "2023-06-13 21:42:54.627749"
            ]
        ],
        "ncr-received": [
            [
                0,
                "2023-06-13 21:42:54.627737"
            ]
        ],
        "queue-mgr-queue-full": [
            [
                0,
                "2023-06-13 21:42:54.627737"
            ]
        ],
        "update-error": [
            [
                0,
                "2023-06-13 21:42:54.627759"
            ]
        ],
        "update-sent": [
            [
                0,
                "2023-06-13 21:42:54.627752"
            ]
        ],
        "update-signed": [

```

(continues on next page)

(continued from previous page)

```

        0,
        "2023-06-13 21:42:54.627753"
    ],
    "update-success": [
        [
            0,
            "2023-06-13 21:42:54.627755"
        ]
    ],
    "update-timeout": [
        [
            0,
            "2023-06-13 21:42:54.627757"
        ]
    ],
    "update-unsigned": [
        [
            0,
            "2023-06-13 21:42:54.627754"
        ]
    ]
  ]
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.175 statistic-remove

This command deletes a single statistic. It takes a single string parameter called name that specifies the statistic name.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-remove command*

Command syntax:

```

{
  "command": "statistic-remove",
  "arguments": {
    "name": "pkt4-received"
  }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

If the specific statistic is found and its removal is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.176 statistic-remove-all

(Deprecated) This command deletes all statistics.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-remove-all command*

Command syntax:

```
{
  "command": "statistic-remove-all",
  "arguments": { }
}
```

If the removal of all statistics is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.177 statistic-reset

This command sets the specified statistic to its neutral value: 0 for integer, 0.0 for float, 0h0m0s0us for time duration, and "" for string type. It takes a single string parameter called name that specifies the statistic name.

Supported by: *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-reset command*

Command syntax:

```
{
  "command": "statistic-reset",
  "arguments": {
    "name": "pkt4-received"
  }
}
```

If the specific statistic is found and the reset is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.178 statistic-reset-all

This command sets all statistics to their neutral values: 0 for integer, 0.0 for float, 0h0m0s0us for time duration, and "" for string type.

Supported by: *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.0.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-reset-all command*

Command syntax:

```
{
  "command": "statistic-reset-all",
  "arguments": { }
}
```

If the operation is successful, the server responds with a status of 0, indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.179 statistic-sample-age-set

This command sets a time-based limit for a single statistic. It takes two parameters: a string called name and an integer value called duration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-sample-age-set command*

Command syntax:

```
{
  "command": "statistic-sample-age-set",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```
    "name": "pkt4-received",
    "duration": 1245
  }
}
```

The server responds with a message about a successfully set limit for the given statistic, with a result of 0 indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.180 statistic-sample-age-set-all

This command sets a time-based limit for all statistics. It takes a single integer parameter called duration.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-sample-age-set-all command*

Command syntax:

```
{
  "command": "statistic-sample-age-set-all",
  "arguments": {
    "duration": 1245
  }
}
```

The server responds with a message about successfully set limits for all statistics, with a result of 0 indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.181 statistic-sample-count-set

This command sets a size-based limit for a single statistic. It takes two parameters: a string called name and an integer value called max-samples.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-sample-count-set command*

Command syntax:

```
{
  "command": "statistic-sample-count-set",
  "arguments": {
    "name": "pkt4-received",
    "max-samples": 100
  }
}
```

The server responds with a message about a successfully set limit for the given statistic, with a result of 0 indicating success, and an empty parameters field. If an error is encountered (e.g. the requested statistic was not found), the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)

- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.182 statistic-sample-count-set-all

This command sets a size-based limit for all statistics. It takes a single integer parameter called max-samples.

Supported by: *kea-dhcp4*, *kea-dhcp6*

Availability: 1.6.0 (built-in)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *statistic-sample-count-set-all command*

Command syntax:

```
{
  "command": "statistic-sample-count-set-all",
  "arguments": {
    "max-samples": 100
  }
}
```

The server responds with a message about successfully set limits for all statistics, with a result of 0 indicating success, and an empty parameters field. If an error is encountered, the server returns a status code of 1 (error) and the text field contains the error description.

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.183 status-get

This command returns server's runtime information. It takes no arguments.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.7.3 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *status-get command*

Command syntax:

```
{
  "command": "status-get"
}
```

Response syntax:

```
{
  "result": <integer>,
  "arguments": {
    "dhcp-state": {
      "disabled-by-db-connection": [],
      "disabled-by-local-command": [],
      "disabled-by-remote-command": [],
      "disabled-by-user": false,
      "globally-disabled": false
    },
    "extended-info-tables": <whether relay information is held in lease tables>,
    "high-availability": [
      {
        "ha-mode": <HA mode configured for this relationship>,
        "ha-servers": {
          "local": {
            "role": <role of this server as in the configuration file>,
            "scopes": <list of scope names served by this server>,
            "server-name": <name of the local server>,
            "state": <HA state name of the server receiving the command>,
            "system-time": <system time in format '%Y-%m-%d %H:%M:%S' on UTC_
↪timezone>
          },
          "remote": {
            "age": <the age of the remote status in seconds>,
            "analyzed-packets": <number of packets sent to the partner_
↪server since communication was interrupted>,
            "clock-skew": <difference in seconds between local and partner_
↪server times>,
            "communication-interrupted": <whether communication did not_
↪happen for more than max-response-delay milliseconds>,
            "connecting-clients": <number of different clients getting a_
↪lease from partner>,
            "in-touch": <indicates if this server communicated with partner>,
            "last-scopes": <list of scopes served by partner>,
            "last-state": <HA state name of the partner>,
            "role": <partner role>,
            "server-name": <name of the partner server>,
            "system-time": <system time in format '%Y-%m-%d %H:%M:%S' on UTC_
↪timezone>,
            "unacked-clients": <number of unacked clients>,
            "unacked-clients-left": <how many more clients have to be_
↪unacked before partner-down state>
          }
        ]
      }
    ],
  },
}
```

(continues on next page)

(continued from previous page)

```

    "multi-threading-enabled": true,
    "packet-queue-size": 64,
    "packet-queue-statistics": [ 1.2, 2.3, 3.4 ],
    "pid": <integer>,
    "reload": <time since reload in seconds>,
    "sockets": {
        "errors": <list of errors received during the last attempt to open all
↪ sockets; only appears when status is failed or retrying>,
        "status": <ready, retrying, or failed>
    },
    "thread-pool-size": 4,
    "uptime": <uptime in seconds>
}

```

If the libdhcp_ha (High Availability) hooks library is loaded by the DHCP server receiving this command the response also includes the HA specific status information of the server receiving the command and its partner's status.

24.184 subnet4-add

This command creates and adds a new subnet to the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-add command*

Command syntax:

```

{
    "command": "subnet4-add",
    "arguments": {
        "subnet4": [ {
            "id": 123,
            "subnet": "10.20.30.0/24",
            ...
        } ]
    }
}

```

Response syntax:

```

{
    "result": 0,
    "text": "IPv4 subnet added",
    "arguments": {
        "subnets": [
            {
                "id": 123,
                "subnet": "10.20.30.0/24"
            }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.185 subnet4-del

This command removes a subnet from the server's configuration. This command has no effect on other configured subnets, but removing a subnet has certain implications which the server's administrator should be aware of.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-del command*

Command syntax:

```

{
  "command": "subnet4-del",
  "arguments": {
    "id": 123
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "IPv4 subnet 192.0.2.0/24 (id 123) deleted",
  "arguments": {
    "subnets": [
      {
        "id": 123,
        "subnet": "192.0.2.0/24"
      }
    ]
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success

- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.186 subnet4-delta-add

This command updates (adds or overwrites) parts of a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp4*

Availability: 2.1.7 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-delta-add command*

Command syntax:

```
{
  "command": "subnet4-delta-add",
  "arguments": {
    "subnet4": [ {
      "id": 123,
      "subnet": "10.20.30.0/24",
      ...
    } ]
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)

- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.187 subnet4-delta-del

This command updates (removes) parts of a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp4*

Availability: 2.1.7 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-delta-del command*

Command syntax:

```
{
  "command": "subnet4-delta-del",
  "arguments": {
    "subnet4": [ {
      "id": 123,
      "subnet": "10.20.30.0/24",
      ...
    } ]
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.188 subnet4-get

This command retrieves detailed information about the specified subnet. This command usually follows `subnet4-list`, which discovers available subnets with their respective subnet identifiers and prefixes.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-get command*

Command syntax:

```
{
  "command": "subnet4-get",
  "arguments": {
    "id": 10
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Info about IPv4 subnet 10.0.0.0/8 (id 10) returned",
  "arguments": {
    "subnets": [
      {
        "subnet": "10.0.0.0/8",
        "id": 1,
        "option-data": [{
          ...
        }],
        ...
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.189 subnet4-list

This command lists all currently configured subnets. The subnets are returned in a brief format, i.e. a subnet identifier, subnet prefix, and shared network name (if available) are included for each subnet.

Supported by: *kea-dhcp4*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-list command*

Command syntax:

```
{
  "command": "subnet4-list"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "2 IPv4 subnets found",
  "arguments": {
    "subnets": [
      {
        "id": 10,
        "subnet": "10.0.0.0/8"
      },
      {
        "id": 100,
        "subnet": "192.0.2.0/24",
        "shared-network-name": "example-shared-network"
      }
    ]
  }
}
```

If no IPv4 subnets are found, an error code is returned along with the error description.

24.190 subnet4-select-test

This command returns the result of DHCPv4 subnet selection

Supported by: *kea-dhcp4*

Availability: 2.7.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-select-test command*

Command syntax:

```
{
  "command": "subnet4-select-test",
  "arguments": {
```

(continues on next page)

(continued from previous page)

```

    "interface": <string>,
    "address": <string>,
    "relay": <string>,
    "local": <string>,
    "remote": <string>,
    "link": <string>,
    "subnet": <string>,
    "classes": [ <strings> ]
  }
}

```

Possible parameters are the incoming interface name, client address, relay address, client classes, etc.

Response syntax:

```
{ "result": 1, "text": <localization result> }
```

3 forms for selected shared network, subnet or nothing..

24.191 subnet4-update

This command updates (overwrites) a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp4*

Availability: 1.6.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet4-update command*

Command syntax:

```

{
  "command": "subnet4-update",
  "arguments": {
    "subnet4": [ {
      "id": 123,
      "subnet": "10.20.30.0/24",
      ...
    } ]
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "IPv4 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 123,
        "subnet": "10.20.30.0/24"
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  ]
}
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.192 subnet4o6-select-test

This command returns the result of DHCPv4o6 subnet selection

Supported by: *kea-dhcp4*

Availability: 2.7.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet4o6-select-test command*

Command syntax:

```

{
  "command": "subnet4o6-select-test",
  "arguments": {
    "interface": <string>,
    "interface-id": <string>,
    "address": <string>,
    "relay": <string>,
    "local": <string>,
    "remote": <string>,
    "link": <string>,
    "subnet": <string>,
    "classes": [ <strings> ]
  }
}

```

Possible parameters are the remote address, interface id option, interface name (the three from the DHCPv6 server), etc.

Response syntax:

```
{ "result": 1, "text": <localization result> }
```

3 forms for selected shared network, subnet or nothing..

24.193 subnet6-add

This command creates and adds a new subnet to the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-add command*

Command syntax:

```
{
  "command": "subnet6-add",
  "arguments": {
    "subnet6": [ {
      "id": 234,
      "subnet": "2001:db8:1::/64",
      ...
    } ]
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet added",
  "arguments": {
    "subnets": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.194 subnet6-del

This command removes a subnet from the server's configuration. This command has no effect on other configured subnets, but removing a subnet has certain implications which the server's administrator should be aware of.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-del command*

Command syntax:

```
{
  "command": "subnet6-del",
  "arguments": {
    "id": 234
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet 2001:db8:1::/64 (id 234) deleted",
  "subnets": [
    {
      "id": 234,
      "subnet": "2001:db8:1::/64"
    }
  ]
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.195 subnet6-delta-add

This command updates (adds or overwrites) parts of a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp6*

Availability: 2.1.7 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-delta-add command*

Command syntax:

```
{
  "command": "subnet6-delta-add",
  "arguments": {
    "subnet6": [ {
      "id": 234,
      "subnet": "2001:db8:1::/64",
      ...
    } ]
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.196 subnet6-delta-del

This command updates (removes) parts of a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp6*

Availability: 2.1.7 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-delta-del command*

Command syntax:

```
{
  "command": "subnet6-delta-del",
```

(continues on next page)

(continued from previous page)

```

    "arguments": {
      "subnet6": [ {
        "id": 234,
        "subnet": "2001:db8:1::/64",
        ...
      } ]
    }
  }
}

```

Response syntax:

```

{
  "result": 0,
  "text": "IPv6 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}

```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.197 subnet6-get

This command retrieves detailed information about the specified subnet. This command usually follows `subnet6-list`, which discovers available subnets with their respective subnet identifiers and prefixes.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-get command*

Command syntax:

```

{
  "command": "subnet6-get",
  "arguments": {
    "id": 11
  }
}

```

(continues on next page)

(continued from previous page)

```
}
}
```

Response syntax:

```
{
  "result": 0,
  "text": "Info about IPv6 subnet 2001:db8:1::/64 (id 11) returned",
  "arguments": {
    "subnets": [
      {
        "subnet": "2001:db8:1::/64",
        "id": 1,
        "option-data": [{
          ...
        }],
        ...
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.198 subnet6-list

This command lists all currently configured subnets. The subnets are returned in a brief format, i.e. a subnet identifier, subnet prefix, and shared network name are included for each subnet.

Supported by: *kea-dhcp6*

Availability: 1.3.0 (*subnet_cmds* hook library)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-list command*

Command syntax:

```
{
  "command": "subnet6-list"
}
```

Response syntax:

```
{
  "result": 0,
  "text": "2 IPv6 subnets found",
  "arguments": {
    "subnets": [
      {
        "id": 11,
        "subnet": "2001:db8:1::/64"
      },
      {
        "id": 233,
        "subnet": "3000::/16",
        "shared-network-name": "example-shared-network"
      }
    ]
  }
}
```

If no IPv6 subnets are found, an error code is returned along with the error description.

24.199 subnet6-select-test

This command returns the result of DHCPv6 subnet selection

Supported by: *kea-dhcp6*

Availability: 2.7.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-select-test command*

Command syntax:

```
{
  "command": "subnet6-select-test",
  "arguments": {
    "interface": <string>,
    "interface-id": <string>,
    "remote": <string>,
    "link": <string>,
    "classes": [ <strings> ]
  }
}
```

Possible parameters are the incoming interface name, remote address, relay link address, client classes, etc.

Response syntax:

```
{ "result": 1, "text": <localization result> }
```

3 forms for selected shared network, subnet or nothing..

24.200 subnet6-update

This command updates (overwrites) a single subnet in the existing server configuration. This operation has no impact on other subnets.

Supported by: *kea-dhcp6*

Availability: 1.6.0 (*subnet_cmds* hook library)

Access: write (*parameter ignored in this Kea version*)

Description and examples: see *subnet6-update command*

Command syntax:

```
{
  "command": "subnet6-update",
  "arguments": {
    "subnet6": [ {
      "id": 234,
      "subnet": "2001:db8:1::/64",
      ...
    } ]
  }
}
```

Response syntax:

```
{
  "result": 0,
  "text": "IPv6 subnet updated",
  "arguments": {
    "subnets": [
      {
        "id": 234,
        "subnet": "2001:db8:1::/64"
      }
    ]
  }
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

24.201 version-get

This command returns extended information about the Kea version that is running. The returned string is the same as if Kea were run with the `-V` command-line option.

Supported by: *kea-ctrl-agent*, *kea-dhcp-ddns*, *kea-dhcp4*, *kea-dhcp6*

Availability: 1.2.0 (built-in)

Access: read (*parameter ignored in this Kea version*)

Description and examples: see *version-get command*

Command syntax:

```
{
  "command": "version-get"
}
```

Response syntax:

```
{
  "result": <integer>,
  "text": "<string>"
}
```

Result is an integer representation of the status. Currently supported statuses are:

- 0 - success
- 1 - error
- 2 - unsupported
- 3 - empty (command was completed successfully, but no data was affected or returned)
- 4 - conflict (command could not apply requested configuration changes because they were in conflict with the server state)

25.1 kea-dhcp4 - DHCPv4 server in Kea

25.1.1 Synopsis

kea-dhcp4 [-v] [-V] [-W] [-d] [-c config-file] [-t config-file] [-p server-port-number] [-P client-port-number]

25.1.2 Description

The kea-dhcp4 daemon provides the DHCPv4 server implementation.

25.1.3 Arguments

The arguments are as follows:

- v**
Displays the Kea version.
- V**
Displays the extended Kea version.
- W**
Displays the configuration report.
- d**
Enables the debug mode with extra verbosity.
- c config-file**
Specifies the configuration file with the configuration for the DHCPv4 server. It may also contain configuration entries for other Kea services.
- t config-file**
Checks the configuration file and reports the first error, if any. Note that not all parameters are completely checked; in particular, service and control channel sockets are not opened, and hook libraries are not loaded.
- T config-file**
Checks the configuration file and reports the first error, if any. It performs extra checks beyond what -t offers, such as establishing database connections (for the lease backend, host reservations backend, configuration backend, and forensic logging backend), loading hook libraries, parsing hook-library configurations, etc. It does not open UNIX or TCP/UDP sockets, nor does it open or rotate files, as any of these actions could interfere with a running process on the same machine.
- p server-port-number**
Specifies the server port number (1-65535) on which the server listens. This is useful for testing purposes only.
- P client-port-number**
Specifies the client port number (1-65535) to which the server responds. This is useful for testing purposes only.

25.1.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.1.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.1.6 History

The b10-dhcp4 daemon was first coded in November 2011 by Tomek Mrugalski.

In mid-2014, Kea was decoupled from the BIND 10 framework and became a standalone DHCP server. The DHCPv4 server binary was renamed to kea-dhcp4. Kea 1.0.0 was released in December 2015.

25.1.7 See Also

kea-dhcp6(8), *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-netconf(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.2 kea-dhcp6 - DHCPv6 server in Kea

25.2.1 Synopsis

kea-dhcp6 [-v] [-V] [-W] [-d] [-c config-file] [-t config-file] [-p server-port-number] [-P client-port-number]

25.2.2 Description

The kea-dhcp6 daemon provides the DHCPv6 server implementation.

25.2.3 Arguments

The arguments are as follows:

- v**
Displays the Kea version.
- V**
Displays the extended Kea version.
- W**
Displays the configuration report.
- d**
Enables the debug mode with extra verbosity.

-c config-file

Specifies the configuration file with the configuration for the DHCPv6 server. It may also contain configuration entries for other Kea services.

-t config-file

Checks the configuration file and reports the first error, if any. Note that not all parameters are completely checked; in particular, service and control channel sockets are not opened, and hook libraries are not loaded.

-T config-file

Checks the configuration file and reports the first error, if any. It performs extra checks beyond what -t offers, such as establishing database connections (for the lease backend, host reservations backend, configuration backend, and forensic logging backend), loading hook libraries, parsing hook-library configurations, etc. It does not open UNIX or TCP/UDP sockets, nor does it open or rotate files, as any of these actions could interfere with a running process on the same machine.

-p server-port-number

Specifies the server port number (1-65535) on which the server listens. This is useful for testing purposes only.

-P client-port-number

Specifies the client port number (1-65535) to which the server responds. This is useful for testing purposes only.

25.2.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.2.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.2.6 History

The **b10-dhcp6** daemon was first coded in June 2011 by Tomek Mrugalski.

In mid-2014, Kea was decoupled from the BIND 10 framework and became a standalone DHCP server. The DHCPv6 server binary was renamed to **kea-dhcp6**. Kea 1.0.0 was released in December 2015.

25.2.7 See Also

kea-dhcp4(8), *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-netconf(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.3 kea-ctrl-agent - Control Agent process in Kea

25.3.1 Synopsis

kea-ctrl-agent [-v] [-V] [-W] [-d] [-c config-file] [-t config-file]

25.3.2 Description

The **kea-ctrl-agent** provides a REST service for controlling Kea services. The received HTTP requests are de-encapsulated and forwarded to the respective Kea services in JSON format. Received JSON responses are encapsulated within HTTP responses and returned to the controlling entity. Some commands may be handled by the Control Agent directly, and not forwarded to any Kea service.

25.3.3 Arguments

The arguments are as follows:

- v**
Displays the Kea version.
- V**
Displays the extended Kea version.
- W**
Displays the configuration report.
- d**
Sets the logging level to debug with extra verbosity. This is primarily for development purposes in stand-alone mode.
- c config-file**
Specifies the file with the configuration for the Control Agent server. It may also contain configuration entries for other Kea services.
- t config-file**
Checks the syntax of the configuration file and reports the first error, if any. Note that not all parameters are completely checked; in particular, service and client sockets are not opened, and hook libraries are not loaded.

25.3.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.3.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.3.6 History

The `kea-ctrl-agent` was first coded in December 2016 by Marcin Siodelski.

25.3.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.4 keactrl - Shell script for managing Kea

25.4.1 Synopsis

keactrl [**command**] [-c keactrl-config-file] [-s server[,server,...]] [-v]

25.4.2 Description

`keactrl` is a shell script which controls the startup, shutdown, and reconfiguration of the Kea servers (`kea-dhcp4`, `kea-dhcp6`, `kea-dhcp-ddns`, `kea-ctrl-agent`, and `kea-netconf`). It also provides a way to check the current status of the servers and determine the configuration files in use.

25.4.3 Configuration File

Depending on the user's requirements, not all of the available servers need be run. The `keactrl` configuration file specifies which servers are enabled and which are disabled. By default the configuration file is `[kea-install-dir]/etc/kea/keactrl.conf`.

See the Kea Administrator Reference Manual for documentation of the parameters in the `keactrl` configuration file.

25.4.4 Options

command

Specifies the command to be issued to the servers. It can be one of the following:

start

Starts the servers.

stop

Stops the servers.

reload

Instructs the servers to re-read the Kea configuration file. This command is not supported by the NETCONF agent.

status

Prints the status of the servers.

-c|--ctrl-config keactrl-config-file

Specifies the `keactrl` configuration file. Without this switch, `keactrl` uses the file `[kea-install-dir]/etc/kea/keactrl.conf`.

-s|--server server[,server,...]

Specifies a subset of the enabled servers to which the command should be issued. The list of servers should be separated by commas, with no intervening spaces. Acceptable values are:

dhcp4

DHCPv4 server (`kea-dhcp4`).

dhcp6

DHCPv6 server (kea-dhcp6).

dhcp_ddns

DHCP DDNS server (kea-dhcp-ddns).

ctrl_agent

Control Agent (kea-ctrl-agent).

netconf

NETCONF agent (kea-netconf).

all

All servers, including NETCONF if it was configured to be built. This is the default.

-v|--version

Displays the Kea version.

-V

Displays the extended Kea version.

25.4.5 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.4.6 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.4.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *kea-netconf(8)*, *perfdhcp(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.5 kea-admin - Shell script for managing Kea databases

25.5.1 Synopsis

kea-admin [command] [backend] [-h database_host] [-P database_port] [-u database_username] [-p [database_password]] [-n database_name] [-d script_directory] [-v] [-x extra_argument [-x extra_argument ...]] [-4 | -6] [-i input_file] [-o output_file] [-y]

25.5.2 Description

`kea-admin` is a shell script that offers database maintenance. In particular, it features database initialization, database version checking, and database schema upgrading.

25.5.3 Arguments

command

Specifies the command to be issued to the servers. It can be one of the following:

db-init

Initializes a new database schema. This is useful during a new Kea installation. The database is initialized to the latest version supported by the version of the software being installed.

db-version

Reports the database backend version number. This is not necessarily equal to the Kea version number, as each backend has its own versioning scheme.

db-upgrade

Conducts a database schema upgrade. This is useful when upgrading Kea.

lease-dump

Dumps the contents of the lease database (for MySQL or PostgreSQL backends) to a CSV (comma-separated values) text file. (Support for the Cassandra backend has been deprecated.) The first line of the file contains the column names. This can be used as a way to switch from a database backend to a memfile backend. It can also be used as a diagnostic tool, as it provides a portable form of the lease data. There are other mandatory arguments that must be used together with this command: either `-4` or `-6` must be specified, and either `-o` or `--output` must be provided.

lease-upload

Uploads leases from a CSV (comma-separated values) text file to a MySQL or a PostgreSQL lease database. The CSV file needs to be in memfile format. There are other mandatory arguments that must be used together with this command: either `-4` or `-6` must be specified, and either `-i` or `--input` must be provided.

stats-recount

Recounts lease statistics for a MySQL or PostgreSQL database.

backend

Specifies the backend type. Currently allowed backends are: `memfile`, `mysql`, and `pgsql`; `cql` has been deprecated.

-h|--host hostname

Specifies the hostname when connecting to a database. The default value is `localhost`.

-i|--input input_file

Specifies the CSV (comma-separated values) text file with the leases to be uploaded; required for `lease-upload`.

-P|--port port

Specifies the port when connecting to a database. If not specified, the default value chosen by the database client is used.

-u|--user username

Specifies the username when connecting to a database. The default value is `keatest`.

-p|--password password

Specifies the password when connecting to a database. If only `-p` or `--password` is given, the user is prompted for a password. If not specified at all, the `KEA_ADMIN_DB_PASSWORD` environment variable is checked for a value and used if it exists. Otherwise the default value of `1234` is used.

-n|--name database-name

Specifies the name of the database to connect to. The default value is `keatest`.

-d|--directory script-directory

Specifies the override scripts directory. That script is used during upgrades, database initialization, and possibly other operations. The default value is (prefix)/share/kea/scripts/.

-o|--output output_file

Specifies the file to which the lease data will be dumped. Required for lease-dump.

-v|--version

Displays the Kea version.

-V

Displays the extended Kea version.

-4

Directs kea-admin to lease-dump the DHCPv4 leases. Incompatible with the -6 option.

-6

Directs kea-admin to lease-dump the DHCPv6 leases. Incompatible with the -4 option.

-x|--extra

Specifies an extra argument to pass to the database command tool e.g. to invoke `mysql` with the `--ssl` argument. This can be repeated to pass more than one argument. Quotes are not preserved. Avoid commands containing spaces.

-y|--yes

Assume yes on overwriting temporary files.

25.5.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.5.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.5.6 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-netconf(8)*, Kea Administrator Reference Manual.

25.6 kea-dhcp-ddns - DHCP-DDNS process in Kea

25.6.1 Synopsis

kea-dhcp-ddns [-v] [-V] [-W] [-d] [-c config-file] [-t config-file]

25.6.2 Description

The `kea-dhcp-ddns` service process requests an update of DNS mapping based on DHCP lease-change events. It runs as a separate process that expects to receive Name Change Requests from Kea DHCP servers.

25.6.3 Arguments

The arguments are as follows:

- v**
Displays the Kea version.
- V**
Displays the extended Kea version.
- W**
Displays the configuration report.
- d**
Sets the logging level to debug with extra verbosity. This is primarily for development purposes in stand-alone mode.
- c config-file**
Specifies the configuration file with the configuration for the DHCP-DDNS server. It may also contain configuration entries for other Kea services.
- t config-file**
Checks the syntax of the configuration file and reports the first error, if any. Note that not all parameters are completely checked; in particular, a service socket is not opened.

25.6.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.6.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.6.6 History

The `b10-dhcp-ddns` process was first coded in May 2013 by Thomas Markwalder.

Kea became a standalone server and the BIND 10 framework was removed. The DHCP-DDNS server binary was renamed to `kea-dhcp-ddns` in July 2014. Kea 1.0.0 was released in December 2015.

25.6.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-netconf(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.7 kea-lfc - Lease File Cleanup process in Kea

25.7.1 Synopsis

kea-lfc [-4**|-6**] [-c config-file] [-p pid-file] [-x previous-file] [-i copy-file] [-o output-file] [-f finish-file] [-v] [-V] [-W] [-d] [-h]

25.7.2 Description

The *kea-lfc* service process removes redundant information from the files used to provide persistent storage for the memfile database backend. The service is written to run as a stand-alone process. While it can be started externally, there is usually no need to do this. It is run periodically by the Kea DHCP servers.

25.7.3 Arguments

The arguments are as follows:

-4 | -6

Indicates the protocol version of the lease files; must be either 4 or 6.

-c config-file

Specifies the file with the configuration for the *kea-lfc* process. It may also contain configuration entries for other Kea services. Currently *kea-lfc* gets all of its arguments from the command line.

-p pid-file

Specifies the PID file. When the *kea-lfc* process starts, it attempts to determine if another instance of the process is already running, by examining the PID file. If one is already running, the new process is terminated. If one is not running, Kea writes its PID into the PID file.

-x previous-file

Specifies the previous or ex-lease file. When *kea-lfc* starts, this is the result of any previous run of *kea-lfc*; when *kea-lfc* finishes, it is the result of the current run. If *kea-lfc* is interrupted before completing, this file may not exist.

-i copy-file

Specifies the input or copy of lease file. Before the DHCP server invokes *kea-lfc*, it moves the current lease file here and then calls *kea-lfc* with this file.

-o output-file

Specifies the output lease file, which is the temporary file *kea-lfc* should use to write the leases. Once this file is finished writing, it is moved to the finish file (see below).

-f finish-file

Specifies the finish or completion file, another temporary file *kea-lfc* uses for bookkeeping. When *kea-lfc* finishes writing the output file, it moves it to this file name. After *kea-lfc* finishes deleting the other files (previous and input), it moves this file to the previous lease file. By moving the files in this fashion, the *kea-lfc* and the DHCP server processes can determine the correct file to use even if one of the processes was interrupted before completing its task.

-v

Displays the Kea version.

-V

Displays the extended Kea version.

- W**
Displays the configuration report.
- d**
Sets the logging level to debug with extra verbosity. This is primarily for development purposes in stand-alone mode.
- h**
Causes the usage string to be printed.

25.7.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.7.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.7.6 History

The `kea-lfc` process was first coded in January 2015 by the ISC Kea/DHCP team.

25.7.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-netconf(8)*, Kea Administrator Reference Manual.

25.8 kea-shell - Text client for kea-dhcp4, kea-dhcp6 and kea-dhcp-ddns servers

25.8.1 Synopsis

kea-shell [-h] [-v] [--host] [--port] [--path] [--ca] [--cert] [--key] [--auth-user] [--auth-password] [--auth-password-file] [--timeout] [--service] [command]

25.8.2 Description

The `kea-shell` provides a REST client for the Kea `kea-dhcp4`, `kea-dhcp6` and `kea-dhcp-ddns` servers. Previously it was meant for the Control Agent (CA) daemon, which is deprecated. It takes commands as a command-line parameter that is sent to the Kea server with proper JSON encapsulation. Optional arguments may be specified on the standard input. The request is sent via HTTP and a response is retrieved, displayed on the standard output. Basic HTTP authentication and HTTPS, i.e. TLS transport, are supported.

25.8.3 Arguments

The arguments are as follows:

-h

Displays help regarding command-line parameters.

-v

Displays the Kea version.

--host

Specifies the host to connect to. The Kea server must be running at the specified host. If not specified, 127.0.0.1 is used.

--port

Specifies the TCP port to connect to. The server must be listening at the specified port. If not specified, 8000 is used.

--path

Specifies the path in the URL to connect to. If not specified, an empty path is used. As Kea servers listen at the empty path, this parameter is useful only with a reverse proxy.

--ca

Specifies the file or directory name of the Certification Authority. If not specified, HTTPS is not used.

--cert

Specifies the file name of the user end-entity public key certificate. If specified, the file name of the user key must also be specified.

--key

Specifies the file name of the user key file. If specified, the file name of the user certificate must also be specified. Encrypted key files are not supported.

--auth-user

Specifies the user ID for basic HTTP authentication. If not specified, or specified as the empty string, authentication is not used.

--auth-password

Specifies the password for basic HTTP authentication. If not specified but the user ID is specified, an empty password is used. If used together with `--auth-password-file`, it is disregarded.

--auth-password-file

Specifies a text file whose first line contains the password to be used for basic HTTP authentication. This flag takes precedence over `--auth-password` flag.

--timeout

Specifies the connection timeout, in seconds. The default is 10.

--service

(Obsolete) Kea `kea-dhcp4`, `kea-dhcp6` and `kea-dhcp-ddns` servers ignore this argument. It was meant for communication with deprecated Control Agent daemon. If used with Kea version < 2.7.2: specifies the service that is the target of a command. If not specified, the Control Agent itself is targeted. May be used more than once to specify multiple targets.

command

Specifies the command to be sent to the server. If not specified, `list-commands` is used.

25.8.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.8.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.8.6 History

The `kea-shell` was first coded in March 2017 by Tomek Mrugalski.

25.8.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.9 kea-netconf - NETCONF agent for configuring Kea

25.9.1 Synopsis

kea-netconf [-v] [-V] [-W] [-d] [-c config-file] [-t config-file]

25.9.2 Description

The `kea-netconf` agent provides a YANG/NETCONF interface for the Kea environment.

25.9.3 Arguments

The arguments are as follows:

- v**
Displays the Kea version.
- V**
Displays the extended Kea version.
- W**
Displays the configuration report.
- d**
Enables the debug mode with extra verbosity.
- c config-file**
Specifies the file with the configuration for the NETCONF agent.

-t config-file

Checks the syntax of the configuration file and reports the first error, if any. Note that not all parameters are completely checked; in particular, service and client sockets are not opened, and hook libraries are not loaded.

25.9.4 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.9.5 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.9.6 History

Early prototypes of `kea-netconf` implementation were written during IETF Hackathons in Berlin, London, and Montreal. An actual production-ready implementation was started in August 2018 by Tomek Mrugalski and Francis Dupont.

25.9.7 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *keactrl(8)*, *perfdhcp(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

25.10 perfdhcp - DHCP benchmarking tool

25.10.1 Synopsis

perfdhcp [-1] [-4 | -6] [-A encapsulation-level] [-b base] [-B] [-c] [-C separator] [-d drop-time] [-D max-drop] [-e lease-type] [-E time-offset] [-f renew-rate] [-F release-rate] [-g thread-mode] [-h] [-i] [-I ip-offset] [-J remote-address-list-file] [-l local-address|interface] [-L local-port] [-M mac-list-file] [-n num-request] [-N remote-port] [-O random-offset] [-o code,hexstring] [--or encapsulation-level:code,hexstring] [-p test-period] [-P preload] [-r rate] [-R num-clients] [-s seed] [-S srvid-offset] [--scenario name] [-t report] [-T template-file] [-u] [-v] [-W exit-wait-time] [-w script_name] [-x diagnostic-selector] [-X xid-offset] [server]

25.10.2 Description

perfdhcp is a DHCP benchmarking tool. It provides a way to measure the performance of DHCP servers by generating large amounts of traffic from multiple simulated clients. It is able to test both IPv4 and IPv6 servers, and provides statistics concerning response times and the number of requests that are dropped.

The tool supports two different scenarios, which offer certain behaviors to be tested. By default (the basic scenario), tests are run using the full four-packet exchange sequence (DORA for DHCPv4, SARR for DHCPv6). An option is provided to run tests using the initial two-packet exchange (DO and SA) instead. It is also possible to configure **perfdhcp** to send DHCPv6 RENEW and RELEASE messages at a specified rate, in parallel with the DHCPv6 four-way exchanges.

By default, if there is no response received with one second, a response is considered lost and `perfdhcp` continues with other transactions.

A second scenario, called *avalanche*, is selected via `--scenario avalanche`. It first sends the number of Discovery or Solicit messages specified by the `-R` option; then a retransmission (with an exponential back-off mechanism) is used for each simulated client, until all requests are answered. It generates a report when all clients receive their addresses, or when it is manually stopped. This scenario attempts to replicate a case where the server is not able to handle the traffic swiftly enough. Real clients will assume the packet or response was lost and will retransmit, further increasing DHCP traffic. This is sometimes called an *avalanche effect*, thus the scenario name. Option `-p` is ignored in the *avalanche* scenario.

When running a performance test, `perfdhcp` exchanges packets with the server under test as quickly as possible, unless the `-r` parameter is used to limit the request rate. The length of the test can be limited by setting a threshold on any or all of the number of requests made by `perfdhcp`, the elapsed time, or the number of requests dropped by the server.

25.10.3 Templates

To allow the contents of packets sent to the server to be customized, `perfdhcp` allows the specification of template files that determine the contents of the packets. For example, the customized packet may contain a DHCPv6 ORO to request a set of options to be returned by the server, or it may contain the Client FQDN option to request that the server perform DNS updates. This may be used to discover performance bottlenecks for different server configurations (e.g. DDNS enabled or disabled).

Up to two template files can be specified on the command line, with each file representing the contents of a particular type of packet, and the type being determined by the test being carried out. For example, if testing DHCPv6:

- With no template files specified on the command line, `perfdhcp` generates both Solicit and Request packets.
- With one template file specified, that file is used as the pattern for Solicit packets: `perfdhcp` generates the Request packets.
- With two template files given on the command line, the first is used as the pattern for Solicit packets, and the second as the pattern for Request packets.

(A similar determination applies to DHCPv4's DHCPDISCOVER and DHCPREQUEST packets.)

The template file holds the DHCP packet, represented as a stream of ASCII hexadecimal digits; it excludes any IP/UDP stack headers. The template file must not contain any characters other than hexadecimal digits and spaces. Spaces are discarded when the template file is parsed; in the file, `12B4` is the same as `12 B4`, which is the same as `1 2 B 4`.

The template files should be used in conjunction with the command-line parameters which specify offsets of the data fields being modified in outbound packets. For example, the `-E time-offset` switch specifies the offset of the DHCPv6 Elapsed Time option in the packet template. If the offset is specified, `perfdhcp` injects the current elapsed-time value into this field before sending the packet to the server.

In many scenarios, `perfdhcp` needs to simulate multiple clients, each having a unique client identifier. Since packets for each client are generated from the same template file, it is necessary to randomize the client identifier (or HW address in DHCPv4) in the packet created from it. The `-O random-offset` option allows specification of the offset in the template where randomization should be performed. It is important to note that this offset points to the end (not the beginning) of the client identifier (or HW address field). The number of bytes being randomized depends on the number of simulated clients. If the number of simulated clients is between 1 and 255, only one byte (to which the randomization offset points) is randomized. If the number of simulated clients is between 256 and 65535, two bytes are randomized. Note that the last two bytes of the client identifier are randomized in this case: the byte which the randomization offset parameter points to, and the one which precedes it (`random-offset - 1`). If the number of simulated clients exceeds 65535, three bytes are randomized, and so on.

`perfdhcp` can simulate traffic from multiple subnets by enabling option `-J` and passing a path to a file that contains v4 or v6 addresses to be used as relays in generated messages. That enables testing of vast numbers of Kea shared networks. While testing DHCPv4, Kea should be started with the `KEA_TEST_SEND_RESPONSES_TO_SOURCE` environment variable, to force Kea to send generated messages to the source address of the incoming packet.

Templates may currently be used to generate packets being sent to the server in 4-way exchanges, i.e. Solicit, Request (DHCPv6) and DHCPDISCOVER, DHCPREQUEST (DHCPv4). They cannot be used when Renew or DHCPRELEASE packets are being sent.

25.10.4 Options

- 1**
Takes the `server-id` option from the first received message.
- 4**
Establishes DHCPv4 operation; this is the default. It is incompatible with the `-6` option.
- 6**
Establishes DHCPv6 operation. It is incompatible with the `-4` option.
- b basetype=value**
Indicates the base MAC or DUID used to simulate different clients. The basetype may be "mac" or "duid". (The keyword "ether" may alternatively used for MAC.) The `-b` option can be specified multiple times. The MAC address must consist of six octets separated by single (:) or double (::) colons; for example: `mac=00:0c:01:02:03:04`. The DUID value is a hexadecimal string; it must be at least six octets long and not longer than 64 bytes, and the length must be less than 128 hexadecimal digits. For example: `duid=01010101010101010111F14`.
- d drop-time**
Specifies the time after which a request is treated as having been lost. The value is given in seconds and may contain a fractional component. The default is 1.
- e lease-type**
Specifies the type of lease being requested from the server. It may be one of the following:
 - address-only**
Only regular addresses (v4 or v6) are requested.
 - prefix-only**
Only IPv6 prefixes are requested.
 - address-and-prefix**
Both IPv6 addresses and prefixes are requested.The `-e prefix-only` and `-e address-and-prefix` forms may not be used with the `-4` option.
- F release-rate**
Specifies the rate at which DHCPv4 DHCPRELEASE or DHCPv6 Release requests are sent to a server. This value is only valid when used in conjunction with the exchange rate (given by `-r rate`). Furthermore, the sum of this value and the renew-rate (given by `-f rate`) must be equal to or less than the exchange rate value.
- f renew-rate**
Specifies the rate at which DHCPv4 DHCPREQUEST or DHCPv6 Renew requests are sent to a server. This value is only valid when used in conjunction with the exchange rate (given by `-r rate`). Furthermore, the sum of this value and the release-rate (given by `-F rate`) must be equal to or less than the exchange rate.
- g thread-mode**
Allows selection of thread-mode, which can be either `single` or `multi`. In multi-thread mode, packets are received in a separate thread, which allows better utilisation of CPUs. In a single-CPU system it is better to run in one thread, to avoid threads blocking each other. If more than one CPU is present in the system, multi-thread mode is the default; otherwise single-thread is the default.
- h**
Prints help and exits.

- i**
Performs only the initial part of the exchange: DISCOVER-OFFER if -4 is selected, Solicit-Advertise if -6 is chosen.
- i is incompatible with the following options: -1, -d, -D, -E, -S, -I and -F. In addition, it cannot be used with multiple instances of -O, -T, and -X.
- J remote-address-list-file**
Specifies a text file that includes multiple addresses, and is designed to test shared networks. If provided, `perfdhcp` randomly chooses one of the addresses for each exchange, to generate traffic from multiple subnets. When testing DHCPv4, it should be started with the `KEA_TEST_SEND_RESPONSES_TO_SOURCE=ENABLE` environment variable; otherwise, `perfdhcp` will not be able to receive responses.
- l local-addr|interface**
For DHCPv4 operation, specifies the local hostname/address to use when communicating with the server. By default, the interface address through which traffic would normally be routed to the server is used. For DHCPv6 operation, specifies the name of the network interface through which exchanges are initiated.
- L local-port**
Specifies the local port to use. This must be zero or a positive integer up to 65535. A value of 0 (the default) allows `perfdhcp` to choose its own port.
- M mac-list-file**
Specifies a text file containing a list of MAC addresses, one per line. If provided, a MAC address is chosen randomly from this list for every new exchange. In DHCPv6, MAC addresses are used to generate DUID-LLs. This parameter must not be used in conjunction with the -b parameter.
- N remote-port**
Specifies the remote port to use. This must be zero or a positive integer up to 65535. A value of 0 (the default) allows `perfdhcp` to choose the standard service port.
- o code,hexstring**
Forces `perfdhcp` to insert the specified extra option (or options if used multiple times) into packets being transmitted. The code specifies the option code, and the hexstring is a hexadecimal string that defines the content of the option. Care should be taken, as `perfdhcp` does not offer any kind of logic behind the options; they are simply inserted into packets and sent as is. Be careful not to duplicate options that are already inserted. For example, to insert the client-class identifier (option code 60) with the string "docsis", use -o 60,646f63736973. The -o may be used multiple times. The protocol family (-4 or -6) must be specified before using -o.
- P preload**
Initiates preload exchanges back-to-back at startup. Must be 0 (the default) or a positive integer.
- r rate**
Initiates the rate of DORA/SARR (or if -i is given, DO/SA) exchanges per second. A periodic report is generated showing the number of exchanges which were not completed, as well as the average response latency. The program continues until interrupted, at which point a final report is generated.
- R num-clients**
Specifies how many different clients are used. With a value of 1 (the default), all requests appear to come from the same client. Must be a positive number.
- s seed**
Specifies the seed for randomization, making runs of `perfdhcp` repeatable. This must be 0 or a positive integer. The value 0 means that a seed is not used; this is the default.
- scenario name**
Specifies the type of scenario, and can be `basic` (the default) or `avalanche`.
- T template-file**
Specifies a file containing the template to use as a stream of hexadecimal digits. This may be specified up to two

times and controls the contents of the packets sent (see the "Templates" section above).

-u

Enables checks for address uniqueness. The lease valid-lifetime should not be shorter than the test duration, and clients should not request an address more than once without releasing it.

-v

Displays the Kea version.

-V

Displays the extended Kea version.

-W exit-wait-time

Specifies the exit-wait-time parameter, which causes `perfdhcp` to wait for a certain amount of time after an exit condition has been met, to receive all packets without sending any new packets. Expressed in microseconds. If not specified, 0 is used (i.e. exit immediately after exit conditions are met).

-w script_name

Specifies the name of the script to be run before/after `perfdhcp`. When called, the script is passed a single parameter, either "start" or "stop", indicating whether it is being called before or after `perfdhcp`.

-x diagnostic-selector

Includes extended diagnostics in the output. This is a string of single keywords specifying the operations for which verbose output is desired. The selector key letters are:

a

Prints the decoded command-line arguments.

e

Prints the exit reason.

i

Prints the rate-processing details.

l

Prints the received leases.

s

Prints the first server ID.

t

When finished, prints timers of all successful exchanges.

T

When finished, prints templates.

-Y seconds

Time in seconds after which `perfdhcp` starts simulating the client waiting longer for server responses. This increases the `secs` field in DHCPv4 and sends increased values in the `Elapsed Time` option in DHCPv6. Must be used with `-y`.

-y seconds

Time in seconds during which `perfdhcp` simulates the client waiting longer for server responses. This increases the `secs` field in DHCPv4 and sends increased values in the `Elapsed Time` option in DHCPv6. Must be used with `-Y`.

25.10.5 DHCPv4-Only Options

The following options only apply for DHCPv4 (i.e. when `-4` is given).

-B

Forces broadcast handling.

25.10.6 DHCPv6-Only Options

The following options only apply for DHCPv6 (i.e. when `-6` is given).

-c

Adds a rapid-commit option (exchanges are Solicit-Advertise).

-A encapsulation-level

Specifies that relayed traffic must be generated. The argument specifies the level of encapsulation, i.e. how many relay agents are simulated. Currently the only supported encapsulation-level value is 1, which means that the generated traffic is equivalent to the amount of traffic passing through a single relay agent.

--or encapsulation-level:code,hexstring

Forces `perfdhcp` to insert the specified extra option (or options if used multiple times) into a relayed DHCPv6 message at the given level of encapsulation; currently, the only supported encapsulation-level value is 1. This option is very similar to `-o`: the code specifies the option code, and the hexstring is a hexadecimal string that defines the content of the option. Care should be taken, as `perfdhcp` does not offer any kind of logic behind the options; they are simply inserted into packets and sent as is. `encapsulation-level` is optional; if it is omitted, the default encapsulation-level value 1 is used. For example, to insert the Subscriber identifier (option code 38) with a value 1234 at the first level of encapsulation, use `--or 38,31323334` or `--or 1:38,31323334`. The `--or` may be used multiple times, and it must be used in conjunction with `-A`.

25.10.7 Template-Related Options

The following options may only be used in conjunction with `-T` and control how `perfdhcp` modifies the template. The options may be specified multiple times on the command line; each occurrence affects the corresponding template file (see "Templates" above).

-E time-offset

Specifies the offset of the secs field (DHCPv4) or Elapsed Time option (DHCPv6) in the second (i.e. Request) template; must be 0 or a positive integer. A value of 0 disables this.

-I ip-offset

Specifies the offset of the IP address (DHCPv4) in the requested-ip option or IA_NA option (DHCPv6) in the second (Request) template.

-O random-offset

Specifies the offset of the last octet to randomize in the template. This must be an integer greater than 3. The `-T` switch must be given to use this option.

-S srvid-offset

Specifies the offset of the server-id option in the second (Request) template. This must be a positive integer, and the switch can only be used when the template option (`-T`) is also given.

-X xid-offset

Specifies the offset of the transaction ID (xid) in the template. This must be a positive integer, and the switch can only be used when the template option (`-T`) is also given.

25.10.8 Options Controlling a Test

-D max-drop

Aborts the test immediately if "max-drop" requests have been dropped. Use `-D 1` to abort if even a single request has been dropped. "max-drop" must be a positive integer. If "max-drop" includes the suffix %, it specifies the maximum percentage of requests that may be dropped before aborting. In this case, testing of the threshold begins after 10 requests are expected to have been received.

-n num-requests

Initiates "num-request" transactions. No report is generated until all transactions have been initiated/waited-for, after which a report is generated and the program terminates.

-p test-period

Sends requests for "test-period", which is specified in the same manner as -d. This can be used as an alternative to -n, or both options can be given, in which case the testing is completed when either limit is reached.

-t interval

Sets the delay (in seconds) between two successive reports.

-C separator

Suppresses the preliminary output and causes the interim data to only contain the values delimited by separator. Used in conjunction with -t to produce easily parsable reports at -t intervals.

25.10.9 Arguments

server

Indicates the server to test, specified as an IP address. In the DHCPv6 case, the special name `all` can be used to refer to `All_DHCP_Relay_Agents_and_Servers` (the multicast address FF02::1:2), or the special name `servers` to refer to `All_DHCP_Servers` (the multicast address FF05::1:3). The server is mandatory except where the -l option is given to specify an interface, in which case it defaults to `all`.

25.10.10 Errors

perfdhcp can report the following errors in the packet exchange:

tooshort

A message was received that was too short.

orphans

A message was received which does not match one sent to the server (i.e. it is a duplicate message, a message that has arrived after an excessive delay, or one that is just not recognized).

locallimit

Local system limits have been reached when sending a message.

25.10.11 Exit Status

perfdhcp exits with one of the following status codes:

0

Success.

1

General error.

2

Error in command-line arguments.

3

No general failures in operation, but one or more exchanges were unsuccessful.

25.10.12 Usage Examples

Here is an example that simulates regular DHCPv4 traffic of 100 DHCPv4 devices (-R 100), 10 packets per second (-r 10), shows the query/response rate details (-xi), shows a report every 2 seconds (-t 2), and sends the packets to the IP 192.0.2.1:

```
sudo perfdhcp -xi -t 2 -r 10 -R 100 192.0.2.1
```

Here's a similar case, but for DHCPv6. Note that the DHCPv6 protocol uses link-local addresses, so the interface (eth0 in this example) must be specified on which to send the traffic. `all` is a convenience alias for

All_DHCP_Relay_Agents_and_Servers (the multicast address FF02::1:2). It is also possible to use the `servers` alias to refer to All_DHCP_Servers (the multicast address FF05::1:3). The default is `all`.

```
sudo perfdhcp -6 -xi -t 1 -r 1 -R 10 -l eth0 all
```

The following examples simulate normal DHCPv4 and DHCPv6 traffic that, after 3 seconds, starts pretending not to receive any responses from the server for 10 seconds. The DHCPv4 protocol signals this by an increased `secs` field, while DHCPv6 uses the `Elapsed Time` option. In real networks, this indicates that clients are not getting responses in a timely matter. This can be used to simulate some HA scenarios, as Kea uses the `secs` field and `Elapsed Time` option value as one of the indicators that the HA partner is not responding. When enabled with `-y` and `-Y`, the `secs` and `Elapsed Time` values increase steadily.

```
sudo perfdhcp -xi -t 1 -r 1 -y 10 -Y 3 192.0.2.1
```

```
sudo perfdhcp -6 -xi -t 1 -r 1 -y 10 -Y 3 2001:db8::1
```

25.10.13 Documentation

Kea comes with an extensive Kea Administrator Reference Manual that covers all aspects of running the Kea software - compilation, installation, configuration, configuration examples, and much more. Kea also features a Kea Messages Manual, which lists all possible messages Kea can print with a brief description for each of them. Both documents are available in various formats (.txt, .html, .pdf) with the Kea distribution. The Kea documentation is available at <https://kea.readthedocs.io>.

Kea source code is documented in the Kea Developer's Guide, available at https://reports.kea.isc.org/dev_guide/.

The Kea project website is available at <https://kea.isc.org>.

25.10.14 Mailing Lists and Support

There are two public mailing lists available for the Kea project. **kea-users** (kea-users at lists.isc.org) is intended for Kea users, while **kea-dev** (kea-dev at lists.isc.org) is intended for Kea developers, prospective contributors, and other advanced users. Both lists are available at <https://lists.isc.org>. The community provides best-effort support on both of those lists.

ISC provides professional support for Kea services. See <https://www.isc.org/kea/> for details.

25.10.15 History

The `perfdhcp` tool was initially coded in October 2011 by John DuBois, Francis Dupont, and Marcin Siodelski of ISC. Kea 1.0.0, which included `perfdhcp`, was released in December 2015.

25.10.16 See Also

kea-dhcp4(8), *kea-dhcp6(8)*, *kea-dhcp-ddns(8)*, *kea-ctrl-agent(8)*, *kea-admin(8)*, *kea-netconf(8)*, *keactrl(8)*, *kea-lfc(8)*, Kea Administrator Reference Manual.

KEA MESSAGES MANUAL

Kea is an open source implementation of the Dynamic Host Configuration Protocol (DHCP) servers, developed and maintained by Internet Systems Consortium (ISC).

This is the reference guide for Kea version 3.0.2. Links to the most up-to-date version of this document (in PDF, HTML, and plain text formats), along with other useful information about Kea, can be found in ISC's [Knowledgebase](#).

Please note that in the messages below, the percent sign (%) followed by a number is used to indicate a placeholder for data that is provided by the Kea code during its operation.

26.1 ALLOC

26.1.1 ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS

%1: ignoring globally reserved address %2, it falls outside %3

Logged at debug log level 40. This debug message is issued when the allocation engine determines that the globally reserved address falls outside the selected subnet or shared-network. The server should ignore the reserved address and attempt a dynamic allocation.

26.1.2 ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS6

%1: ignoring globally reserved address %2, it falls outside %3

Logged at debug log level 40. This debug message is issued when the allocation engine determines that the globally reserved address falls outside the selected subnet or shared-network. The server should ignore the reserved address and attempt a dynamic allocation.

26.1.3 ALLOC_ENGINE_LEASE_RECLAIMED

successfully reclaimed lease %1

Logged at debug log level 40. This debug message is logged when the allocation engine successfully reclaims a lease. The lease is now available for assignment.

26.1.4 ALLOC_ENGINE_V4_ALLOC_ERROR

%1: error during attempt to allocate an IPv4 address: %2

An error occurred during an attempt to allocate an IPv4 address, the reason for the failure being contained in the message. The server will return a message to the client refusing a lease. The first argument includes the client identification information.

26.1.5 ALLOC_ENGINE_V4_ALLOC_FAIL

```
%1: failed to allocate an IPv4 address after %2 attempt(s)
```

This is an old warning message issued when the allocation engine fails to allocate a lease for a client. This message includes a number of lease allocation attempts that the engine made before giving up. If the number of attempts is 0 because the engine was unable to use any of the address pools for the particular client, this message is not logged. Even though, several more detailed logs precede this message, it was left for backward compatibility. This message may indicate that your address pool is too small for the number of clients you are trying to service and should be expanded. Alternatively, if you know that the number of concurrently active clients is less than the addresses you have available, you may want to consider reducing the lease lifetime. This way, addresses allocated to clients that are no longer active on the network will become available sooner.

26.1.6 ALLOC_ENGINE_V4_ALLOC_FAIL_CLASSES

```
%1: Failed to allocate an IPv4 address for client with classes: %2
```

This warning message is printed when Kea failed to allocate an address and the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned. One of them may be a case when all pools require packet to belong to certain classes and the incoming packet didn't belong to any of them. Another case where this information may be useful is to point out that the pool reserved to a given class has ran out of addresses. When you see this message, you may consider checking your pool size and your classification definitions.

26.1.7 ALLOC_ENGINE_V4_ALLOC_FAIL_NO_POOLS

```
%1: no pools were available for the address allocation
```

This warning message is issued when the allocation engine fails to allocate a lease because it could not use any configured pools for the particular client. It is also possible that all of the subnets from which the allocation engine attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects that some clients should be assigned dynamic addresses. A subnet may lack any pools only when all clients should be assigned reserved IP addresses. Suppose the subnets connected to a shared network or a single subnet to which the client belongs have pools configured. In that case, this message is an indication that none of the pools could be used for the client because the client does not belong to appropriate client classes.

26.1.8 ALLOC_ENGINE_V4_ALLOC_FAIL_SHARED_NETWORK

```
%1: failed to allocate an IPv4 address in the shared network %2: %3 subnets have no
↪available addresses, %4 subnets have no matching pools
```

This warning message is issued when the allocation engine fails to allocate a lease for a client connected to a shared network. The shared network should contain at least one subnet, but typically it aggregates multiple subnets. This log message indicates that the allocation engine could not find and allocate any suitable lease in any of the subnets within the shared network. The first argument includes the client identification information. The second argument specifies the shared network name. The remaining two arguments provide additional information useful for debugging why the allocation engine could not assign a lease. The allocation engine tries to allocate addresses from different subnets in the shared network, and it may fail for some subnets because there are no leases available in those subnets or the free leases are reserved to other clients. The number of such subnets is specified in the third argument. For other subnets the allocation may fail because their pools may not be available to the particular client. These pools are guarded by client classes that the client does not belong to. The fourth argument specifies the number of such subnets. By looking at the values in the third and fourth argument, an operator can identify the situations when there are no addresses left in some of the pools. He or she can also identify a client classification misconfigurations causing some clients to be refused the service.

26.1.9 ALLOC_ENGINE_V4_ALLOC_FAIL_SUBNET

```
%1: failed to allocate an IPv4 lease in the subnet %2, subnet-id %3, shared network %4
```

This warning message is issued when the allocation engine fails to allocate a lease for a client connected to a subnet. The first argument includes the client identification information. The second and third arguments identify the subnet. The fourth argument specifies the shared network, if the subnet belongs to a shared network. There are many reasons for failing lease allocations. One of them may be the pools exhaustion or existing reservations for the free leases. However, in some cases, the allocation engine may fail to find a suitable pool for the client when the pools are only available to certain client classes, but the requesting client does not belong to them. Further log messages provide more information to distinguish between these different cases.

26.1.10 ALLOC_ENGINE_V4_DECLINED_RECOVERED

```
IPv4 address %1 was recovered after %2 seconds of probation-period
```

This informational message indicates that the specified address was reported as duplicate (client sent DECLINE) and the server marked this address as unavailable for a period of time. This time now has elapsed and the address has been returned to the available pool. This step concludes the decline recovery process.

26.1.11 ALLOC_ENGINE_V4_DISCOVER_ADDRESS_CONFLICT

```
%1: conflicting reservation for address %2 with existing lease %3
```

This warning message is issued when the DHCP server finds that the address reserved for the client can't be offered because this address is currently allocated to another client. The server will try to allocate a different address to the client to use until the conflict is resolved. The first argument includes the client identification information.

26.1.12 ALLOC_ENGINE_V4_DISCOVER_HR

```
client %1 sending DHCPDISCOVER has reservation for the address %2
```

Logged at debug log level 40. This message is issued when the allocation engine determines that the client sending the DHCPDISCOVER has a reservation for the specified address. The allocation engine will try to offer this address to the client.

26.1.13 ALLOC_ENGINE_V4_LEASES_RECLAMATION_COMPLETE

```
reclaimed %1 leases in %2
```

Logged at debug log level 40. This debug message is logged when the allocation engine completes reclamation of a set of expired leases. The maximum number of leases to be reclaimed in a single pass of the lease reclamation routine is configurable using 'max-reclaim-leases' parameter. However, the number of reclaimed leases may also be limited by the timeout value, configured with 'max-reclaim-time'. The message includes the number of reclaimed leases and the total time.

26.1.14 ALLOC_ENGINE_V4_LEASES_RECLAMATION_FAILED

```
reclamation of expired leases failed: %1
```

This error message is issued when the reclamation of the expired leases failed. The error message is displayed.

26.1.15 ALLOC_ENGINE_V4_LEASES_RECLAMATION_SLOW

expired leases still exist after %1 reclamations

This warning message is issued when the server has been unable to reclaim all expired leases in a specified number of consecutive attempts. This indicates that the value of "reclaim-timer-wait-time" may be too high. However, if this is just a short burst of leases' expirations the value does not have to be modified and the server should deal with this in subsequent reclamation attempts. If this is a result of a permanent increase of the server load, the value of "reclaim-timer-wait-time" should be decreased, or the values of "max-reclaim-leases" and "max-reclaim-time" should be increased to allow processing more leases in a single cycle. Alternatively, these values may be set to 0 to remove the limitations on the number of leases and duration. However, this may result in longer periods of server's unresponsiveness to DHCP packets, while it processes the expired leases.

26.1.16 ALLOC_ENGINE_V4_LEASES_RECLAMATION_START

starting reclamation of expired leases (limit = %1 leases or %2 milliseconds)

Logged at debug log level 40. This debug message is issued when the allocation engine starts the reclamation of the expired leases. The maximum number of leases to be reclaimed and the timeout is included in the message. If any of these values is 0, it means "unlimited".

26.1.17 ALLOC_ENGINE_V4_LEASES_RECLAMATION_TIMEOUT

timeout of %1 ms reached while reclaiming IPv4 leases

Logged at debug log level 40. This debug message is issued when the allocation engine hits the timeout for performing reclamation of the expired leases. The reclamation will now be interrupted and all leases which haven't been reclaimed, because of the timeout, will be reclaimed when the next scheduled reclamation is started. The argument is the timeout value expressed in milliseconds.

26.1.18 ALLOC_ENGINE_V4_LEASE_RECLAIM

%1: reclaiming expired lease for address %2

Logged at debug log level 40. This debug message is issued when the server begins reclamation of the expired DHCPv4 lease. The first argument specifies the client identification information. The second argument holds the leased IPv4 address.

26.1.19 ALLOC_ENGINE_V4_LEASE_RECLAMATION_FAILED

failed to reclaim the lease %1: %2

This error message is logged when the allocation engine fails to reclaim an expired lease. The reason for the failure is included in the message. The error may be triggered in the lease expiration hook or while performing the operation on the lease database.

26.1.20 ALLOC_ENGINE_V4_NO_MORE_EXPIRED_LEASES

all expired leases have been reclaimed

Logged at debug log level 40. This debug message is issued when the server reclaims all expired DHCPv4 leases in the database.

26.1.21 ALLOC_ENGINE_V4_OFFER_EXISTING_LEASE

allocation engine will try to offer existing lease to the client %1

Logged at debug log level 40. This message is issued when the allocation engine determines that the client has a lease in the lease database, it doesn't have reservation for any other lease, and the leased address is not reserved for any other client. The allocation engine will try to offer the same lease to the client.

26.1.22 ALLOC_ENGINE_V4_OFFER_NEW_LEASE

allocation engine will try to offer new lease to the client %1

Logged at debug log level 40. This message is issued when the allocation engine will try to offer a new lease to the client. This is the case when the client doesn't have any existing lease, it has no reservation or the existing or reserved address is leased to another client. Also, the client didn't specify a hint, or the address in the hint is in use.

26.1.23 ALLOC_ENGINE_V4_OFFER_REQUESTED_LEASE

allocation engine will try to offer requested lease %1 to the client %2

Logged at debug log level 40. This message is issued when the allocation engine will try to offer the lease specified in the hint. This situation may occur when: (a) client doesn't have any reservations, (b) client has reservation but the reserved address is leased to another client.

26.1.24 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE

begin deletion of reclaimed leases expired more than %1 seconds ago

Logged at debug log level 40. This debug message is issued when the allocation engine begins deletion of the reclaimed leases which have expired more than a specified number of seconds ago. This operation is triggered periodically according to the "flush-reclaimed-timer-wait-time" parameter. The "hold-reclaimed-time" parameter defines a number of seconds for which the leases are stored before they are removed.

26.1.25 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE_COMPLETE

successfully deleted %1 expired-reclaimed leases

Logged at debug log level 40. This debug message is issued when the server successfully deletes "expired-reclaimed" leases from the lease database. The number of deleted leases is included in the log message.

26.1.26 ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE_FAILED

deletion of expired-reclaimed leases failed: %1

This error message is issued when the deletion of "expired-reclaimed" leases from the database failed. The error message is appended to the log message.

26.1.27 ALLOC_ENGINE_V4_REQUEST_ADDRESS_RESERVED

%1: requested address %2 is reserved

Logged at debug log level 40. This message is issued when the allocation engine refused to allocate address requested by the client because this address is reserved for another client. The first argument includes the client identification information.

26.1.28 ALLOC_ENGINE_V4_REQUEST_ALLOC_REQUESTED

```
%1: trying to allocate requested address %2
```

Logged at debug log level 40. This message is issued when the allocation engine is trying to allocate (or reuse an expired) address which has been requested by the client. The first argument includes the client identification information.

26.1.29 ALLOC_ENGINE_V4_REQUEST_EXTEND_LEASE

```
%1: extending lifetime of the lease for address %2
```

Logged at debug log level 40. This message is issued when the allocation engine determines that the client already has a lease whose lifetime can be extended, and which can be returned to the client. The first argument includes the client identification information.

26.1.30 ALLOC_ENGINE_V4_REQUEST_INVALID

```
client %1 having a reservation for address %2 is requesting invalid address %3
```

Logged at debug log level 40. This message is logged when the client, having a reservation for one address, is requesting a different address. The client is only allowed to do this when the reserved address is in use by another client. However, the allocation engine has determined that the reserved address is available and the client should request the reserved address.

26.1.31 ALLOC_ENGINE_V4_REQUEST_IN_USE

```
%1: requested address %2 is in use
```

Logged at debug log level 40. This message is issued when the client is requesting or has a reservation for an address which is in use. The first argument includes the client identification information.

26.1.32 ALLOC_ENGINE_V4_REQUEST_OUT_OF_POOL

```
client %1, which doesn't have a reservation, requested address %2 out of the dynamic pool
```

Logged at debug log level 40. This message is issued when the client has requested allocation of the address which doesn't belong to any address pool from which addresses are dynamically allocated. The client also doesn't have reservation for this address. This address could only be allocated if the client had reservation for it.

26.1.33 ALLOC_ENGINE_V4_REQUEST_PICK_ADDRESS

```
client %1 hasn't specified an address - picking available address from the pool
```

Logged at debug log level 40. This message is logged when the client hasn't specified any preferred address (the client should always do it, but Kea tries to be forgiving). The allocation engine will try to pick an available address from the dynamic pool and allocate it to the client.

26.1.34 ALLOC_ENGINE_V4_REQUEST_REMOVE_LEASE

```
%1: removing previous client's lease %2
```

Logged at debug log level 40. This message is logged when the allocation engine removes previous lease for the client because the client has been allocated new one.

26.1.35 ALLOC_ENGINE_V4_REQUEST_USE_HR

```
client %1 hasn't requested specific address, using reserved address %2
```

Logged at debug log level 40. This message is issued when the client is not requesting any specific address but the allocation engine has determined that there is a reservation for this client. The allocation engine will try to allocate the reserved address.

26.1.36 ALLOC_ENGINE_V4_REUSE_EXPIRED_LEASE_DATA

```
%1: reusing expired lease, updated lease information: %2
```

Logged at debug log level 55. This message is logged when the allocation engine is reusing an existing lease. The details of the updated lease are printed. The first argument includes the client identification information.

26.1.37 ALLOC_ENGINE_V6_ALLOC_ERROR

```
%1: error during attempt to allocate an IPv6 address: %2
```

An error occurred during an attempt to allocate an IPv6 address, the reason for the failure being contained in the message. The server will return a message to the client refusing a lease. The first argument includes the client identification information.

26.1.38 ALLOC_ENGINE_V6_ALLOC_FAIL

```
%1: failed to allocate an IPv6 lease after %2 attempt(s)
```

This is an old warning message issued when the allocation engine fails to allocate a lease for a client. This message includes a number of lease allocation attempts that the engine made before giving up. If the number of attempts is 0 because the engine was unable to use any of the pools for the particular client, this message is not logged. Even though, several more detailed logs precede this message, it was left for backward compatibility. This message may indicate that your pool is too small for the number of clients you are trying to service and should be expanded. Alternatively, if you know that the number of concurrently active clients is less than the leases you have available, you may want to consider reducing the lease lifetime. This way, leases allocated to clients that are no longer active on the network will become available sooner.

26.1.39 ALLOC_ENGINE_V6_ALLOC_FAIL_CLASSES

```
%1: Failed to allocate an IPv6 address for client with classes: %2
```

This warning message is printed when Kea failed to allocate an address and the client's packet belongs to one or more classes. There may be several reasons why a lease was not assigned. One of them may be a case when all pools require packet to belong to certain classes and the incoming packet didn't belong to any of them. Another case where this information may be useful is to point out that the pool reserved to a given class has ran out of addresses. When you see this message, you may consider checking your pool size and your classification definitions.

26.1.40 ALLOC_ENGINE_V6_ALLOC_FAIL_NO_POOLS

```
%1: no pools were available for the lease allocation
```

This warning message is issued when the allocation engine fails to allocate a lease because it could not use any configured pools for the particular client. It is also possible that all of the subnets from which the allocation engine attempted to assign an address lack address pools. In this case, it should be considered misconfiguration if an operator expects

that some clients should be assigned dynamic addresses. A subnet may lack any pools only when all clients should be assigned reserved leases. Suppose the subnets connected to a shared network or a single subnet to which the client belongs have pools configured. In that case, this message is an indication that none of the pools could be used for the client because the client does not belong to appropriate client classes.

26.1.41 ALLOC_ENGINE_V6_ALLOC_FAIL_SHARED_NETWORK

```
%1: failed to allocate a lease in the shared network %2: %3 subnets have no available_
↳leases, %4 subnets have no matching pools
```

This warning message is issued when the allocation engine fails to allocate a lease for a client connected to a shared network. The shared network should contain at least one subnet, but typically it aggregates multiple subnets. This log message indicates that the allocation engine could not find and allocate any suitable lease in any of the subnets within the shared network. The first argument includes the client identification information. The second argument specifies the shared network name. The remaining two arguments provide additional information useful for debugging why the allocation engine could not assign a lease. The allocation engine tries to allocate leases from different subnets in the shared network, and it may fail for some subnets because there are no leases available in those subnets or the free leases are reserved to other clients. The number of such subnets is specified in the third argument. For other subnets the allocation may fail because their pools may not be available to the particular client. These pools are guarded by client classes that the client does not belong to. The fourth argument specifies the number of such subnets. By looking at the values in the third and fourth argument, an operator can identify the situations when there are no leases left in some of the pools. He or she can also identify client classification misconfigurations causing some clients to be refused the service.

26.1.42 ALLOC_ENGINE_V6_ALLOC_FAIL_SUBNET

```
%1: failed to allocate an IPv6 lease in the subnet %2, subnet-id %3, shared network %4
```

This warning message is issued when the allocation engine fails to allocate a lease for a client connected to a subnet. The first argument includes the client identification information. The second and third arguments identify the subnet. The fourth argument specifies the shared network, if the subnet belongs to a shared network. There are many reasons for failing lease allocations. One of them may be the pools exhaustion or existing reservations for the free leases. However, in some cases, the allocation engine may fail to find a suitable pool for the client when the pools are only available to certain client classes, but the requesting client does not belong to them. Further log messages provide more information to distinguish between these different cases.

26.1.43 ALLOC_ENGINE_V6_ALLOC_HR_LEASE_EXISTS

```
%1: lease type %2 for reserved address/prefix %3 already exists
```

Logged at debug log level 40. This debug message is issued when the allocation engine determines that the lease for the IPv6 address or prefix has already been allocated for the client and the client can continue using it. The first argument includes the client identification information.

26.1.44 ALLOC_ENGINE_V6_ALLOC_LEASES_HR

```
leases and static reservations found for client %1
```

Logged at debug log level 40. This message is logged when the allocation engine is in the process of allocating leases for the client, it found existing leases and static reservations for the client. The allocation engine will verify if existing leases match reservations. Those leases that are reserved for other clients and those that are not reserved for the client will be removed. All leases matching the reservations will be renewed and returned.

26.1.45 ALLOC_ENGINE_V6_ALLOC_LEASES_NO_HR

no reservations found but leases exist for client %1

Logged at debug log level 40. This message is logged when the allocation engine is in the process of allocating leases for the client, there are no static reservations, but lease(s) exist for the client. The allocation engine will remove leases which are reserved for other clients, and return all remaining leases to the client.

26.1.46 ALLOC_ENGINE_V6_ALLOC_NO_LEASES_HR

no leases found but reservations exist for client %1

Logged at debug log level 40. This message is logged when the allocation engine is in the process of allocating leases for the client. It hasn't found any existing leases for this client, but the client appears to have static reservations. The allocation engine will try to allocate the reserved resources for the client.

26.1.47 ALLOC_ENGINE_V6_ALLOC_NO_V6_HR

%1: unable to allocate reserved leases - no IPv6 reservations

Logged at debug log level 40. This message is logged when the allocation engine determines that the client has no IPv6 reservations and thus the allocation engine will have to try to allocate allocating leases from the dynamic pool or stop the allocation process if none can be allocated. The first argument includes the client identification information.

26.1.48 ALLOC_ENGINE_V6_ALLOC_UNRESERVED

no static reservations available - trying to dynamically allocate leases for client %1

Logged at debug log level 40. This debug message is issued when the allocation engine will attempt to allocate leases from the dynamic pools. This may be due to one of (a) there are no reservations for this client, (b) there are reservations for the client but they are not usable because the addresses are in use by another client or (c) we had a reserved lease but that has now been allocated to another client.

26.1.49 ALLOC_ENGINE_V6_CALCULATED_PREFERRED_LIFETIME

%1: using a calculated preferred-lifetime of %2

Logged at debug log level 40. This debug message indicates that the preferred-lifetime being returned to the client is defaulting to 62.5% of the valid-lifetime. This may occur if either the preferred-lifetime has not been explicitly configured, or the configured value is larger than the valid-lifetime. The arguments detail the client and the preferred-lifetime that will be used.

26.1.50 ALLOC_ENGINE_V6_DECLINED_RECOVERED

IPv6 address %1 was recovered after %2 seconds of probation-period

This informational message indicates that the specified address was reported as duplicate (client sent DECLINE) and the server marked this address as unavailable for a period of time. This time now has elapsed and the address has been returned to the available pool. This step concludes the decline recovery process.

26.1.51 ALLOC_ENGINE_V6_EXPIRED_HINT_RESERVED

%1: expired lease for the client's hint %2 is reserved for another client

Logged at debug log level 40. This message is logged when the allocation engine finds that the expired lease for the client's hint can't be reused because it is reserved for another client. The first argument includes the client identification information.

26.1.52 ALLOC_ENGINE_V6_EXTEND_ALLOC_UNRESERVED

allocate new (unreserved) leases for the renewing client %1

Logged at debug log level 40. This debug message is issued when the allocation engine is trying to allocate new leases for the renewing client because it was unable to renew any of the existing client's leases, e.g. because leases are reserved for another client or for any other reason.

26.1.53 ALLOC_ENGINE_V6_EXTEND_ERROR

%1: allocation engine experienced error with attempting to extend lease lifetime: %2

This error message indicates that an error was experienced during Renew or Rebind processing. Additional explanation is provided with this message. Depending on its nature, manual intervention may be required to continue processing messages from this particular client; other clients will be unaffected. The first argument includes the client identification information.

26.1.54 ALLOC_ENGINE_V6_EXTEND_LEASE

%1: extending lifetime of the lease type %2, address %3

Logged at debug log level 50. This debug message is issued when the allocation engine is trying to extend lifetime of the lease. The first argument includes the client identification information.

26.1.55 ALLOC_ENGINE_V6_EXTEND_LEASE_DATA

%1: detailed information about the lease being extended: %2

Logged at debug log level 55. This debug message prints detailed information about the lease which lifetime is being extended (renew or rebind). The first argument includes the client identification information.

26.1.56 ALLOC_ENGINE_V6_EXTEND_NEW_LEASE_DATA

%1: new lease information for the lease being extended: %2

Logged at debug log level 55. This debug message prints updated information about the lease to be extended. If the lease update is successful, the information printed by this message will be stored in the database. The first argument includes the client identification information.

26.1.57 ALLOC_ENGINE_V6_HINT_RESERVED

%1: lease for the client's hint %2 is reserved for another client

Logged at debug log level 40. This message is logged when the allocation engine cannot allocate the lease using the client's hint because the lease for this hint is reserved for another client. The first argument includes the client identification information.

26.1.58 ALLOC_ENGINE_V6_HR_ADDR_GRANTED

```
reserved address %1 was assigned to client %2
```

This informational message signals that the specified client was assigned the address reserved for it.

26.1.59 ALLOC_ENGINE_V6_HR_PREFIX_GRANTED

```
reserved prefix %1/%2 was assigned to client %3
```

This informational message signals that the specified client was assigned the prefix reserved for it.

26.1.60 ALLOC_ENGINE_V6_LEASES_RECLAMATION_COMPLETE

```
reclaimed %1 leases in %2
```

Logged at debug log level 40. This debug message is logged when the allocation engine completes reclamation of a set of expired leases. The maximum number of leases to be reclaimed in a single pass of the lease reclamation routine is configurable using 'max-reclaim-leases' parameter. However, the number of reclaimed leases may also be limited by the timeout value, configured with 'max-reclaim-time'. The message includes the number of reclaimed leases and the total time.

26.1.61 ALLOC_ENGINE_V6_LEASES_RECLAMATION_FAILED

```
reclamation of expired leases failed: %1
```

This error message is issued when the reclamation of the expired leases failed. The error message is displayed.

26.1.62 ALLOC_ENGINE_V6_LEASES_RECLAMATION_SLOW

```
expired leases still exist after %1 reclamations
```

This warning message is issued when the server has been unable to reclaim all expired leases in a specified number of consecutive attempts. This indicates that the value of "reclaim-timer-wait-time" may be too high. However, if this is just a short burst of leases' expirations the value does not have to be modified and the server should deal with this in subsequent reclamation attempts. If this is a result of a permanent increase of the server load, the value of "reclaim-timer-wait-time" should be decreased, or the values of "max-reclaim-leases" and "max-reclaim-time" should be increased to allow processing more leases in a single cycle. Alternatively, these values may be set to 0 to remove the limitations on the number of leases and duration. However, this may result in longer periods of server's unresponsiveness to DHCP packets, while it processes the expired leases.

26.1.63 ALLOC_ENGINE_V6_LEASES_RECLAMATION_START

```
starting reclamation of expired leases (limit = %1 leases or %2 milliseconds)
```

Logged at debug log level 40. This debug message is issued when the allocation engine starts the reclamation of the expired leases. The maximum number of leases to be reclaimed and the timeout is included in the message. If any of these values is 0, it means "unlimited".

26.1.64 ALLOC_ENGINE_V6_LEASES_RECLAMATION_TIMEOUT

```
timeout of %1 ms reached while reclaiming IPv6 leases
```

Logged at debug log level 40. This debug message is issued when the allocation engine hits the timeout for performing reclamation of the expired leases. The reclamation will now be interrupted and all leases which haven't been reclaimed, because of the timeout, will be reclaimed when the next scheduled reclamation is started. The argument is the timeout value expressed in milliseconds.

26.1.65 ALLOC_ENGINE_V6_LEASE_RECLAIM

```
%1: reclaiming expired lease for prefix %2/%3
```

Logged at debug log level 40. This debug message is issued when the server begins reclamation of the expired DHCPv6 lease. The reclaimed lease may either be an address lease or delegated prefix. The first argument provides the client identification information. The other arguments specify the prefix and the prefix length for the lease. The prefix length for address lease is equal to 128.

26.1.66 ALLOC_ENGINE_V6_LEASE_RECLAMATION_FAILED

```
failed to reclaim the lease %1: %2
```

This error message is logged when the allocation engine fails to reclaim an expired lease. The reason for the failure is included in the message. The error may be triggered in the lease expiration hook or while performing the operation on the lease database.

26.1.67 ALLOC_ENGINE_V6_NO_MORE_EXPIRED_LEASES

```
all expired leases have been reclaimed
```

Logged at debug log level 40. This debug message is issued when the server reclaims all expired DHCPv6 leases in the database.

26.1.68 ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE

```
begin deletion of reclaimed leases expired more than %1 seconds ago
```

Logged at debug log level 40. This debug message is issued when the allocation engine begins deletion of the reclaimed leases which have expired more than a specified number of seconds ago. This operation is triggered periodically according to the "flush-reclaimed-timer-wait-time" parameter. The "hold-reclaimed-time" parameter defines a number of seconds for which the leases are stored before they are removed.

26.1.69 ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE_COMPLETE

```
successfully deleted %1 expired-reclaimed leases
```

Logged at debug log level 40. This debug message is issued when the server successfully deletes "expired-reclaimed" leases from the lease database. The number of deleted leases is included in the log message.

26.1.70 ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE_FAILED

deletion of expired-reclaimed leases failed: %1

This error message is issued when the deletion of "expired-reclaimed" leases from the database failed. The error message is appended to the log message.

26.1.71 ALLOC_ENGINE_V6_RENEW_HR

allocating leases reserved for the client %1 as a result of Renew

Logged at debug log level 40. This debug message is issued when the allocation engine tries to allocate reserved leases for the client sending a Renew message. The server will also remove any leases that the client is trying to renew that are not reserved for the client.

26.1.72 ALLOC_ENGINE_V6_RENEW_REMOVE_RESERVED

%1: checking if existing client's leases are reserved for another client

Logged at debug log level 40. This message is logged when the allocation engine finds leases for the client and will check if these leases are reserved for another client. If they are, they will not be renewed for the client requesting their renewal. The first argument includes the client identification information.

26.1.73 ALLOC_ENGINE_V6_REUSE_EXPIRED_LEASE_DATA

%1: reusing expired lease, updated lease information: %2

Logged at debug log level 55. This message is logged when the allocation engine is reusing an existing lease. The details of the updated lease are printed. The first argument includes the client identification information.

26.1.74 ALLOC_ENGINE_V6_REVOKED_ADDR_LEASE

%1: address %2 was revoked from client %3 as it is reserved for client %4

This informational message is an indication that the specified IPv6 address was used by client A but it is now reserved for client B. Client A has been told to stop using it so that it can be leased to client B. This is a normal occurrence during conflict resolution, which can occur in cases such as the system administrator adding a reservation for an address that is currently in use by another client. The server will fully recover from this situation, but clients will change their addresses.

26.1.75 ALLOC_ENGINE_V6_REVOKED_PREFIX_LEASE

%1: prefix %2/%3 was revoked from client %4 as it is reserved for client %5

This informational message is an indication that the specified IPv6 prefix was used by client A but it is now reserved for client B. Client A has been told to stop using it so that it can be leased to client B. This is a normal occurrence during conflict resolution, which can occur in cases such as the system administrator adding a reservation for an address that is currently in use by another client. The server will fully recover from this situation, but clients will change their prefixes.

26.1.76 ALLOC_ENGINE_V6_REVOKED_SHARED_ADDR_LEASE

```
%1: address %2 was revoked from client %3 as it is reserved for %4 other clients
```

This informational message is an indication that the specified IPv6 address was used by client A but it is now reserved for multiple other clients. Client A has been told to stop using it so that it can be leased to one of the clients having the reservation for it. This is a normal occurrence during conflict resolution, which can occur in cases such as the system administrator adding reservations for an address that is currently in use by another client. The server will fully recover from this situation, but clients will change their addresses.

26.1.77 ALLOC_ENGINE_V6_REVOKED_SHARED_PREFIX_LEASE

```
%1: prefix %2/%3 was revoked from client %4 as it is reserved for %5 other clients
```

This informational message is an indication that the specified IPv6 prefix was used by client A but it is now reserved for multiple other clients. Client A has been told to stop using it so that it can be leased to one of the clients having the reservation for it. This is a normal occurrence during conflict resolution, which can occur in cases such as the system administrator adding reservations for an address that is currently in use by another client. The server will fully recover from this situation, but clients will change their prefixes.

26.2 ASIODNS

26.2.1 ASIODNS_FETCH_COMPLETED

```
upstream fetch to %1(%2) has now completed
```

Logged at debug log level 70. A debug message, this records that the upstream fetch (a query made by the resolver on behalf of its client) to the specified address has completed.

26.2.2 ASIODNS_FETCH_STOPPED

```
upstream fetch to %1(%2) has been stopped
```

Logged at debug log level 40. An external component has requested the halting of an upstream fetch. This is an allowed operation, and the message should only appear if debug is enabled.

26.2.3 ASIODNS_OPEN_SOCKET

```
error %1 opening %2 socket to %3(%4)
```

The asynchronous I/O code encountered an error when trying to open a socket of the specified protocol in order to send a message to the target address. The number of the system error that caused the problem is given in the message.

26.2.4 ASIODNS_READ_DATA

```
error %1 reading %2 data from %3(%4)
```

The asynchronous I/O code encountered an error when trying to read data from the specified address on the given protocol. The number of the system error that caused the problem is given in the message.

26.2.5 ASIODNS_READ_TIMEOUT

```
receive timeout while waiting for data from %1(%2)
```

Logged at debug log level 50. An upstream fetch from the specified address timed out. This may happen for any number of reasons and is most probably a problem at the remote server or a problem on the network. The message will only appear if debug is enabled.

26.2.6 ASIODNS_SEND_DATA

```
error %1 sending data using %2 to %3(%4)
```

The asynchronous I/O code encountered an error when trying to send data to the specified address on the given protocol. The number of the system error that caused the problem is given in the message.

26.2.7 ASIODNS_UNKNOWN_ORIGIN

```
unknown origin for ASIO error code %1 (protocol: %2, address %3)
```

An internal consistency check on the origin of a message from the asynchronous I/O module failed. This may indicate an internal error; please submit a bug report.

26.2.8 ASIODNS_UNKNOWN_RESULT

```
unknown result (%1) when IOFetch::stop() was executed for I/O to %2(%3)
```

An internal error indicating that the termination method of the resolver's upstream fetch class was called with an unknown result code (which is given in the message). Please submit a bug report.

26.3 BAD

26.3.1 BAD_CLIENT_CREDENTIALS

```
bad client credentials: %1
```

This error message is issued when the client credential processing failed, including when the credential remaining lifetime is shorter than the TKEY lifetime. The argument details the error.

26.4 BOOTP

26.4.1 BOOTP_BOOTP_QUERY

```
recognized a BOOTP query: %1
```

Logged at debug log level 40. This debug message is printed when the BOOTP query was recognized. The BOOTP client class was added and the message type set to DHCPREQUEST. The query client and transaction identification are displayed.

26.4.2 BOOTP_LOAD

Bootp hooks library has been loaded

This info message indicates that the Bootp hooks library has been loaded.

26.4.3 BOOTP_PACKET_OPTIONS_SKIPPED

an error unpacking an option, caused subsequent options to be skipped: %1

Logged at debug log level 40. A debug message issued when an option failed to unpack correctly, making it impossible to unpack the remaining options in the DHCPv4 query. The server will still attempt to service the packet. The sole argument provides a reason for unpacking error.

26.4.4 BOOTP_PACKET_PACK

%1: preparing on-wire format of the packet to be sent

Logged at debug log level 40. This debug message is issued when the server starts preparing the on-wire format of the packet to be sent back to the client. The argument specifies the client and the transaction identification information.

26.4.5 BOOTP_PACKET_PACK_FAIL

%1: preparing on-wire-format of the packet to be sent failed %2

This error message is issued when preparing an on-wire format of the packet has failed. The first argument identifies the client and the BOOTP transaction. The second argument includes the error string.

26.4.6 BOOTP_PACKET_UNPACK_FAILED

failed to parse query from %1 to %2, received over interface %3, reason: %4

Logged at debug log level 40. This debug message is issued when received DHCPv4 query is malformed and can't be parsed by the buffer4_receive callout. The query will be dropped by the server. The first three arguments specify source IP address, destination IP address and the interface. The last argument provides a reason for failure.

26.4.7 BOOTP_UNLOAD

Bootp hooks library has been unloaded

This info message indicates that the Bootp hooks library has been unloaded.

26.5 BULK

26.5.1 BULK_LEASE_QUERY4_UNSUPPORTED_MSG_TYPE

Dropping packet with an unsupported DHCPv4 message type %1 received from: %2

This error message is issued when a DHCPv4 packet type that the lease query hook does not support has been received. The first argument is the unsupported message type, the second the remote address of the connection which will be closed.

26.5.2 BULK_LEASE_QUERY6_UNSUPPORTED_MSG_TYPE

Dropping packet with an unsupported DHCPv6 message type %1 received from: %2

This error message is issued when a DHCPv6 packet type that the lease query hook does not support has been received. The first argument is the unsupported message type, the second the remote address of the connection which will be closed.

26.5.3 BULK_LEASE_QUERY_AT_MAX_CONCURRENT_QUERIES

Queuing query from: %1, details: %2, connection already has %3 queries in progress

Logged at debug log level 40. This debug message is issued when a requester sends a bulk lease query on a connection that already has the maximum number of queries allowed in progress. The first argument is the requester's address, the second details the query which has been queued, and the third is the value of max-concurrent-queries.

26.5.4 BULK_LEASE_QUERY_DEQUEUED

A query from %1, details: %2, dequeued.

Logged at debug log level 40. This debug message is issued when a query has been dequeued and will be processed. The first argument is the requester's address, the second details the query which has been dequeued.

26.5.5 BULK_LEASE_QUERY_DUPLICATE_XID

Dropping query from: %1, transaction id %2 is a duplicate

This warning message is issued when a requester sends a bulk lease query with the same transaction id while that requester already has a query with the same transaction id in-progress. The first argument is the requester's address, the second the duplicated transaction id.

26.5.6 BULK_LEASE_QUERY_EMPTY_REQUEST

A bulk lease query packet received from %1 is empty.

This error message is issued when received bulk lease query packet with no payload. The argument contains the remote address of the connection which will be closed.

26.5.7 BULK_LEASE_QUERY_INVALID_REQUEST

A bulk lease query packet received from %1 is invalid, query: %2, error: %3

This error message is issued when received invalid bulk lease query packet. The first argument is the remote address of the connection which will be closed, the second is the query, the last one is the error message.

26.5.8 BULK_LEASE_QUERY_LISTENER_START_FAILED

Bulk lease query listener thread pool could not be started, error %1

This error message is emitted when the bulk lease query listener's could (re)started following a reconfiguration event. This most likely cause would be a runtime configuration error, such an IP address that is invalid or already in-use as the service address. The argument details the error.

26.5.9 BULK_LEASE_QUERY_PAUSE_CHECK_PERMISSIONS_FAILED

An unexpected error occurred while checking pause permissions, error %1

This error message is emitted when attempting to pause Bulk Lease Query's listener. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.5.10 BULK_LEASE_QUERY_PAUSE_LISTENER_FAILED

Listener could not be paused, error %1

This error message is emitted when attempting to pause Bulk Lease Query's listener. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.5.11 BULK_LEASE_QUERY_PAUSE_LISTENER_ILLEGAL

Pausing multi-threaded processing failed: %1

This error message is emitted when attempting to pause the bulk lease query listener's thread pool from a worker thread. This error indicates that an action attempted on listener thread is trying to use a critical section which would result in a dead-lock. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.5.12 BULK_LEASE_QUERY_PROCESSING_UNEXPECTED_FAILURE

A bulk lease query packet processing throws unexpected exception: %1

This error message is issued when bulk lease query processing throws. The exception is displayed.

26.5.13 BULK_LEASE_QUERY_QUERY_RECEIVED

A bulk lease query packet received from %1, details: %2

Logged at debug log level 40. This debug message is issued when a bulk lease query query has been received. The first argument is the address that sent the packet, the second details the packet.

26.5.14 BULK_LEASE_QUERY_REJECTED_CONNECTION

A new bulk lease query connection from %1 was rejected: %2

This debug message is issued when a new bulk lease query connection was rejected. The client address and the error message are displayed.

26.5.15 BULK_LEASE_QUERY_RESPONSE_SEND_ERROR

A bulk lease query response could not be sent to: %1, response: %2, error: %3

This debug message is issued when the server when an attempt to send a query response failed. The first argument is there address to which the response was destined, the second contains the response details, the third is the error explanation.

26.5.16 BULK_LEASE_QUERY_RESPONSE_SENT

A bulk lease query response sent to %1, details: %2

Logged at debug log level 40. This debug message is issued when a bulk lease query response has been sent. The first argument is the address that the packet has been sent to, the second details the packet.

26.5.17 BULK_LEASE_QUERY_RESUME_LISTENER_FAILED

Listener could not be resumed, error %1

This error message is emitted when attempting to resume Bulk Lease Query's listener. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.5.18 BULK_LEASE_QUERY_UNPACK_ERROR

A bulk lease query packet received from %1, could not be unpacked, error: %2

This error message is issued when received bulk lease query is malformed and could not be unpacked. The first argument is the remote address of the connection which will be closed, the second is the error explanation.

26.6 CB

26.6.1 CB_CMDS_CLASS4_DEL_HANDLER_FAILED

remote-class4-del command failed: %1

This error message is issued to indicate that the remote-class4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.2 CB_CMDS_CLASS4_GET_ALL_HANDLER_FAILED

remote-class4-get-all command failed: %1

This error message is issued to indicate that the remote-class4-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.3 CB_CMDS_CLASS4_GET_HANDLER_FAILED

remote-class4-get command failed: %1

This error message is issued to indicate that the remote-class4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.4 CB_CMDS_CLASS4_SET_HANDLER_FAILED

remote-class4-set command failed: %1

This error message is issued to indicate that the remote-class4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.5 CB_CMDS_CLASS6_DEL_HANDLER_FAILED

```
remote-class6-del command failed: %1
```

This error message is issued to indicate that the remote-class6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.6 CB_CMDS_CLASS6_GET_ALL_HANDLER_FAILED

```
remote-class6-get-all command failed: %1
```

This error message is issued to indicate that the remote-class6-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.7 CB_CMDS_CLASS6_GET_HANDLER_FAILED

```
remote-class6-get command failed: %1
```

This error message is issued to indicate that the remote-class6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.8 CB_CMDS_CLASS6_SET_HANDLER_FAILED

```
remote-class6-set command failed: %1
```

This error message is issued to indicate that the remote-class6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.9 CB_CMDS_DEINIT_OK

```
unloading cb_cmds hooks library successful
```

This informational message indicates that the Config Commands hooks library has been unloaded successfully.

26.6.10 CB_CMDS_GLOBAL_PARAMETER4_DEL_HANDLER_FAILED

```
remote-global-parameter4-del command failed: %1
```

This error message is issued to indicate that the remote-global-parameter4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.11 CB_CMDS_GLOBAL_PARAMETER4_GET_ALL_HANDLER_FAILED

```
remote-global-parameter4-get-all command failed: %1
```

This error message is issued to indicate that the remote-global-parameter4-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.12 CB_CMDS_GLOBAL_PARAMETER4_GET_HANDLER_FAILED

```
remote-global-parameter4-get command failed: %1
```

This error message is issued to indicate that the remote-global-parameter4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.13 CB_CMDS_GLOBAL_PARAMETER4_SET_HANDLER_FAILED

```
remote-global-parameter4-set command failed: %1
```

This error message is issued to indicate that the remote-global-parameter4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.14 CB_CMDS_GLOBAL_PARAMETER6_DEL_HANDLER_FAILED

```
remote-global-parameter6-del command failed: %1
```

This error message is issued to indicate that the remote-global-parameter6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.15 CB_CMDS_GLOBAL_PARAMETER6_GET_ALL_HANDLER_FAILED

```
remote-global-parameter6-get-all command failed: %1
```

This error message is issued to indicate that the remote-global-parameter6-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.16 CB_CMDS_GLOBAL_PARAMETER6_GET_HANDLER_FAILED

```
remote-global-parameter6-get command failed: %1
```

This error message is issued to indicate that the remote-global-parameter6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.17 CB_CMDS_GLOBAL_PARAMETER6_SET_HANDLER_FAILED

```
remote-global-parameter6-set command failed: %1
```

This error message is issued to indicate that the remote-global-parameter6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.18 CB_CMDS_INIT_OK

```
loading cb_cmds hooks library successful
```

This informational message indicates that the Config Commands hooks library has been loaded successfully. Enjoy!

26.6.19 CB_CMDS_NETWORK4_DEL_HANDLER_FAILED

```
remote-network4-del command failed: %1
```

This error message is issued to indicate that the remote-network4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.20 CB_CMDS_NETWORK4_GET_HANDLER_FAILED

```
remote-network4-get command failed: %1
```

This error message is issued to indicate that the remote-network4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.21 CB_CMDS_NETWORK4_LIST_HANDLER_FAILED

```
remote-network4-list command failed: %1
```

This error message is issued to indicate that the remote-network4-list command handler failed while processing the command. The argument provides the reason for failure.

26.6.22 CB_CMDS_NETWORK4_SET_HANDLER_FAILED

```
remote-network4-set command failed: %1
```

This error message is issued to indicate that the remote-network4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.23 CB_CMDS_NETWORK6_DEL_HANDLER_FAILED

```
remote-network6-del command failed: %1
```

This error message is issued to indicate that the remote-network6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.24 CB_CMDS_NETWORK6_GET_HANDLER_FAILED

```
remote-network6-get command failed: %1
```

This error message is issued to indicate that the remote-network6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.25 CB_CMDS_NETWORK6_LIST_HANDLER_FAILED

```
remote-network6-list command failed: %1
```

This error message is issued to indicate that the remote-network6-list command handler failed while processing the command. The argument provides the reason for failure.

26.6.26 CB_CMDS_NETWORK6_SET_HANDLER_FAILED

```
remote-network6-set command failed: %1
```

This error message is issued to indicate that the remote-network6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.27 CB_CMDS_OPTION4_GLOBAL_DEL_HANDLER_FAILED

```
remote-global-option4-del command failed: %1
```

This error message is issued to indicate that the remote-global-option4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.28 CB_CMDS_OPTION4_GLOBAL_GET_ALL_HANDLER_FAILED

```
remote-global-option4-get-all command failed: %1
```

This error message is issued to indicate that the remote-global-option4-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.29 CB_CMDS_OPTION4_GLOBAL_GET_HANDLER_FAILED

```
remote-global-option4-get command failed: %1
```

This error message is issued to indicate that the remote-global-option4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.30 CB_CMDS_OPTION4_GLOBAL_SET_HANDLER_FAILED

```
remote-global-option4-set command failed: %1
```

This error message is issued to indicate that the remote-global-option4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.31 CB_CMDS_OPTION4_NETWORK_DEL_HANDLER_FAILED

```
remote-network-option4-del command failed: %1
```

This error message is issued to indicate that the remote-network-option4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.32 CB_CMDS_OPTION4_NETWORK_SET_HANDLER_FAILED

```
remote-network-option4-set command failed: %1
```

This error message is issued to indicate that the remote-network-option4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.33 CB_CMDS_OPTION4_POOL_DEL_HANDLER_FAILED

```
remote-pool-option4-del command failed: %1
```

This error message is issued to indicate that the remote-pool-option4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.34 CB_CMDS_OPTION4_POOL_SET_HANDLER_FAILED

```
remote-pool-option4-set command failed: %1
```

This error message is issued to indicate that the remote-pool-option4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.35 CB_CMDS_OPTION4_SUBNET_DEL_HANDLER_FAILED

```
remote-subnet-option4-del command failed: %1
```

This error message is issued to indicate that the remote-subnet-option4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.36 CB_CMDS_OPTION4_SUBNET_SET_HANDLER_FAILED

```
remote-subnet-option4-set command failed: %1
```

This error message is issued to indicate that the remote-subnet-option4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.37 CB_CMDS_OPTION6_GLOBAL_DEL_HANDLER_FAILED

```
remote-global-option6-del command failed: %1
```

This error message is issued to indicate that the remote-global-option6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.38 CB_CMDS_OPTION6_GLOBAL_GET_ALL_HANDLER_FAILED

```
remote-global-option6-get-all command failed: %1
```

This error message is issued to indicate that the remote-global-option6-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.39 CB_CMDS_OPTION6_GLOBAL_GET_HANDLER_FAILED

```
remote-global-option6-get command failed: %1
```

This error message is issued to indicate that the remote-global-option6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.40 CB_CMDS_OPTION6_GLOBAL_SET_HANDLER_FAILED

```
remote-global-option6-set command failed: %1
```

This error message is issued to indicate that the remote-global-option6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.41 CB_CMDS_OPTION6_NETWORK_DEL_HANDLER_FAILED

```
remote-network-option6-del command failed: %1
```

This error message is issued to indicate that the remote-network-option6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.42 CB_CMDS_OPTION6_NETWORK_SET_HANDLER_FAILED

```
remote-network-option6-set command failed: %1
```

This error message is issued to indicate that the remote-network-option6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.43 CB_CMDS_OPTION6_PD_POOL_DEL_HANDLER_FAILED

```
remote-pd-pool-option6-del command failed: %1
```

This error message is issued to indicate that the remote-pd-pool-option6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.44 CB_CMDS_OPTION6_PD_POOL_SET_HANDLER_FAILED

```
remote-pd-pool-option6-set command failed: %1
```

This error message is issued to indicate that the remote-pd-pool-option6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.45 CB_CMDS_OPTION6_POOL_DEL_HANDLER_FAILED

```
remote-pool-option6-del command failed: %1
```

This error message is issued to indicate that the remote-pool-option6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.46 CB_CMDS_OPTION6_POOL_SET_HANDLER_FAILED

```
remote-pool-option6-set command failed: %1
```

This error message is issued to indicate that the remote-pool-option6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.47 CB_CMDS_OPTION6_SUBNET_DEL_HANDLER_FAILED

```
remote-subnet-option6-del command failed: %1
```

This error message is issued to indicate that the remote-subnet-option6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.48 CB_CMDS_OPTION6_SUBNET_SET_HANDLER_FAILED

```
remote-subnet-option6-set command failed: %1
```

This error message is issued to indicate that the remote-subnet-option6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.49 CB_CMDS_OPTION_DEF4_DEL_HANDLER_FAILED

```
remote-option-def4-del command failed: %1
```

This error message is issued to indicate that the remote-option-def4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.50 CB_CMDS_OPTION_DEF4_GET_ALL_HANDLER_FAILED

```
remote-option-def4-get-all command failed: %1
```

This error message is issued to indicate that the remote-option-def4-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.51 CB_CMDS_OPTION_DEF4_GET_HANDLER_FAILED

```
remote-option-def4-get command failed: %1
```

This error message is issued to indicate that the remote-option-def4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.52 CB_CMDS_OPTION_DEF4_SET_HANDLER_FAILED

```
remote-option-def4-set command failed: %1
```

This error message is issued to indicate that the remote-option-def4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.53 CB_CMDS_OPTION_DEF6_DEL_HANDLER_FAILED

```
remote-option-def6-del command failed: %1
```

This error message is issued to indicate that the remote-option-def6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.54 CB_CMDS_OPTION_DEF6_GET_ALL_HANDLER_FAILED

```
remote-option-def6-get-all command failed: %1
```

This error message is issued to indicate that the remote-option-def6-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.55 CB_CMDS_OPTION_DEF6_GET_HANDLER_FAILED

```
remote-option-def6-get command failed: %1
```

This error message is issued to indicate that the remote-option-def6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.56 CB_CMDS_OPTION_DEF6_SET_HANDLER_FAILED

```
remote-option-def6-set command failed: %1
```

This error message is issued to indicate that the remote-option-def6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.57 CB_CMDS_SERVER4_DEL_HANDLER_FAILED

```
remote-server4-del command failed: %1
```

This error message is issued to indicate that the remote-server4-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.58 CB_CMDS_SERVER4_GET_ALL_HANDLER_FAILED

```
remote-server4-get-all command failed: %1
```

This error message is issued to indicate that the remote-server4-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.59 CB_CMDS_SERVER4_GET_HANDLER_FAILED

```
remote-server4-get command failed: %1
```

This error message is issued to indicate that the remote-server4-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.60 CB_CMDS_SERVER4_SET_HANDLER_FAILED

```
remote-server4-set command failed: %1
```

This error message is issued to indicate that the remote-server4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.61 CB_CMDS_SERVER6_DEL_HANDLER_FAILED

```
remote-server6-del command failed: %1
```

This error message is issued to indicate that the remote-server6-del command handler failed while processing the command. The argument provides the reason for failure.

26.6.62 CB_CMDS_SERVER6_GET_ALL_HANDLER_FAILED

```
remote-server6-get-all command failed: %1
```

This error message is issued to indicate that the remote-server6-get-all command handler failed while processing the command. The argument provides the reason for failure.

26.6.63 CB_CMDS_SERVER6_GET_HANDLER_FAILED

```
remote-server6-get command failed: %1
```

This error message is issued to indicate that the remote-server6-get command handler failed while processing the command. The argument provides the reason for failure.

26.6.64 CB_CMDS_SERVER6_SET_HANDLER_FAILED

```
remote-server6-set command failed: %1
```

This error message is issued to indicate that the remote-server6-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.65 CB_CMDS_SUBNET4_DEL_BY_ID_HANDLER_FAILED

```
remote-subnet4-del-by-id command failed: %1
```

This error message is issued to indicate that the remote-subnet4-del-by-id command handler failed while processing the command. The argument provides the reason for failure.

26.6.66 CB_CMDS_SUBNET4_DEL_BY_PREFIX_HANDLER_FAILED

```
remote-subnet4-del-by-prefix command failed: %1
```

This error message is issued to indicate that the remote-subnet4-del-by-prefix command handler failed while processing the command. The argument provides the reason for failure.

26.6.67 CB_CMDS_SUBNET4_GET_BY_ID_HANDLER_FAILED

```
remote-subnet4-get-by-id command failed: %1
```

This error message is issued to indicate that the remote-subnet4-get-by-id command handler failed while processing the command. The argument provides the reason for failure.

26.6.68 CB_CMDS_SUBNET4_GET_BY_PREFIX_HANDLER_FAILED

```
remote-subnet4-get-by-prefix command failed: %1
```

This error message is issued to indicate that the remote-subnet4-get-by-prefix command handler failed while processing the command. The argument provides the reason for failure.

26.6.69 CB_CMDS_SUBNET4_LIST_HANDLER_FAILED

```
remote-subnet4-list command failed: %1
```

This error message is issued to indicate that the remote-subnet4-list command handler failed while processing the command. The argument provides the reason for failure.

26.6.70 CB_CMDS_SUBNET4_SET_HANDLER_FAILED

```
remote-subnet4-set command failed: %1
```

This error message is issued to indicate that the remote-subnet4-set command handler failed while processing the command. The argument provides the reason for failure.

26.6.71 CB_CMDS_SUBNET6_DEL_BY_ID_HANDLER_FAILED

```
remote-subnet6-del-by-id command failed: %1
```

This error message is issued to indicate that the remote-subnet6-del-by-id command handler failed while processing the command. The argument provides the reason for failure.

26.6.72 CB_CMDS_SUBNET6_DEL_BY_PREFIX_HANDLER_FAILED

```
remote-subnet6-del-by-prefix command failed: %1
```

This error message is issued to indicate that the remote-subnet6-del-by-prefix command handler failed while processing the command. The argument provides the reason for failure.

26.6.73 CB_CMDS_SUBNET6_GET_BY_ID_HANDLER_FAILED

```
remote-subnet6-get-by-id command failed: %1
```

This error message is issued to indicate that the remote-subnet6-get-by-id command handler failed while processing the command. The argument provides the reason for failure.

26.6.74 CB_CMDS_SUBNET6_GET_BY_PREFIX_HANDLER_FAILED

```
remote-subnet6-get-by-prefix command failed: %1
```

This error message is issued to indicate that the remote-subnet6-get-by-prefix command handler failed while processing the command. The argument provides the reason for failure.

26.6.75 CB_CMDS_SUBNET6_LIST_HANDLER_FAILED

```
remote-subnet6-list command failed: %1
```

This error message is issued to indicate that the remote-subnet6-list command handler failed while processing the command. The argument provides the reason for failure.

26.6.76 CB_CMDS_SUBNET6_SET_HANDLER_FAILED

```
remote-subnet6-set command failed: %1
```

This error message is issued to indicate that the remote-subnet6-set command handler failed while processing the command. The argument provides the reason for failure.

26.7 CLASS

26.7.1 CLASS_CMDS_CLASS_ADD

```
class added: %1
```

The class-add command has been successful. Definition of the added class is logged.

26.7.2 CLASS_CMDS_CLASS_ADD_FAILED

```
failed to add a new class: %1
```

The class-add command has failed. The reason for failure is provided within the error message.

26.7.3 CLASS_CMDS_CLASS_ADD_HANDLER_FAILED

```
failed to run handler for 'class-add' command
```

This error message is logged when the class_cmds hooks library experiences an unexpected error when 'class-add' command handler fails to execute.

26.7.4 CLASS_CMDS_CLASS_DEL

```
class deleted: %1
```

The class-del command has been successful. The name of the deleted class is logged.

26.7.5 CLASS_CMDS_CLASS_DEL_EMPTY

```
class not deleted (not found): %1
```

The class-del command found no matching class. The name of the to be deleted class is logged.

26.7.6 CLASS_CMDS_CLASS_DEL_FAILED

```
failed to delete a class: %1
```

The class-del command has failed. The reason for failure is provided within the error message.

26.7.7 CLASS_CMDS_CLASS_DEL_HANDLER_FAILED

```
failed to run handler for 'class-del' command
```

This error message is logged when the class_cmds hooks library experiences an unexpected error when 'class-del' command handler fails to execute.

26.7.8 CLASS_CMDS_CLASS_GET

```
successfully retrieved a class: %1
```

The class-get command has been successful. The name of the retrieved class is logged.

26.7.9 CLASS_CMDS_CLASS_GET_EMPTY

specified class was not found: %1

The class-get command found no matching class. The name of the searched class is logged.

26.7.10 CLASS_CMDS_CLASS_GET_FAILED

failed to retrieve a class: %1

The class-get command has failed. The reason for failure is provided within the error message.

26.7.11 CLASS_CMDS_CLASS_GET_HANDLER_FAILED

failed to run handler for 'class-get' command

This error message is logged when the class_cmds hooks library experiences an unexpected error when 'class-get' command handler fails to execute.

26.7.12 CLASS_CMDS_CLASS_LIST

successfully retrieved classes names

The class-list command has been successful.

26.7.13 CLASS_CMDS_CLASS_LIST_EMPTY

no class was found

The class-list command found no class.

26.7.14 CLASS_CMDS_CLASS_LIST_FAILED

failed to retrieve classes names: %1

The class-list command has failed. The reason for failure is provided within the error message.

26.7.15 CLASS_CMDS_CLASS_LIST_HANDLER_FAILED

failed to run handler for 'class-list' command

This error message is logged when the class_cmds hooks library experiences an unexpected error when 'class-list' command handler fails to execute.

26.7.16 CLASS_CMDS_CLASS_UPDATE

class updated: %1

The class-update command has been successful. Parameters of the updated class are logged.

26.7.17 CLASS_CMDS_CLASS_UPDATE_EMPTY

```
class not updated (not found): %1
```

The class-update command found no matching class. The name of the to be updated class is logged.

26.7.18 CLASS_CMDS_CLASS_UPDATE_FAILED

```
failed to update a class: %1
```

The class-update command has failed. The reason for failure is provided within the error message.

26.7.19 CLASS_CMDS_CLASS_UPDATE_HANDLER_FAILED

```
failed to run handler for 'class-update' command
```

This error message is logged when the class_cmds hooks library experiences an unexpected error when 'class-update' command handler fails to execute.

26.7.20 CLASS_CMDS_DEINIT_OK

```
unloading Class Commands hooks library successful
```

This info message indicates that the Class Commands hooks library has been removed successfully.

26.7.21 CLASS_CMDS_INIT_FAILED

```
loading Class Commands hooks library failed: %1
```

This error message indicates an error during loading the Class Commands hooks library. The details of the error are provided as argument of the log message.

26.7.22 CLASS_CMDS_INIT_OK

```
loading Class Commands hooks library successful
```

This info message indicates that the Class Commands hooks library has been loaded successfully. Enjoy!

26.8 COMMAND

26.8.1 COMMAND_ACCEPTOR_START

```
Starting to accept connections via unix domain socket bound to %1
```

This informational message is issued when the Kea server starts an acceptor via which it is going to accept new control connections. The acceptor is bound to the endpoint associated with the filename provided as an argument. If starting the acceptor fails, subsequent error messages will provide a reason for failure.

26.8.2 COMMAND_DEREGISTERED

```
Command %1 deregistered
```

Logged at debug log level 10. This debug message indicates that the daemon stopped supporting specified command. This command can no longer be issued. If the command socket is open and this command is issued, the daemon will not be able to process it.

26.8.3 COMMAND_EXTENDED_REGISTERED

```
Command %1 registered
```

Logged at debug log level 10. This debug message indicates that the daemon started supporting specified command. The handler for the registered command includes a parameter holding entire command to be processed.

26.8.4 COMMAND_HTTP_LISTENER_COMMAND_REJECTED

```
Command HTTP listener rejected command '%1' from '%2'
```

Logged at debug log level 10. This debug messages is issued when a command is rejected. Arguments detail the command and the address the request was received from.

26.8.5 COMMAND_HTTP_LISTENER_STARTED

```
Command HTTP listener started with %1 threads, listening on address: %2 port: %3, use_↵  
↵TLS: %4
```

Logged at debug log level 10. This debug messages is issued when an HTTP listener has been started to accept connections from Command API clients through which commands can be received and responses sent. Arguments detail the number of threads that the listener is using, the address and port at which it is listening, and if HTTPS/TLS is used or not.

26.8.6 COMMAND_HTTP_LISTENER_STOPPED

```
Command HTTP listener for address: %1 port: %2 stopped.
```

Logged at debug log level 10. This debug messages is issued when the Command HTTP listener, listening at the given address and port, has completed shutdown.

26.8.7 COMMAND_HTTP_LISTENER_STOPPING

```
Stopping Command HTTP listener for address: %1 port: %2
```

Logged at debug log level 10. This debug messages is issued when the Command HTTP listener, listening at the given address and port, has begun to shutdown.

26.8.8 COMMAND_PROCESS_ERROR1

```
Error while processing command: %1
```

This warning message indicates that the server encountered an error while processing received command. Additional information will be provided, if available. Additional log messages may provide more details.

26.8.9 COMMAND_PROCESS_ERROR2

Error while processing command: %1

This warning message indicates that the server encountered an error while processing received command. The difference, compared to COMMAND_PROCESS_ERROR1 is that the initial command was well formed and the error occurred during logic processing, not the command parsing. Additional information will be provided, if available. Additional log messages may provide more details.

26.8.10 COMMAND_RECEIVED

Received command '%1'

This informational message indicates that a command was received over command socket. The nature of this command and its possible results will be logged with separate messages.

26.8.11 COMMAND_REGISTERED

Command %1 registered

Logged at debug log level 10. This debug message indicates that the daemon started supporting specified command. If the command socket is open, this command can now be issued.

26.8.12 COMMAND_RESPONSE_ERROR

Server failed to generate response for command: %1

This error message indicates that the server failed to generate response for specified command. This likely indicates a server logic error, as the server is expected to generate valid responses for all commands, even malformed ones.

26.8.13 COMMAND_SOCKET_ACCEPT_FAIL

Failed to accept incoming connection on command socket %1: %2

This error indicates that the server detected incoming connection and executed accept system call on said socket, but this call returned an error. Additional information may be provided by the system as second parameter.

26.8.14 COMMAND_SOCKET_CLOSED_BY_FOREIGN_HOST

Closed command socket %1 by foreign host, %2

This is an information message indicating that the command connection has been closed by a command control client, and whether any partially read data was discarded.

26.8.15 COMMAND_SOCKET_CONNECTION_CANCEL_FAIL

Failed to cancel read operation on socket %1: %2

This error message is issued to indicate an error to cancel asynchronous read of the control command over the control socket. The cancel operation is performed when the timeout occurs during communication with a client. The error message includes details about the reason for failure.

26.8.16 COMMAND_SOCKET_CONNECTION_CLOSED

Closed socket %1 for existing command connection

Logged at debug log level 10. This is a debug message indicating that the socket created for handling client's connection is closed. This usually means that the client disconnected, but may also mean a timeout.

26.8.17 COMMAND_SOCKET_CONNECTION_CLOSE_FAIL

Failed to close command connection: %1

This error message is issued when an error occurred when closing a command connection and/or removing it from the connections pool. The detailed error is provided as an argument.

26.8.18 COMMAND_SOCKET_CONNECTION_OPENED

Opened socket %1 for incoming command connection

Logged at debug log level 10. This is a debug message indicating that a new incoming command connection was detected and a dedicated socket was opened for that connection.

26.8.19 COMMAND_SOCKET_CONNECTION_SHUTDOWN_FAIL

Encountered error %1 while trying to gracefully shutdown socket

This message indicates an error while trying to gracefully shutdown command connection. The type of the error is included in the message.

26.8.20 COMMAND_SOCKET_CONNECTION_TIMEOUT

Timeout occurred for connection over socket %1

This is an informational message that indicates that the timeout has occurred for one of the command channel connections. The response sent by the server indicates a timeout and is then closed.

26.8.21 COMMAND_SOCKET_READ

Received %1 bytes over command socket %2

Logged at debug log level 10. This debug message indicates that specified number of bytes was received over command socket identified by specified file descriptor.

26.8.22 COMMAND_SOCKET_READ_FAIL

Encountered error %1 while reading from command socket %2

This error message indicates that an error was encountered while reading from command socket.

26.8.23 COMMAND_SOCKET_WRITE

Sent response of %1 bytes (%2 bytes left to send) over command socket %3

Logged at debug log level 10. This debug message indicates that the specified number of bytes was sent over command socket identifier by the specified file descriptor.

26.8.24 COMMAND_SOCKET_WRITE_FAIL

```
Error while writing to command socket %1 : %2
```

This error message indicates that an error was encountered while attempting to send a response to the command socket.

26.8.25 COMMAND_WATCH_SOCKET_CLEAR_ERROR

```
watch socket failed to clear: %1
```

This error message is issued when the command manager was unable to reset the ready status after completing a send. This is a programmatic error that should be reported. The command manager may or may not continue to operate correctly.

26.8.26 COMMAND_WATCH_SOCKET_CLOSE_ERROR

```
watch socket failed to close: %1
```

This error message is issued when command manager attempted to close the socket used for indicating the ready status for send operations. This should not have any negative impact on the operation of the command manager as it happens when the connection is being terminated.

26.8.27 COMMAND_WATCH_SOCKET_MARK_READY_ERROR

```
watch socket failed to mark ready: %1
```

This error message is issued when the command manager was unable to set ready status after scheduling asynchronous send. This is programmatic error that should be reported. The command manager may or may not continue to operate correctly.

26.9 CONFIG

26.9.1 CONFIG_BACKENDS_REGISTERED

```
the following config backend types are available: %1
```

This informational message lists all possible config backends that could be used in config-database[s].

26.10 CTRL

26.10.1 CTRL_AGENT_COMMAND_FORWARDED

```
command %1 successfully forwarded to the service %2 from remote address %3
```

This informational message is issued when the CA successfully forwards the control message to the specified Kea service and receives a response.

26.10.2 CTRL_AGENT_COMMAND_FORWARD_BEGIN

```
begin forwarding command %1 to service %2
```

Logged at debug log level 10. This debug message is issued when the Control Agent starts forwarding a received command to one of the Kea servers.

26.10.3 CTRL_AGENT_COMMAND_FORWARD_FAILED

```
failed forwarding command %1: %2
```

Logged at debug log level 10. This debug message is issued when the Control Agent failed forwarding a received command to one of the Kea servers. The second argument provides the details of the error.

26.10.4 CTRL_AGENT_COMMAND_RECEIVED

```
command %1 received from remote address %2
```

This informational message is issued when the CA receives a control message, whether it is destined to the control agent itself, or to be forwarded on.

26.10.5 CTRL_AGENT_CONFIG_CHECK_FAIL

```
Control Agent configuration check failed: %1
```

This error message indicates that the CA had failed configuration check. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.10.6 CTRL_AGENT_CONFIG_FAIL

```
Control Agent configuration failed: %1
```

This error message indicates that the CA had failed configuration attempt. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.10.7 CTRL_AGENT_CONFIG_SYNTAX_WARNING

```
Control Agent configuration syntax warning: %1
```

This warning message indicates that the CA configuration had a minor syntax error. The error was displayed and the configuration parsing resumed.

26.10.8 CTRL_AGENT_FAILED

```
application experienced a fatal error: %1
```

This is a fatal error message issued when the Control Agent application encounters an unrecoverable error from within the event loop.

26.10.9 CTRL_AGENT_HTTPS_SERVICE_REUSE_FAILED

```
failed to reuse HTTPS service bound to address: %1 port: %2
```

This error message indicates that the server has failed reusing existing HTTPS service on the specified address and port. The server can not swith from HTTPS to HTTP sockets using the same address and port.

26.10.10 CTRL_AGENT_HTTPS_SERVICE_STARTED

```
HTTPS service bound to address: %1 port: %2
```

This informational message indicates that the server has started HTTPS service on the specified address and port. All control commands should be sent to this address and port over a TLS channel.

26.10.11 CTRL_AGENT_HTTPS_SERVICE_UPDATED

```
reused HTTPS service bound to address: %1 port: %2 and updated TLS settings
```

This informational message indicates that the server has reused existing HTTPS service on the specified address and port. Note that any change in the TLS setup has been applied.

26.10.12 CTRL_AGENT_HTTP_SERVICE_REUSE_FAILED

```
failed to reused HTTP service bound to address: %1 port: %2
```

This error message indicates that the server has failed reusing existing HTTP service on the specified address and port. The server can not switch from HTTP to HTTPS sockets using the same address and port.

26.10.13 CTRL_AGENT_HTTP_SERVICE_STARTED

```
HTTP service bound to address: %1 port: %2
```

This informational message indicates that the server has started HTTP service on the specified address and port. All control commands should be sent to this address and port.

26.10.14 CTRL_AGENT_HTTP_SERVICE_UPDATED

```
reused HTTP service bound to address: %1 port: %2
```

This informational message indicates that the server has reused existing HTTP service on the specified address and port.

26.10.15 CTRL_AGENT_IS_DEPRECATED

```
Kea Control Agent is deprecated. Its function has been moved to Kea servers.
```

This warning message indicates that the Control Agent has been deprecated. All its function has been moved to Kea servers.

26.10.16 CTRL_AGENT_RUN_EXIT

```
application is exiting the event loop
```

Logged at debug log level 0. This is a debug message issued when the Control Agent exits its event loop.

26.10.17 CTRL_AGENT_SECURITY_CHECKS_DISABLED

```
Invoked with command line option -X, Security checks are disabled!!
```

This warning is emitted when internal security checks normally performed by kea-ctrl-agent have been disabled via command line option '-X'. This means the server is not enforcing restrictions on resource paths or permissions. This mode of operation may expose your environment to security vulnerabilities and should only be used after consideration.

26.10.18 CTRL_AGENT_STARTED

```
Kea Control Agent version %1 started
```

This informational message indicates that the Control Agent has processed all configuration information and is ready to begin processing. The version is also printed.

26.11 DATABASE

26.11.1 DATABASE_INVALID_ACCESS

```
invalid database access string: %1
```

This is logged when an attempt has been made to parse a database access string and the attempt ended in error. The access string in question - which should be of the form 'keyword=value keyword=value...' is included in the message.

26.11.2 DATABASE_MYSQL_COMMIT

```
committing to MySQL database
```

The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the MySQL settings, the committal may not include a write to disk.

26.11.3 DATABASE_MYSQL_FATAL_ERROR

```
Unrecoverable MySQL error occurred: %1 for <%2>, reason: %3 (error code: %4).
```

An error message indicating that communication with the MySQL database server has been lost. If automatic recovery has been enabled, then the server will attempt to recover connectivity. If not, then the server will exit with a non-zero exit code. The cause of such an error is most likely a network issue or the MySQL server has gone down.

26.11.4 DATABASE_MYSQL_INITIALIZE_SCHEMA

```
Initializing the MySQL schema with command: %1.
```

This is logged before running the kea-admin command to automatically initialize the schema from Kea after getting the schema version initially failed. The full kea-admin command is shown.

26.11.5 DATABASE_MYSQL_ROLLBACK

```
rolling back MySQL database
```

The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.11.6 DATABASE_MYSQL_START_TRANSACTION

starting new MySQL transaction

A debug message issued when a new MySQL transaction is being started. This message is typically not issued when inserting data into a single table because the server doesn't explicitly start transactions in this case. This message is issued when data is inserted into multiple tables with multiple INSERT statements and there may be a need to rollback the whole transaction if any of these INSERT statements fail.

26.11.7 DATABASE_PGSQL_COMMIT

committing to PostgreSQL database

The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the PostgreSQL settings, the committal may not include a write to disk.

26.11.8 DATABASE_PGSQL_CREATE_SAVEPOINT

creating a new PostgreSQL savepoint: %1

The code is issuing a call to create a savepoint within the current transaction. Database modifications made up to this point will be preserved should a subsequent call to rollback to this savepoint occurs prior to the transaction being committed.

26.11.9 DATABASE_PGSQL_DEALLOC_ERROR

An error occurred deallocating SQL statements while closing the PostgreSQL lease.
↪ database: %1

This is an error message issued when a DHCP server (either V4 or V6) experienced an error freeing database SQL resources as part of closing its connection to the PostgreSQL database. The connection is closed as part of normal server shutdown. This error is most likely a programmatic issue that is highly unlikely to occur or negatively impact server operation.

26.11.10 DATABASE_PGSQL_FATAL_ERROR

Unrecoverable PostgreSQL error occurred: Statement: <%1>, reason: %2 (error code: %3).

An error message indicating that communication with the PostgreSQL database server has been lost. If automatic recovery has been enabled, then the server will attempt to recover the connectivity. If not, then the server will exit with a non-zero exit code. The cause of such an error is most likely a network issue or the PostgreSQL server has gone down.

26.11.11 DATABASE_PGSQL_INITIALIZE_SCHEMA

Initializing the PostgreSQL schema with command: %1.

This is logged before running the kea-admin command to automatically initialize the schema from Kea after getting the schema version initially failed. The full kea-admin command is shown.

26.11.12 DATABASE_PGSQL_ROLLBACK

rolling back PostgreSQL database

The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.11.13 DATABASE_PGSQL_ROLLBACK_SAVEPOINT

rolling back PostgreSQL database to savepoint: \$1

The code is issuing a call to rollback to the given savepoint. Any database modifications that were made after the savepoint was created will be rolled back and not committed to the database.

26.11.14 DATABASE_PGSQL_START_TRANSACTION

starting a new PostgreSQL transaction

A debug message issued when a new PostgreSQL transaction is being started. This message is typically not issued when inserting data into a single table because the server doesn't explicitly start transactions in this case. This message is issued when data is inserted into multiple tables with multiple INSERT statements and there may be a need to rollback the whole transaction if any of these INSERT statements fail.

26.11.15 DATABASE_PGSQL_TCP_USER_TIMEOUT_UNSUPPORTED

tcp_user_timeout is not supported in this PostgreSQL version

This warning message is issued when a user has configured the tcp_user_timeout parameter in the connection to the PostgreSQL database but the installed database does not support this parameter. It is supported by the PostgreSQL version 12 or later. The parameter setting will be ignored.

26.11.16 DATABASE_TO_JSON_BOOLEAN_ERROR

Internal logic error: invalid boolean value found in database connection parameters: %1=
↪%2

This error message is printed when conversion to JSON of the internal state is requested, but the connection string contains a boolean parameter with invalid value. It is a programming error. The software will continue operation, but the returned JSON data will be syntactically valid, but incomplete. The culprit parameter will not be converted.

26.11.17 DATABASE_TO_JSON_INTEGER_ERROR

Internal logic error: invalid integer value found in database connection parameters: %1=
↪%2

This error message is printed when conversion to JSON of the internal state is requested, but the connection string contains the integer parameter with a wrong value. It is a programming error. The software will continue operation, but the returned JSON data will be syntactically valid, but incomplete. The culprit parameter will not be converted.

26.11.18 DATABASE_TO_JSON_UNKNOWN_TYPE_ERROR

Internal logic error: unknown element found in database connection parameters: %1=%2

This error message is printed when conversion to JSON of the internal state is requested, but the connection string contains unrecognized parameter. It is a programming error. The software will continue operation, but the returned JSON data will be syntactically valid, but incomplete. The unknown parameter will not be converted.

26.12 DCTL

26.12.1 DCTL_ALREADY_RUNNING

```
%1 already running? %2
```

This is an error message that occurs when a module encounters a pre-existing PID file which contains the PID of a running process. This most likely indicates an attempt to start a second instance of a module using the same configuration file. It is possible, though unlikely, that the PID file is a remnant left behind by a server crash or power failure and the PID it contains refers to a process other than Kea process. In such an event, it would be necessary to manually remove the PID file. The first argument is the process name, the second contains the PID and PID file.

26.12.2 DCTL_CFG_FILE_RELOAD_ERROR

```
configuration reload failed: %1, reverting to current configuration.
```

This is an error message indicating that the application attempted to reload its configuration from file and encountered an error. This is likely due to invalid content in the configuration file. The application should continue to operate under its current configuration.

26.12.3 DCTL_CFG_FILE_RELOAD_SIGNAL_RECVD

```
OS signal %1 received, reloading configuration from file: %2
```

This is an informational message indicating the application has received a signal instructing it to reload its configuration from file.

26.12.4 DCTL_CONFIG_CHECK_COMPLETE

```
server has completed configuration check: %1, result: %2
```

This is an informational message announcing the successful processing of a new configuration check is complete. The result of that check is printed. This informational message is printed when configuration check is requested.

26.12.5 DCTL_CONFIG_COMPLETE

```
server has completed configuration: %1
```

This is an informational message announcing the successful processing of a new configuration. It is output during server startup, and when an updated configuration is committed by the administrator. Additional information may be provided.

26.12.6 DCTL_CONFIG_DEPRECATED

```
server configuration includes a deprecated object: %1
```

This error message is issued when the configuration includes a deprecated object (i.e. a top level element) which will be ignored.

26.12.7 DCTL_CONFIG_FETCH

Fetching configuration data from config backends.

This is an informational message emitted when the Kea server is about to begin retrieving configuration data from one or more configuration backends.

26.12.8 DCTL_CONFIG_FILE_LOAD_FAIL

%1 reason: %2

This fatal error message indicates that the application attempted to load its initial configuration from file and has failed. The service will exit.

26.12.9 DCTL_CONFIG_START

parsing new configuration: %1

Logged at debug log level 10. A debug message indicating that the application process has received an updated configuration and has passed it to its configuration manager for parsing.

26.12.10 DCTL_DB_OPEN_CONNECTION_WITH_RETRY_FAILED

Failed to connect to database: %1 with error: %2

This is an informational message issued when the server failed to connect to the configuration database. The operation started a retry to connect procedure. The database access string with password redacted is logged, along with the error and details for the reconnect procedure.

26.12.11 DCTL_DEPRECATED_OUTPUT_OPTIONS

The output_options parameter is deprecated. Use output-options parameter instead.

This warning message is displayed when deprecated output_options is used instead of output-options.

26.12.12 DCTL_DEVELOPMENT_VERSION

This software is a development branch of Kea. It is not recommended for production use.

This warning message is displayed when the version is a development (vs stable) one: the second number of the version is odd.

26.12.13 DCTL_INIT_PROCESS

%1 initializing the application

Logged at debug log level 0. This debug message is issued just before the controller attempts to create and initialize its application instance.

26.12.14 DCTL_INIT_PROCESS_FAIL

```
%1 application initialization failed: %2
```

This error message is issued if the controller could not initialize the application and will exit.

26.12.15 DCTL_NOT_RUNNING

```
%1 application instance is not running
```

A warning message is issued when an attempt is made to shut down the application when it is not running.

26.12.16 DCTL_OPEN_CONFIG_DB

```
Opening configuration database: %1
```

This message is printed when the Kea server is attempting to open a configuration database. The database access string with password redacted is logged.

26.12.17 DCTL_PARSER_FAIL

```
Parser error: %1
```

On receipt of a new configuration, the server failed to create a parser to decode the contents of the named configuration element, or the creation succeeded but the parsing actions and committal of changes failed. The reason for the failure is given in the message.

26.12.18 DCTL_PID_FILE_ERROR

```
%1 could not create a PID file: %2
```

This is an error message that occurs when the server is unable to create its PID file. The log message should contain details sufficient to determine the underlying cause. The most likely culprits are that some portion of the pathname does not exist or a permissions issue. The default path is determined by `--localstatedir` and `--prefix` meson setup options but may be overridden by setting environment variable `KEA_PIDFILE_DIR`. The first argument is the process name.

26.12.19 DCTL_PROCESS_FAILED

```
%1 application execution failed: %2
```

The controller has encountered a fatal error while running the application and is terminating. The reason for the failure is included in the message.

26.12.20 DCTL_RUN_PROCESS

```
%1 starting application event loop
```

Logged at debug log level 0. This debug message is issued just before the controller invokes the application run method.

26.12.21 DCTL_SHUTDOWN

```
%1 has shut down, pid: %2, version: %3
```

Logged at debug log level 0. This is an informational message indicating that the service has shut down. The argument specifies a name of the service.

26.12.22 DCTL_SHUTDOWN_SIGNAL_RECVD

```
OS signal %1 received, starting shutdown
```

Logged at debug log level 0. This is a debug message indicating the application has received a signal instructing it to shutdown.

26.12.23 DCTL_STANDALONE

```
%1 skipping message queue, running standalone
```

Logged at debug log level 0. This is a debug message indicating that the controller is running in the application in standalone mode. This means it will not connected to the Kea message queue. Standalone mode is only useful during program development, and should not be used in a production environment.

26.12.24 DCTL_STARTING

```
%1 starting, pid: %2, version: %3 (%4)
```

This is an informational message issued when controller for the service first starts. Version is also reported.

26.12.25 DCTL_UNLOAD_LIBRARIES_ERROR

```
error unloading hooks libraries during shutdown: %1
```

This error message indicates that during shutdown, unloading hooks libraries failed to close them. If the list of libraries is empty it is a programmatic error in the server code. If it is not empty it could be a programmatic error in one of the hooks libraries which could lead to a crash during finalization.

26.12.26 DCTL_UNSUPPORTED_SIGNAL

```
ignoring reception of unsupported signal: %1
```

This is a debug message indicating that the application received an unsupported signal. This is a programming error indicating that the application has registered to receive the signal but no associated processing logic has been added.

26.13 DDNS

26.13.1 DDNS_TUNING4_CALCULATED_HOSTNAME

```
Replacing host name: %1, with calculated host name: %2, for query: %3
```

Logged at debug log level 40. This debug message is emitted when the DDNS Tuning hooks library has calculated a new host name for the given client query. The original host name and the calculated host name are provided

26.13.2 DDNS_TUNING4_PROCESS_ERROR

```
An error occurred processing query %1: %2
```

This error message indicates an error during processing of a query by the DDNS Tuning hooks library. The client identification information from the query and the details of the error are provided as arguments of the log message.

26.13.3 DDNS_TUNING4_SKIPPING_DDNS

```
Client is a member matches SKIP_DDNS class, skipping DDNS updates, query: %1
```

Logged at debug log level 40. This debug message is emitted when during the ddns4-update callout, if the Client query matches the SKIP_DDNS class. The kea-dhcp4 server will not send DDNS update requests to kea-dhcp-ddns for this query. The query is shown in the log message.

26.13.4 DDNS_TUNING6_CALCULATED_HOSTNAME

```
Replacing host name: %1, with calculated host name: %2, for query: %3
```

Logged at debug log level 40. This debug message is emitted when the DDNS Tuning hooks library has calculated a new host name for the given client query. The original host name and the calculated host name are provided

26.13.5 DDNS_TUNING6_PROCESS_ERROR

```
An error occurred processing query %1: %2
```

This error message indicates an error during processing of a query by the DDNS Tuning hooks library. The client identification information from the query and the details of the error are provided as arguments of the log message.

26.13.6 DDNS_TUNING6_SKIPPING_DDNS

```
Client is a member matches SKIP_DDNS class, skipping DDNS updates, query: %1
```

Logged at debug log level 40. This debug message is emitted when during the ddns6-update callout, if the Client query matches the SKIP_DDNS class. The kea-dhcp6 server will not send DDNS update requests to kea-dhcp-ddns for this query. The query is shown in the log message.

26.13.7 DDNS_TUNING_GLOBAL_EXPR_SET

```
Global hostname expression set to: %1
```

This message indicates that the global hostname expression has been set and will be used everywhere, unless overridden by subnet level parameters.

26.13.8 DDNS_TUNING_LOAD_ERROR

```
loading DDNS Tuning hooks library failed: %1
```

This error message indicates an error during loading the DDNS Tuning hooks library. The details of the error are provided as argument of the log message.

26.13.9 DDNS_TUNING_LOAD_OK

```
DDNS Tuning hooks library loaded successfully.
```

This info message indicates that the DDNS Tuning hooks library has been loaded successfully.

26.13.10 DDNS_TUNING_SUBNET_EXPRESSION_PARSE

```
Parsing subnet expression (%1) for for subnet %2
```

Logged at debug log level 40. This debug message is emitted when the DDNS Tuning hooks library is attempting to parse the hostname expression for a subnet. Parsing occurs after configuration events (e.g. reconfigure command, config back end updates, subnet command updates).

26.13.11 DDNS_TUNING_SUBNET_EXPRESSION_PARSE_ERROR

```
An error occurred while parsing the hostname expression for subnet %1, %2
```

This error indicates that the hostname expression assigned to a subnet is not a valid expression. DHCP service will continue as though the subnet did not specify an expression. The subnet and the parsing error are included in the log message.

26.13.12 DDNS_TUNING_SUBNET_EXPR_CACHED

```
Using subnet expression stored in a cache for subnet %1
```

Logged at debug log level 40. The expression for this subnet has previously been used and cached. Using the cached version.

26.13.13 DDNS_TUNING_UNLOAD

```
DDNS Tuning hooks library has been unloaded
```

This info message indicates that the DDNS Tuning hooks library has been unloaded.

26.14 DHCP4

26.14.1 DHCP4_ADDITIONAL_CLASS_EVAL_ERROR

```
%1: Expression '%2' evaluated to %3
```

This error message indicates that a problem was encountered while evaluating the expression of an additional client class. A description of the problem is printed.

26.14.2 DHCP4_ADDITIONAL_CLASS_EVAL_RESULT

```
%1: Expression '%2' evaluated to %3
```

Logged at debug log level 50. This debug message indicates that the expression of an additional client class has been successfully evaluated. The client class name and the result value of the evaluation are printed.

26.14.3 DHCP4_ADDITIONAL_CLASS_NO_TEST

```
additional class %1 has no test expression, adding it to client's classes unconditionally
```

Logged at debug log level 40. This debug message informs that a class was listed for additional evaluation but its definition does not include a test expression to evaluate. The class is unconditionally added to the query.

26.14.4 DHCP4_ADDITIONAL_CLASS_UNDEFINED

```
additional class %1 has no definition
```

Logged at debug log level 40. This debug message informs that a class is listed for additional evaluation but has no definition. The class is ignored.

26.14.5 DHCP4_ALREADY_RUNNING

```
%1 already running? %2
```

This is an error message that occurs when the DHCPv4 server encounters a pre-existing PID file which contains the PID of a running process. This most likely indicates an attempt to start a second instance of the server using the same configuration file. It is possible, though unlikely that the PID file is a remnant left behind by a server crash or power failure and the PID it contains refers to a process other than the server. In such an event, it would be necessary to manually remove the PID file. The first argument is the DHCPv4 process name, the second contains the PID and PID file.

26.14.6 DHCP4_BUFFER_RECEIVED

```
received buffer from %1:%2 to %3:%4 over interface %5
```

Logged at debug log level 40. This debug message is logged when the server has received a packet over the socket. When the message is logged the contents of the received packet hasn't been parsed yet. The only available information is the interface and the source and destination IPv4 addresses/ports.

26.14.7 DHCP4_BUFFER_RECEIVE_FAIL

```
error on attempt to receive packet: %1
```

The DHCPv4 server tried to receive a packet but an error occurred during this attempt. The reason for the error is included in the message.

26.14.8 DHCP4_BUFFER_UNPACK

```
parsing buffer received from %1 to %2 over interface %3
```

Logged at debug log level 50. This debug message is issued when the server starts parsing the received buffer holding the DHCPv4 message. The arguments specify the source and destination IPv4 addresses as well as the interface over which the buffer has been received.

26.14.9 DHCP4_BUFFER_WAIT_SIGNAL

```
signal received while waiting for next packet
```

Logged at debug log level 50. This debug message is issued when the server was waiting for the packet, but the wait has been interrupted by the signal received by the process. The signal will be handled before the server starts waiting for next packets.

26.14.10 DHCP4_CB_ON_DEMAND_FETCH_UPDATES_FAIL

```
error on demand attempt to fetch configuration updates from the configuration_
↪ backend(s): %1
```

This error message is issued when the server attempted to fetch configuration updates from the database and this on demand attempt failed. The sole argument which is returned to the config-backend-pull command caller too contains the reason for failure.

26.14.11 DHCP4_CB_PERIODIC_FETCH_UPDATES_FAIL

```
error on periodic attempt to fetch configuration updates from the configuration_
↪ backend(s): %1
```

This error message is issued when the server attempted to fetch configuration updates from the database and this periodic attempt failed. The server will re-try according to the configured value of the config-fetch-wait-time parameter. The sole argument contains the reason for failure.

26.14.12 DHCP4_CB_PERIODIC_FETCH_UPDATES_RETRIES_EXHAUSTED

```
maximum number of configuration fetch attempts: 10, has been exhausted without success
```

This error indicates that the server has made a number of unsuccessful periodic attempts to fetch configuration updates from a configuration backend. The server will continue to operate but won't make any further attempts to fetch configuration updates. The administrator must fix the configuration in the database and reload (or restart) the server.

26.14.13 DHCP4_CLASSES_ASSIGNED

```
%1: client packet has been assigned on %2 message to the following classes: %3
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified classes. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument specifies the DHCPv4 message type. The third argument includes all classes to which the packet has been assigned.

26.14.14 DHCP4_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION

```
%1: client packet has been assigned to the following classes: %2
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified classes. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument includes all classes to which the packet has been assigned.

26.14.15 DHCP4_CLASS_ASSIGNED

```
%1: client packet has been assigned to the following class: %2
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified class. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument includes the new class to which the packet has been assigned.

26.14.16 DHCP4_CLASS_UNCONFIGURED

```
%1: client packet belongs to an unconfigured class: %2
```

Logged at debug log level 40. This debug message informs that incoming packet belongs to a class which cannot be found in the configuration. Either a hook written before the classification was added to Kea is used, or class naming is inconsistent.

26.14.17 DHCP4_CLIENTID_IGNORED_FOR_LEASES

```
%1: not using client identifier for lease allocation for subnet %2
```

Logged at debug log level 50. This debug message is issued when the server is processing the DHCPv4 message for which client identifier will not be used when allocating new lease or renewing existing lease. The server is explicitly configured to not use client identifier to lookup existing leases for the client and will not record client identifier in the lease database. This mode of operation is useful when clients don't use stable client identifiers, e.g. multi stage booting. The first argument includes the client and transaction identification information. The second argument specifies the identifier of the subnet where the client is connected and for which this mode of operation is configured on the server.

26.14.18 DHCP4_CLIENT_FQDN_DATA

```
%1: Client sent FQDN option: %2
```

Logged at debug log level 55. This debug message includes the detailed information extracted from the Client FQDN option sent in the query. The first argument includes the client and transaction identification information. The second argument specifies the detailed information about the FQDN option received by the server.

26.14.19 DHCP4_CLIENT_FQDN_PROCESS

```
%1: processing Client FQDN option
```

Logged at debug log level 50. This debug message is issued when the server starts processing the Client FQDN option sent in the client's query. The argument includes the client and transaction identification information.

26.14.20 DHCP4_CLIENT_FQDN_SCRUBBED_EMPTY

```
%1: sanitizing client's FQDN option '%2' yielded an empty string
```

Logged at debug log level 50. This debug message is issued when the result of sanitizing the FQDN option(81) sent by the client is an empty string. When this occurs the server will ignore the FQDN option. The arguments include the client and the FQDN option it sent.

26.14.21 DHCP4_CLIENT_HOSTNAME_DATA

```
%1: client sent Hostname option: %2
```

Logged at debug log level 55. This debug message includes the detailed information extracted from the Hostname option sent in the query. The first argument includes the client and transaction identification information. The second argument specifies the hostname carried in the Hostname option sent by the client.

26.14.22 DHCP4_CLIENT_HOSTNAME_MALFORMED

```
%1: client hostname option malformed: %2
```

Logged at debug log level 50. This debug message is issued when the DHCP server was unable to process the the hostname option sent by the client because the content is malformed. The first argument includes the client and transaction identification information. The second argument contains a description of the data error.

26.14.23 DHCP4_CLIENT_HOSTNAME_PROCESS

```
%1: processing client's Hostname option
```

Logged at debug log level 50. This debug message is issued when the server starts processing the Hostname option sent in the client's query. The argument includes the client and transaction identification information.

26.14.24 DHCP4_CLIENT_HOSTNAME_SCRUBBED_EMPTY

```
%1: sanitizing client's Hostname option '%2' yielded an empty string
```

Logged at debug log level 50. This debug message is issued when the result of sanitizing the hostname option(12) sent by the client is an empty string. When this occurs the server will ignore the hostname option. The arguments include the client and the hostname option it sent.

26.14.25 DHCP4_CLIENT_NAME_PROC_FAIL

```
%1: failed to process the fqdn or hostname sent by a client: %2
```

Logged at debug log level 55. This debug message is issued when the DHCP server was unable to process the FQDN or Hostname option sent by a client. This is likely because the client's name was malformed or due to internal server error. The first argument contains the client and transaction identification information. The second argument holds the detailed description of the error.

26.14.26 DHCP4_CONFIG_COMPLETE

```
DHCPv4 server has completed configuration: %1
```

This is an informational message announcing the successful processing of a new configuration. It is output during server startup, and when an updated configuration is committed by the administrator. Additional information may be provided.

26.14.27 DHCP4_CONFIG_LOAD_FAIL

```
configuration error using file: %1, reason: %2
```

This error message indicates that the DHCPv4 configuration has failed. If this is an initial configuration (during server's startup) the server will fail to start. If this is a dynamic reconfiguration attempt the server will continue to use an old configuration.

26.14.28 DHCP4_CONFIG_PACKET_QUEUE

```
DHCPv4 packet queue info after configuration: %1
```

This informational message is emitted during DHCPv4 server configuration, immediately after configuring the DHCPv4 packet queue. The information shown depends upon the packet queue type selected.

26.14.29 DHCP4_CONFIG_RECEIVED

```
received configuration %1
```

Logged at debug log level 10. A debug message listing the configuration received by the DHCPv4 server. The source of that configuration depends on used configuration backend.

26.14.30 DHCP4_CONFIG_START

```
DHCPv4 server is processing the following configuration: %1
```

Logged at debug log level 10. This is a debug message that is issued every time the server receives a configuration. That happens at start up and also when a server configuration change is committed by the administrator.

26.14.31 DHCP4_CONFIG_SYNTAX_WARNING

```
configuration syntax warning: %1
```

This warning message indicates that the DHCPv4 configuration had a minor syntax error. The error was displayed and the configuration parsing resumed.

26.14.32 DHCP4_CONFIG_UNRECOVERABLE_ERROR

```
DHCPv4 server new configuration failed with an error which cannot be recovered
```

This fatal error message is issued when a new configuration raised an error which cannot be recovered. A correct configuration must be applied as soon as possible as the server is no longer working. The configuration can be fixed in several ways. If the control channel is open, config-set with a valid configuration can be used. Alternatively, the original config file on disk could be fixed and SIGHUP signal could be sent (or the config-reload command issued). Finally, the server could be restarted completely.

26.14.33 DHCP4_CONFIG_UNSUPPORTED_OBJECT

```
DHCPv4 server configuration includes an unsupported object: %1
```

This error message is issued when the configuration includes an unsupported object (i.e. a top level element).

26.14.34 DHCP4_DB_RECONNECT_DISABLED

```
database reconnect is disabled: retries left: %1, reconnect wait time: %2, manager ID: ↪%3, timer: %4
```

This is an informational message indicating that connectivity to either the lease or host database or both and that automatic reconnect is not enabled.

26.14.35 DHCP4_DB_RECONNECT_FAILED

```
maximum number of database reconnect attempts: %1, has been exhausted without success, ↪  
↪manager ID: %2, timer: %3
```

This error indicates that the server failed to reconnect to the lease and/or host database(s) after making the maximum configured number of reconnect attempts. This might cause the server to shut down as specified in the configuration. Loss of connectivity is typically a network or database server issue.

26.14.36 DHCP4_DB_RECONNECT_LOST_CONNECTION

```
database connection lost: manager ID: %1, timer: %2.
```

This info message indicates that the connection has been lost and the dhcp service might have been disabled, as specified in the configuration, in order to try to recover the connection.

26.14.37 DHCP4_DB_RECONNECT_NO_DB_CTL

```
unexpected error in database reconnect
```

This is an error message indicating a programmatic error that should not occur. It prohibits the server from attempting to reconnect to its databases if connectivity is lost, and the server exits. This error should be reported.

26.14.38 DHCP4_DB_RECONNECT_SUCCEEDED

```
database connection recovered: manager ID: %1, timer: %2.
```

This info message indicates that the connection has been recovered and the dhcp service has been restored.

26.14.39 DHCP4_DDNS_REQUEST_SEND_FAILED

```
failed sending a request to kea-dhcp-ddns, error: %1, ncr: %2
```

This error message indicates that DHCP4 server attempted to send a DDNS update request to the DHCP-DDNS server. This is most likely a configuration or networking error.

26.14.40 DHCP4_DECLINE_FAIL

```
%1: error on decline lease for address %2: %3
```

This error message indicates that the software failed to decline a lease from the lease database due to an error during a database operation. The first argument includes the client and the transaction identification information. The second argument holds the IPv4 address which decline was attempted. The last one contains the reason for failure.

26.14.41 DHCP4_DECLINE_LEASE

```
Received DHCPDECLINE for addr %1 from client %2. The lease will be unavailable for %3_↵
↵seconds.
```

This informational message is printed when a client received an address, but discovered that it is being used by some other device and notified the server by sending a DHCPDECLINE message. The server checked that this address really was leased to the client and marked this address as unusable for a certain amount of time. This message may indicate a misconfiguration in a network, as there is either a buggy client or more likely a device that is using an address that it is not supposed to. The server will fully recover from this situation, but if the underlying problem of a misconfigured or rogue device is not solved, this address may be declined again in the future.

26.14.42 DHCP4_DECLINE_LEASE_MISMATCH

```
Received DHCPDECLINE for addr %1 from client %2, but the data doesn't match: received_↵
↵hwaddr: %3, lease hwaddr: %4, received client-id: %5, lease client-id: %6
```


This informational message means that a client attempted to report his address as declined (i.e. used by unknown entity). The server has information about a lease for that address, but the client's hardware address or client identifier does not match the server's stored information. The client's request will be ignored.

26.14.43 DHCP4_DECLINE_LEASE_NOT_FOUND

```
Received DHCPDECLINE for addr %1 from client %2, but no such lease found.
```

This warning message indicates that a client reported that his address was detected as a duplicate (i.e. another device in the network is using this address). However, the server does not have a record for this address. This may indicate a client's error or a server's purged database.

26.14.44 DHCP4_DEFERRED_OPTION_MISSING

```
%1: cannot find deferred option code %2 in the query
```

Logged at debug log level 50. This debug message is printed when a deferred option cannot be found in the query.

26.14.45 DHCP4_DEFERRED_OPTION_UNPACK_FAIL

```
%1: An error unpacking the deferred option %2: %3
```

Logged at debug log level 50. A debug message issued when deferred unpacking of an option failed, making it to be left unpacked in the packet. The first argument is the option code, the second the error.

26.14.46 DHCP4_DEVELOPMENT_VERSION

```
This software is a development branch of Kea. It is not recommended for production use.
```

This warning message is displayed when the version is a development (vs stable) one: the second number of the version is odd.

26.14.47 DHCP4_DHCP4O6_BAD_PACKET

```
%1: received malformed DHCPv4o6 packet: %2
```

Logged at debug log level 50. A malformed DHCPv4o6 packet was received.

26.14.48 DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_DROP

```
%1: packet was dropped, because a callout set the next step to 'drop'
```

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet4_select hook point sets the next step to 'drop' value. For this particular hook point, the setting to that value instructs the server to drop the received packet. The argument specifies the client and transaction identification information.

26.14.49 DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_SKIP

```
%1: no subnet was selected, because a callout set the next skip flag
```

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet4_select hook point sets the next step to SKIP value. For this particular hook point, the setting of the flag instructs the server not to choose a subnet, an action that severely limits further processing; the server will be only able to offer global options - no addresses will be assigned. The argument specifies the client and transaction identification information.

26.14.50 DHCP4_DHCP4O6_PACKET_RECEIVED

```
received DHCPv4o6 packet from DHCPv4 server (type %1) for %2 on interface %3
```

Logged at debug log level 40. This debug message is printed when the server is receiving a DHCPv4o6 from the DHCPv4 server over inter-process communication.

26.14.51 DHCP4_DHCP4O6_PACKET_SEND

```
%1: trying to send packet %2 (type %3) to %4 port %5 on interface %6 encapsulating %7:  
→%8 (type %9)
```

Logged at debug log level 40. The arguments specify the client identification information (HW address and client identifier), DHCPv6 message name and type, source IPv6 address and port, and interface name, DHCPv4 client identification, message name and type.

26.14.52 DHCP4_DHCP4O6_PACKET_SEND_FAIL

```
%1: failed to send DHCPv4o6 packet: %2
```

This error is output if the IPv4 DHCP server fails to send an DHCPv4o6 message to the IPv6 DHCP server. The reason for the error is included in the message.

26.14.53 DHCP4_DHCP4O6_RECEIVE_FAIL

```
failed to receive DHCPv4o6: %1
```

Logged at debug log level 50. This debug message indicates the inter-process communication with the DHCPv6 server failed. The reason for the error is included in the message.

26.14.54 DHCP4_DHCP4O6_RECEIVING

```
receiving DHCPv4o6 packet from DHCPv6 server
```

Logged at debug log level 50. This debug message is printed when the server is receiving a DHCPv4o6 from the DHCPv6 server over inter-process communication socket.

26.14.55 DHCP4_DHCP4O6_RESPONSE_DATA

```
%1: responding with packet %2 (type %3), packet details: %4
```

Logged at debug log level 55. A debug message including the detailed data about the packet being sent to the DHCPv6 server to be forwarded to the client. The first argument contains the client and the transaction identification information. The second and third argument contains the packet name and type respectively. The fourth argument contains detailed packet information.

26.14.56 DHCP4_DHCP4O6_SUBNET_DATA

```
%1: the selected subnet details: %2
```

Logged at debug log level 55. This debug message includes the details of the subnet selected for the client. The first argument includes the client and the transaction identification information. The second arguments includes the subnet details.

26.14.57 DHCP4_DHCP4O6_SUBNET_SELECTED

%1: the subnet with ID %2 was selected for client assignments

Logged at debug log level 45. This is a debug message noting the selection of a subnet to be used for address and option assignment. Subnet selection is one of the early steps in the processing of incoming client message. The first argument includes the client and the transaction identification information. The second argument holds the selected subnet id.

26.14.58 DHCP4_DHCP4O6_SUBNET_SELECTION_FAILED

%1: failed to select subnet for the client

Logged at debug log level 50. This debug message indicates that the server failed to select the subnet for the client which has sent a message to the server. The server will not be able to offer any lease to the client and will drop its message if the received message was DHCPDISCOVER, and will send DHCPNAK if the received message was DHCPREQUEST. The argument includes the client and the transaction identification information.

26.14.59 DHCP4_DISCOVER

%1: server is processing DHCPDISCOVER with hint=%2

Logged at debug log level 50. This is a debug message that indicates the processing of a received DHCPDISCOVER message. The first argument contains the client and the transaction identification information. The second argument may hold the hint for the server about the address that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.14.60 DHCP4_DYNAMIC_RECONFIGURATION

initiate server reconfiguration using file: %1, after receiving SIGHUP signal or config-
↪reload command

This is the info message logged when the DHCPv4 server starts reconfiguration as a result of receiving SIGHUP signal or config-reload command.

26.14.61 DHCP4_DYNAMIC_RECONFIGURATION_FAIL

dynamic server reconfiguration failed with file: %1

This is a fatal error message logged when the dynamic reconfiguration of the DHCP server failed.

26.14.62 DHCP4_DYNAMIC_RECONFIGURATION_SUCCESS

dynamic server reconfiguration succeeded with file: %1

This is info message logged when the dynamic reconfiguration of the DHCP server succeeded.

26.14.63 DHCP4_EMPTY_HOSTNAME

%1: received empty hostname from the client, skipping processing of this option

Logged at debug log level 50. This debug message is issued when the server received an empty Hostname option from a client. Server does not process empty Hostname options and therefore option is skipped. The argument holds the client and transaction identification information.

26.14.64 DHCP4_FLEX_ID

%1: flexible identifier generated for incoming packet: %2

Logged at debug log level 40. This debug message is printed when host reservation type is set to flexible identifier and the expression specified in its configuration generated (was evaluated to) an identifier for incoming packet. This debug message is mainly intended as a debugging assistance for flexible identifier.

26.14.65 DHCP4_GENERATE_FQDN

%1: client did not send a FQDN or hostname; FQDN will be generated for the client

Logged at debug log level 55. This debug message is issued when the server did not receive a Hostname option from the client and hostname generation is enabled. This provides a means to create DNS entries for unsophisticated clients.

26.14.66 DHCP4_HOOK_BUFFER_RCVD_DROP

received buffer from %1 to %2 over interface %3 was dropped because a callout set the `drop` flag

Logged at debug log level 15. This debug message is printed when a callout installed on `buffer4_receive` hook point set the drop flag. For this particular hook point, the setting of the flag by a callout instructs the server to drop the packet. The arguments specify the source and destination IPv4 address as well as the name of the interface over which the buffer has been received.

26.14.67 DHCP4_HOOK_BUFFER_RCVD_SKIP

received buffer from %1 to %2 over interface %3 is not parsed because a callout set the `next step` to SKIP.

Logged at debug log level 50. This debug message is printed when a callout installed on `buffer4_receive` hook point set the next step to SKIP. For this particular hook point, this value set by a callout instructs the server to not parse the buffer because it was already parsed by the hook. The arguments specify the source and destination IPv4 address as well as the name of the interface over which the buffer has been received.

26.14.68 DHCP4_HOOK_BUFFER_SEND_SKIP

%1: prepared response is dropped because a callout set the next step to SKIP.

Logged at debug log level 40. This debug message is printed when a callout installed on `buffer4_send` hook point set the next step to SKIP. For this particular hook point, the SKIP value set by a callout instructs the server to drop the packet. Server completed all the processing (e.g. may have assigned, updated or released leases), but the response will not be sent to the client.

26.14.69 DHCP4_HOOK_DDNS_UPDATE

A hook has updated the DDNS parameters: hostname %1=>%2, forward update %3=>%4, reverse `update` %5=>%6

Logged at debug log level 15. This message indicates that there was a hook called on `ddns4_update` hook point and that hook updated the DDNS update parameters: hostname, or whether to conduct forward (A record) or reverse (PTR record) DDNS updates.

26.14.70 DHCP4_HOOK_DECLINE_SKIP

Decline4 hook callouts set status to DROP, ignoring packet.

Logged at debug log level 15. This message indicates that the server received DHCPDECLINE message, it was verified to be correct and matching server's lease information. The server called hooks for decline4 hook point and one of the callouts set next step status to DROP. The server will now abort processing of the packet as if it was never received. The lease will continue to be assigned to this client.

26.14.71 DHCP4_HOOK_LEASE4_OFFER_ARGUMENT_MISSING

hook callouts did not set an argument as expected %1 for %2

This error message is printed when none of the callouts installed on the lease4_offer hook point set an expected argument in the callout status. This is a programming error in the installed hook libraries. Details of the argument and the query in process at the time are provided log arguments.

26.14.72 DHCP4_HOOK_LEASE4_OFFER_DROP

%1: packet is dropped, because a callout set the next step to DROP

This debug message is printed when a callout installed on the lease4_offer hook point sets the next step to DROP.

26.14.73 DHCP4_HOOK_LEASE4_OFFER_PARK

%1: packet is parked, because a callout set the next step to PARK

This debug message is printed when a callout installed on the lease4_offer hook point sets the next step to PARK.

26.14.74 DHCP4_HOOK_LEASE4_OFFER_PARKING_LOT_FULL

The parked-packet-limit %1, has been reached, dropping query: %2

This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp4 is using hook libraries (e.g. ping-check) that implement the "lease4_offer" callout and client queries are arriving faster than those callouts can fulfill them.

26.14.75 DHCP4_HOOK_LEASE4_RELEASE_SKIP

%1: lease was not released because a callout set the next step to SKIP

Logged at debug log level 15. This debug message is printed when a callout installed on lease4_release hook point set the next step status to SKIP. For this particular hook point, the value set by a callout instructs the server to not release a lease.

26.14.76 DHCP4_HOOK_LEASES4_COMMITTED_DROP

%1: packet is dropped, because a callout set the next step to DROP

This debug message is printed when a callout installed on the leases4_committed hook point sets the next step to DROP.

26.14.77 DHCP4_HOOK_LEASES4_COMMITTED_PARK

%1: packet is parked, because a callout set the next step to PARK

This debug message is printed when a callout installed on the leases4_committed hook point sets the next step to PARK.

26.14.78 DHCP4_HOOK_LEASES4_COMMITTED_PARKING_LOT_FULL

The parked-packet-limit %1, has been reached, dropping query: %2

This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp4 is using hook libraries (e.g. HA) that implement the "leases4-committed" callout and client queries are arriving faster than those callouts can fulfill them.

26.14.79 DHCP4_HOOK_PACKET_RCVD_SKIP

%1: packet is dropped, because a callout set the next step to SKIP

Logged at debug log level 40. This debug message is printed when a callout installed on the pkt4_receive hook point sets the next step to SKIP. For this particular hook point, the value setting of the flag instructs the server to drop the packet.

26.14.80 DHCP4_HOOK_PACKET_SEND_DROP

%1: prepared DHCPv4 response was not sent because a callout set the next step to DROP

Logged at debug log level 15. This debug message is printed when a callout installed on the pkt4_send hook point set the next step to DROP. For this particular hook point, the setting of the value by a callout instructs the server to drop the packet. This effectively means that the client will not get any response, even though the server processed client's request and acted on it (e.g. possibly allocated a lease). The argument specifies the client and transaction identification information.

26.14.81 DHCP4_HOOK_PACKET_SEND_SKIP

%1: prepared response is not sent, because a callout set the next step to SKIP

Logged at debug log level 40. This debug message is printed when a callout installed on the pkt4_send hook point sets the next step to SKIP. For this particular hook point, this setting instructs the server to drop the packet. This means that the client will not get any response, even though the server processed client's request and acted on it (e.g. possibly allocated a lease).

26.14.82 DHCP4_HOOK_SUBNET4_SELECT_4O6_PARKING_LOT_FULL

The parked-packet-limit %1, has been reached, dropping query: %2

Logged at debug log level 15. This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp4 is using hook libraries (e.g. radius) that implement the "subnet4_select" callout and DHCP4O6 client queries are arriving faster than those callouts can fulfill them.

26.14.83 DHCP4_HOOK_SUBNET4_SELECT_DROP

`%1: packet was dropped, because a callout set the next step to 'drop'`

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet4_select hook point sets the next step to 'drop' value. For this particular hook point, the setting to that value instructs the server to drop the received packet. The argument specifies the client and transaction identification information.

26.14.84 DHCP4_HOOK_SUBNET4_SELECT_PARK

`%1: packet was parked`

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet4_select hook point set the park flag. The argument holds the client and transaction identification information.

26.14.85 DHCP4_HOOK_SUBNET4_SELECT_PARKING_LOT_FULL

`The parked-packet-limit %1, has been reached, dropping query: %2`

Logged at debug log level 15. This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp4 is using hook libraries (e.g. radius) that implement the "subnet4_select" callout and client queries are arriving faster than those callouts can fulfill them.

26.14.86 DHCP4_HOOK_SUBNET4_SELECT_SKIP

`%1: no subnet was selected, because a callout set the next skip flag`

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet4_select hook point sets the next step to SKIP value. For this particular hook point, the setting of the flag instructs the server not to choose a subnet, an action that severely limits further processing; the server will be only able to offer global options - no addresses will be assigned. The argument specifies the client and transaction identification information.

26.14.87 DHCP4_HOOK_SUBNET6_SELECT_PARKING_LOT_FULL

`The parked-packet-limit %1, has been reached, dropping query: %2`

Logged at debug log level 15. This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp4 is using hook libraries (e.g. radius) that implement the "subnet6_select" callout and client queries are arriving faster than those callouts can fulfill them.

26.14.88 DHCP4_INFORM_DIRECT_REPLY

`%1: DHCPACK in reply to the DHCPINFORM will be sent directly to %2 over %3`

Logged at debug log level 50. This debug message is issued when the DHCPACK will be sent directly to the client, rather than via a relay. The first argument contains the client and transaction identification information. The second argument contains the client's IPv4 address to which the response will be sent. The third argument contains the local interface name.

26.14.89 DHCP4_INIT_FAIL

```
failed to initialize Kea server: %1
```

The server has failed to initialize. This may be because the configuration was not successful, or it encountered any other critical error on startup. Attached error message provides more details about the issue.

26.14.90 DHCP4_INIT_REBOOT

```
%1: client is in INIT-REBOOT state and requests address %2
```

This informational message is issued when the client is in the INIT-REBOOT state and is requesting an IPv4 address it is using to be allocated for it. The first argument includes the client and transaction identification information. The second argument specifies the requested IPv4 address.

26.14.91 DHCP4_LEASE_ALLOC

```
%1: lease %2 has been allocated for %3 seconds
```

This informational message indicates that the server successfully granted a lease in response to client's DHCPREQUEST message. The lease information will be sent to the client in the DHCPACK message. The first argument contains the client and the transaction identification information. The second argument contains the allocated IPv4 address. The third argument is the validity lifetime.

26.14.92 DHCP4_LEASE_OFFER

```
%1: lease %2 will be offered
```

This informational message indicates that the server has found the lease to be offered to the client. It is up to the client to choose one server out of those which offered leases and continue allocation with that server. The first argument specifies the client and the transaction identification information. The second argument specifies the IPv4 address to be offered.

26.14.93 DHCP4_LEASE_QUERY_PACKET_PACK_FAILED

```
preparing on-wire-format of the packet to be sent: %1, failed %2
```

This error message is issued when preparing an on-wire format of the packet has failed. The first argument provides packet details. the second the second explains the nature of the error.

26.14.94 DHCP4_LEASE_QUERY_PACKET_UNPACK_FAILED

```
failed to parse query from %1 to %2, received over interface %3, reason: %4
```

Logged at debug log level 40. This debug message is issued when received DHCPLEASEQUERY is malformed and can't be parsed by the buffer4_receive callout. The query will be dropped by the server. The first three arguments specify source IP address, destination IP address and the interface. The last argument provides a reason for failure.

26.14.95 DHCP4_LEASE_QUERY_PROCESS_FAILED

```
processing failed for lease query: %1, reason: %2
```

Logged at debug log level 40. This error message is issued when the server encountered an error processing a DHCPLEASEQUERY. The first argument provides query details, the second an explanation of the error.

26.14.96 DHCP4_LEASE_QUERY_RECEIVED

```
received query: %1
```

Logged at debug log level 40. This debug message is printed when a DHCPLEASEQUERY query has been received. The argument provides query details.

26.14.97 DHCP4_LEASE_QUERY_RESPONSE_SENT

```
response: %1, sent to %2:%3
```

Logged at debug log level 40. This debug message is printed when a response to a DHCPLEASEQUERY has been sent to a requester. The first argument provides response details, the second and third arguments are the IP address and port to which the response was sent.

26.14.98 DHCP4_LEASE_QUERY_SEND_FAILED

```
unable to send response: %1, iface: %2, address %3:%4 error: %5
```

This error message is issued when the server was unable to send a lease query response back to the requester. a DHCPLEASEQUERY. The first argument provides query details, followed by the output interface, IP address and port, and finally the error itself.

26.14.99 DHCP4_LEASE_REUSE

```
%1: lease %2 has been reused for %3 seconds
```

This informational message indicates that the server successfully reused a lease in response to client's message. The lease information will be sent to the client in the DHCPACK message. The first argument contains the client and the transaction identification information. The second argument contains the allocated IPv4 address. The third argument is the validity lifetime.

26.14.100 DHCP4_MULTI_THREADING_INFO

```
enabled: %1, number of threads: %2, queue size: %3
```

This is a message listing some information about the multi-threading parameters with which the server is running.

26.14.101 DHCP4_NCR_CREATION_FAILED

```
%1: failed to generate name change requests for DNS: %2
```

This message indicates that server was unable to generate NameChangeRequests which should be sent to the kea-dhcp_ddns module to create new DNS records for the lease being acquired or to update existing records for the renewed lease. The first argument contains the client and transaction identification information. The second argument includes the reason for the failure.

26.14.102 DHCP4_NOT_RUNNING

```
DHCPv4 server is not running
```

A warning message is issued when an attempt is made to shut down the DHCPv4 server but it is not running.

26.14.103 DHCP4_NO_LEASE_INIT_REBOOT

%1: no lease for address %2 requested by INIT-REBOOT client

Logged at debug log level 50. This debug message is issued when the client being in the INIT-REBOOT state requested an IPv4 address but this client is unknown. The server will not respond. The first argument includes the client and the transaction id identification information. The second argument includes the IPv4 address requested by the client.

26.14.104 DHCP4_OPEN_SOCKET

opening service sockets on port %1

Logged at debug log level 0. A debug message issued during startup, this indicates that the DHCPv4 server is about to open sockets on the specified port.

26.14.105 DHCP4_OPEN_SOCKETS_FAILED

maximum number of open service sockets attempts: %1, has been exhausted without success

This error indicates that the server failed to bind service sockets after making the maximum configured number of reconnect attempts. This might cause the server to shut down as specified in the configuration.

26.14.106 DHCP4_OPEN_SOCKETS_NO_RECONNECT_CTL

unexpected error in bind service sockets.

This is an error message indicating a programmatic error that should not occur. It prohibits the server from attempting to bind to its service sockets if they are unavailable, and the server exits. This error should be reported.

26.14.107 DHCP4_PACKET_DROP_0001

%1: failed to parse packet from %2 to %3, received over interface %4, reason: %5, %6

Logged at debug log level 15. The DHCPv4 server has received a packet that it is unable to interpret. The reason why the packet is invalid is included in the message.

26.14.108 DHCP4_PACKET_DROP_0002

%1, from interface %2: no suitable subnet configured for a direct client

Logged at debug log level 15. This info message is logged when received a message from a directly connected client but there is no suitable subnet configured for the interface on which this message has been received. The IPv4 address assigned on this interface must belong to one of the configured subnets. Otherwise received message is dropped.

26.14.109 DHCP4_PACKET_DROP_0003

%1, from interface %2: it contains a foreign server identifier

Logged at debug log level 15. This debug message is issued when received DHCPv4 message is dropped because it is addressed to a different server, i.e. a server identifier held by this message doesn't match the identifier used by our server. The arguments of this message hold the name of the transaction id and interface on which the message has been received.

26.14.110 DHCP4_PACKET_DROP_0004

%1, from interface %2: missing msg-type option

Logged at debug log level 15. This is a debug message informing that incoming DHCPv4 packet did not have mandatory DHCP message type option and thus was dropped. The arguments specify the client and transaction identification information, as well as the interface on which the message has been received.

26.14.111 DHCP4_PACKET_DROP_0005

%1: unrecognized type %2 in option 53

Logged at debug log level 15. This debug message indicates that the message type carried in DHCPv4 option 53 is unrecognized by the server. The valid message types are listed on the IANA website: <http://www.iana.org/assignments/bootp-dhcp-parameters/bootp-dhcp-parameters.xhtml#message-type-53>. The message will not be processed by the server. The arguments specify the client and transaction identification information, as well as the received message type.

26.14.112 DHCP4_PACKET_DROP_0006

%1: unsupported DHCPv4 message type %2

Logged at debug log level 15. This debug message indicates that the message type carried in DHCPv4 option 53 is valid but the message will not be processed by the server. This includes messages being normally sent by the server to the client, such as DHCPOFFER, DHCPACK, DHCPNAK etc. The first argument specifies the client and transaction identification information. The second argument specifies the message type.

26.14.113 DHCP4_PACKET_DROP_0007

%1: failed to process packet: %2

Logged at debug log level 15. This is a general catch-all message indicating that the processing of a received packet failed. The reason is given in the message. The server will not send a response but will instead ignore the packet. The first argument contains the client and transaction identification information. The second argument includes the details of the error.

26.14.114 DHCP4_PACKET_DROP_0008

%1: DHCP service is globally disabled

Logged at debug log level 15. This debug message is issued when a packet is dropped because the DHCP service has been temporarily disabled. This affects all received DHCP packets. The service may be enabled by the "dhcp-enable" control command or automatically after a specified amount of time since receiving "dhcp-disable" command.

26.14.115 DHCP4_PACKET_DROP_0009

%1: Option 53 missing (no DHCP message type), is this a BOOTP packet?

Logged at debug log level 15. This debug message is issued when a packet is dropped because it did contain option 53 and thus has no DHCP message type. The most likely explanation is that it was BOOTP packet.

26.14.116 DHCP4_PACKET_DROP_0010

dropped as member of the special class 'DROP': %1, %2

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified into the special class 'DROP' and dropped. The packet details are displayed.

26.14.117 DHCP4_PACKET_DROP_0011

dropped as sent by the same client than a packet being processed by another thread:
↳ dropped %1, %2 by thread %3 as duplicate of %4, %5 processed by thread %6

Logged at debug log level 15. Currently multi-threading processing avoids races between packets sent by a client using the same client id option by dropping new packets until processing is finished. Packet details and thread identifiers are included for both packets in this warning message.

26.14.118 DHCP4_PACKET_DROP_0012

dropped as sent by the same client than a packet being processed by another thread:
↳ dropped %1, %2 by thread %3 as duplicate of %4, %5 processed by thread %6

Logged at debug log level 15. Currently multi-threading processing avoids races between packets sent by a client using the same hardware address by dropping new packets until processing is finished. Packet details and thread identifiers are included for both packets in this warning message.

26.14.119 DHCP4_PACKET_DROP_0013

dropped as member of the special class 'DROP' after host reservation lookup: %1, %2

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified after host reservation lookup into the special class 'DROP' and dropped. The packet details are displayed.

26.14.120 DHCP4_PACKET_DROP_0014

dropped as member of the special class 'DROP' after early global host reservations
↳ lookup: %1, %2

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified after early global host reservations lookup into the special class 'DROP' and dropped. The packet details are displayed.

26.14.121 DHCP4_PACKET_NAK_0001

%1: failed to select a subnet for incoming packet, src %2, type %3

This error message is output when a packet was received from a subnet for which the DHCPv4 server has not been configured. The most probable cause is a misconfiguration of the server. The first argument contains the client and transaction identification information. The second argument contains the source IPv4 address of the packet. The third argument contains the name of the received packet.

26.14.122 DHCP4_PACKET_NAK_0002

```
%1: invalid address %2 requested by INIT-REBOOT
```

Logged at debug log level 50. This debug message is issued when the client being in the INIT-REBOOT state requested an IPv4 address which is not assigned to him. The server will respond to this client with DHCPNAK. The first argument contains the client and the transaction identification information. The second arguments holds the IPv4 address requested by the client.

26.14.123 DHCP4_PACKET_NAK_0003

```
%1: failed to advertise a lease, client sent ciaddr %2, requested-ip-address %3
```

Logged at debug log level 50. This message indicates that the server has failed to offer a lease to the specified client after receiving a DISCOVER message from it. There are many possible reasons for such a failure. The first argument contains the client and the transaction identification information. The second argument contains the IPv4 address in the ciaddr field. The third argument contains the IPv4 address in the requested-ip-address option (if present).

26.14.124 DHCP4_PACKET_NAK_0004

```
%1: failed to grant a lease, client sent ciaddr %2, requested-ip-address %3
```

Logged at debug log level 50. This message indicates that the server failed to grant a lease to the specified client after receiving a REQUEST message from it. There are many possible reasons for such a failure. Additional messages will indicate the reason. The first argument contains the client and the transaction identification information. The second argument contains the IPv4 address in the ciaddr field. The third argument contains the IPv4 address in the requested-ip-address option (if present).

26.14.125 DHCP4_PACKET_OPTIONS_SKIPPED

```
%1: An error unpacking an option, caused subsequent options to be skipped: %2
```

Logged at debug log level 50. A debug message issued when an option failed to unpack correctly, making it impossible to unpack the remaining options in the packet. The server will server will still attempt to service the packet.

26.14.126 DHCP4_PACKET_PACK

```
%1: preparing on-wire format of the packet to be sent
```

Logged at debug log level 50. This debug message is issued when the server starts preparing the on-wire format of the packet to be sent back to the client. The argument specifies the client and the transaction identification information.

26.14.127 DHCP4_PACKET_PACK_FAIL

```
%1: preparing on-wire-format of the packet to be sent failed %2
```

This error message is issued when preparing an on-wire format of the packet has failed. The first argument identifies the client and the DHCP transaction. The second argument includes the error string.

26.14.128 DHCP4_PACKET_PROCESS_EXCEPTION

%1: exception occurred during packet processing

This error message indicates that a non-standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation.

26.14.129 DHCP4_PACKET_PROCESS_EXCEPTION_MAIN

exception occurred during packet processing

This error message indicates that a non-standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation. This error message may appear in main server processing loop.

26.14.130 DHCP4_PACKET_PROCESS_STD_EXCEPTION

%1: exception occurred during packet processing: %2

This error message indicates that a standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation.

26.14.131 DHCP4_PACKET_PROCESS_STD_EXCEPTION_MAIN

exception occurred during packet processing: %1

This error message indicates that a standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation. This error message may appear in main server processing loop.

26.14.132 DHCP4_PACKET_QUEUE_FULL

multi-threading packet queue is full

Logged at debug log level 40. A debug message noting that the multi-threading packet queue is full so the oldest packet of the queue was dropped to make room for the received one.

26.14.133 DHCP4_PACKET_RECEIVED

%1: %2 (type %3) received from %4 to %5 on interface %6

An INFO message noting that the server has received the specified type of packet on the specified interface. The first argument specifies the client and transaction identification information. The second and third argument specify the name of the DHCPv4 message and its numeric type respectively. The remaining arguments specify the source IPv4 address, destination IPv4 address and the name of the interface on which the message has been received.

26.14.134 DHCP4_PACKET_SEND

%1: trying to send packet %2 (type %3) from %4:%5 to %6:%7 on interface %8

An INFO message noting that the server is attempting to send the specified type of packet. The arguments specify the client identification information (HW address and client identifier), DHCP message name and type, source IPv4 address and port, destination IPv4 address and port and the interface name. This debug message is issued when the

server is trying to send the response to the client. When the server is using an UDP socket to send the packet there are cases when this operation may be unsuccessful and no error message will be displayed. One such situation occurs when the server is unicasting the response to the 'ciaddr' of a DHCPINFORM message. This often requires broadcasting an ARP message to obtain the link layer address of the unicast destination. If broadcast ARP messages are blocked in the network, according to the firewall policy, the ARP message will not cause a response. Consequently, the response to the DHCPINFORM will not be sent. Since the ARP communication is under the OS control, Kea is not notified about the drop of the packet which it is trying to send and it has no means to display an error message.

26.14.135 DHCP4_PACKET_SEND_FAIL

```
%1: failed to send DHCPv4 packet: %2
```

This error is output if the DHCPv4 server fails to send an assembled DHCP message to a client. The first argument includes the client and the transaction identification information. The second argument includes the reason for failure.

26.14.136 DHCP4_PARSER_COMMIT_EXCEPTION

```
parser failed to commit changes
```

On receipt of message containing details to a change of the DHCPv4 server configuration, a set of parsers were successfully created, but one of them failed to commit its changes due to a low-level system exception being raised. Additional messages may be output indicating the reason.

26.14.137 DHCP4_PARSER_COMMIT_FAIL

```
parser failed to commit changes: %1
```

On receipt of message containing details to a change of the DHCPv4 server configuration, a set of parsers were successfully created, but one of them failed to commit its changes. The reason for the failure is given in the message.

26.14.138 DHCP4_PARSER_EXCEPTION

```
failed to create or run parser for configuration element %1
```

On receipt of message containing details to a change of its configuration, the DHCPv4 server failed to create a parser to decode the contents of the named configuration element, or the creation succeeded but the parsing actions and committal of changes failed. The message has been output in response to a non-Kea exception being raised. Additional messages may give further information.

26.14.139 DHCP4_PARSER_FAIL

```
failed to create or run parser for configuration element %1: %2
```

On receipt of message containing details to a change of its configuration, the DHCPv4 server failed to create a parser to decode the contents of the named configuration element, or the creation succeeded but the parsing actions and committal of changes failed. The reason for the failure is given in the message.

26.14.140 DHCP4_POST_ALLOCATION_NAME_UPDATE_FAIL

```
%1: failed to update hostname %2 in a lease after address allocation: %3
```

This message indicates the failure when trying to update the lease and/or options in the server's response with the hostname generated by the server or reserved for the client belonging to a shared network. The latter is the case when the server dynamically switches to another subnet (than initially selected for allocation) from the same shared network.

26.14.141 DHCP4_QUERY_DATA

```
%1, packet details: %2
```

Logged at debug log level 55. A debug message printing the details of the received packet. The first argument includes the client and the transaction identification information.

26.14.142 DHCP4_QUERY_LABEL

```
received query: %1
```

This information message indicates that a query was received. It displays the client and the transaction identification information.

26.14.143 DHCP4_RECLAIM_EXPIRED_LEASES_FAIL

```
failed to reclaim expired leases: %1
```

This error message indicates that the reclaim expired leases operation failed and provides the cause of failure.

26.14.144 DHCP4_RECLAIM_EXPIRED_LEASES_SKIPPED

```
dhcp6 service is currently disabled. Try again in %1 seconds.
```

Logged at debug log level 40. This debug message is emitted when lease reclamation was scheduled to begin but skipped because DHCPv6 service was disabled. Reclamation will continue to be scheduled according to the configured value of reclaim-timer-wait-time.

26.14.145 DHCP4_RECOVERED_STASHED_RELAY_AGENT_INFO

```
recovered for query %1 relay agent option from lease %2: %3
```

Logged at debug log level 55. This debug message indicates that agent options were stashed in the lease for the client address of the request and were recovered. The first argument includes the request information, the second the client address and the last argument the content of the dhcp-agent-options option.

26.14.146 DHCP4_RELEASE

```
%1: address %2 was released properly.
```

Logged at debug log level 50. This informational message indicates that an address was released properly. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second argument includes the released IPv4 address.

26.14.147 DHCP4_RELEASE_DELETED

```
%1: address %2 was deleted on release.
```

This informational message indicates that an address was deleted on release. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second argument includes the released IPv4 address.

26.14.148 DHCP4_RELEASE_EXCEPTION

```
%1: while trying to release address %2 an exception occurred: %3
```

This message is output when an error was encountered during an attempt to process a DHCPRELEASE message. The error will not affect the client, which does not expect any response from the server for DHCPRELEASE messages. Depending on the nature of problem, it may affect future server operation. The first argument includes the client and the transaction identification information. The second argument includes the IPv4 address which release was attempted. The last argument includes the detailed error description.

26.14.149 DHCP4_RELEASE_EXPIRED

```
%1: address %2 expired on release.
```

This informational message indicates that an address expired on release. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second argument includes the released IPv4 address.

26.14.150 DHCP4_RELEASE_FAIL

```
%1: failed to remove lease for address %2
```

Logged at debug log level 50. This error message indicates that the software failed to remove a lease from the lease database. It is probably due to an error during a database operation: resolution will most likely require administrator intervention (e.g. check if DHCP process has sufficient privileges to update the database). It may also be triggered if a lease was manually removed from the database during RELEASE message processing. The first argument includes the client and the transaction identification information. The second argument holds the IPv4 address which release was attempted.

26.14.151 DHCP4_RELEASE_FAIL_NO_LEASE

```
%1: client is trying to release non-existing lease %2
```

Logged at debug log level 50. This debug message is printed when client attempts to release a lease, but no such lease is known to the server. The first argument contains the client and transaction identification information. The second argument contains the IPv4 address which the client is trying to release.

26.14.152 DHCP4_RELEASE_FAIL_WRONG_CLIENT

```
%1: client is trying to release the lease %2 which belongs to a different client
```

Logged at debug log level 50. This debug message is issued when a client is trying to release the lease for the address which is currently used by another client, i.e. the 'client identifier' or 'chaddr' doesn't match between the client and the lease. The first argument includes the client and the transaction identification information. The second argument specifies the leased address.

26.14.153 DHCP4_REQUEST

```
%1: server is processing DHCPREQUEST with hint=%2
```

Logged at debug log level 50. This is a debug message that indicates the processing of a received DHCPREQUEST message. The first argument contains the client and the transaction identification information. The second argument may hold the hint for the server about the address that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.14.154 DHCP4_RESERVATIONS_LOOKUP_FIRST_ENABLED

Multi-threading is enabled and host reservations lookup is always performed first.

This is a message informing that host reservations lookup is performed before lease lookup when multi-threading is enabled overwriting configured value.

26.14.155 DHCP4_RESERVED_HOSTNAME_ASSIGNED

%1: server assigned reserved hostname %2

Logged at debug log level 55. This debug message is issued when the server found a hostname reservation for a client and uses this reservation in a hostname option sent back to this client. The reserved hostname is qualified with a value of 'ddns-qualifying-suffix' parameter, if this parameter is specified.

26.14.156 DHCP4_RESPONSE_DATA

%1: responding with packet %2 (type %3), packet details: %4

Logged at debug log level 55. A debug message including the detailed data about the packet being sent to the client. The first argument contains the client and the transaction identification information. The second and third argument contains the packet name and type respectively. The fourth argument contains detailed packet information.

26.14.157 DHCP4_RESPONSE_FQDN_DATA

%1: including FQDN option in the server's response: %2

Logged at debug log level 55. This debug message is issued when the server is adding the Client FQDN option in its response to the client. The first argument includes the client and transaction identification information. The second argument includes the details of the FQDN option being included. Note that the name carried in the FQDN option may be modified by the server when the lease is acquired for the client.

26.14.158 DHCP4_RESPONSE_HOSTNAME_DATA

%1: including Hostname option in the server's response: %2

Logged at debug log level 55. This debug message is issued when the server is adding the Hostname option in its response to the client. The first argument includes the client and transaction identification information. The second argument includes the details of the FQDN option being included. Note that the name carried in the Hostname option may be modified by the server when the lease is acquired for the client.

26.14.159 DHCP4_RESPONSE_HOSTNAME_GENERATE

%1: server has generated hostname %2 for the client

Logged at debug log level 50. This debug message includes the auto-generated hostname which will be used for the client which message is processed. Hostnames may need to be generated when required by the server's configuration or when the client hasn't supplied its hostname. The first argument includes the client and the transaction identification information. The second argument holds the generated hostname.

26.14.160 DHCP4_SECURITY_CHECKS_DISABLED

```
Invoked with command line option -X, Security checks are disabled!!
```

This warning is emitted when internal security checks normally performed by kea-dhcp4 have been disabled via command line option '-X'. This means the server is not enforcing restrictions on resource paths or permissions. This mode of operation may expose your environment to security vulnerabilities and should only be used after careful consideration.

26.14.161 DHCP4_SERVER_FAILED

```
server failed: %1
```

The DHCPv4 server has encountered a fatal error and is terminating. The reason for the failure is included in the message.

26.14.162 DHCP4_SERVER_INITIATED_DECLINE

```
%1: Lease for addr %2 has been found to be already in use. The lease will be unavailable_↵  
↵for %3 seconds.
```

This informational message is printed when the server has detected via ICMP ECHO (i.e. ping check) or other means that a lease which should be free to offer is actually in use. This message may indicate a misconfiguration in a network or more likely a device that is using an address that it is not supposed to use. The server will fully recover from this situation, but if the underlying problem of a misconfigured or rogue device is not solved, this address may be declined again in the future.

26.14.163 DHCP4_SERVER_INITIATED_DECLINE_ADD_FAILED

```
%1: error adding a lease for address %2
```

This error message indicates that the server failed to add a DECLINED lease to the lease store. The first argument includes the client and the transaction identification information. The second argument holds the IPv4 address for which the decline was attempted.

26.14.164 DHCP4_SERVER_INITIATED_DECLINE_RESOURCE_BUSY

```
%1: error declining a lease for address %2
```

This error message indicates that while one server thread was attempting to mark a lease as DECLINED, it was already locked by another thread. The first argument includes the client and the transaction identification information. The second argument holds the IPv4 address for which the decline was attempted.

26.14.165 DHCP4_SERVER_INITIATED_DECLINE_UPDATE_FAILED

```
%1: error updating lease for address %2
```

This error message indicates that the server failed to update a lease in the lease store to the DECLINED state. The first argument includes the client and the transaction identification information. The second argument holds the IPv4 address for which the decline was attempted.

26.14.166 DHCP4_SHUTDOWN

```
server shutdown
```

Logged at debug log level 40. The DHCPv4 server has terminated normally.

26.14.167 DHCP4_SHUTDOWN_REQUEST

```
shutdown of server requested
```

Logged at debug log level 40. This debug message indicates that a shutdown of the DHCPv4 server has been requested via a call to the 'shutdown' method of the core Dhcpv4Srv object.

26.14.168 DHCP4_SRV_CONSTRUCT_ERROR

```
error creating Dhcpv4Srv object, reason: %1
```

This error message indicates that during startup, the construction of a core component within the DHCPv4 server (the Dhcpv4 server object) has failed. As a result, the server will exit. The reason for the failure is given within the message.

26.14.169 DHCP4_SRV_D2STOP_ERROR

```
error stopping IO with DHCP_DDNS during shutdown: %1
```

This error message indicates that during shutdown, an error occurred while stopping IO between the DHCPv4 server and the DHCP_DDNS server. This is probably due to a programmatic error is not likely to impact either server upon restart. The reason for the failure is given within the message.

26.14.170 DHCP4_SRV_DHCP4O6_ERROR

```
error stopping IO with DHCPv4o6 during shutdown: %1
```

This error message indicates that during shutdown, an error occurred while stopping IO between the DHCPv4 server and the DHCPv4o6 server. This is probably due to a programmatic error is not likely to impact either server upon restart. The reason for the failure is given within the message.

26.14.171 DHCP4_SRV_UNLOAD_LIBRARIES_ERROR

```
error unloading hooks libraries during shutdown: %1
```

This error message indicates that during shutdown, unloading hooks libraries failed to close them. If the list of libraries is empty it is a programmatic error in the server code. If it is not empty it could be a programmatic error in one of the hooks libraries which could lead to a crash during finalization.

26.14.172 DHCP4_STARTED

```
Kea DHCPv4 server version %1 started
```

This informational message indicates that the DHCPv4 server has processed all configuration information and is ready to process DHCPv4 packets. The version is also printed.

26.14.173 DHCP4_STARTING

```
Kea DHCPv4 server version %1 (%2) starting
```

This informational message indicates that the DHCPv4 server has processed any command-line switches and is starting. The version is also printed.

26.14.174 DHCP4_START_INFO

```
pid: %1, server port: %2, client port: %3, verbose: %4
```

Logged at debug log level 0. This is a debug message issued during the DHCPv4 server startup. It lists some information about the parameters with which the server is running.

26.14.175 DHCP4_SUBNET_DATA

```
%1: the selected subnet details: %2
```

Logged at debug log level 55. This debug message includes the details of the subnet selected for the client. The first argument includes the client and the transaction identification information. The second arguments includes the subnet details.

26.14.176 DHCP4_SUBNET_DYNAMICALLY_CHANGED

```
%1: changed selected subnet %2 to subnet %3 from shared network %4 for client assignments
```

Logged at debug log level 45. This debug message indicates that the server is using another subnet than initially selected for client assignments. This newly selected subnet belongs to the same shared network as the original subnet. Some reasons why the new subnet was selected include: address pool exhaustion in the original subnet or the fact that the new subnet includes some static reservations for this client.

26.14.177 DHCP4_SUBNET_SELECTED

```
%1: the subnet with ID %2 was selected for client assignments
```

Logged at debug log level 45. This is a debug message noting the selection of a subnet to be used for address and option assignment. Subnet selection is one of the early steps in the processing of incoming client message. The first argument includes the client and the transaction identification information. The second argument holds the selected subnet id.

26.14.178 DHCP4_SUBNET_SELECTION_FAILED

```
%1: failed to select subnet for the client
```

Logged at debug log level 50. This debug message indicates that the server failed to select the subnet for the client which has sent a message to the server. The server will not be able to offer any lease to the client and will drop its message if the received message was DHCPDISCOVER, and will send DHCPNAK if the received message was DHCPREQUEST. The argument includes the client and the transaction identification information.

26.14.179 DHCP4_TESTING_MODE_SEND_TO_SOURCE_ENABLED

```
All packets will be send to source address of an incoming packet - use only for testing
```

This message is printed then `KEA_TEST_SEND_RESPONSES_TO_SOURCE` environment variable is set. It's causing Kea to send packets to source address of incoming packet. Usable just in testing environment to simulate multiple subnet traffic from single source.

26.14.180 DHCP4_UNKNOWN_ADDRESS_REQUESTED

```
%1: client requested an unknown address, client sent ciaddr %2, requested-ip-address %3
```

Logged at debug log level 50. This message indicates that the client requested an address that does not belong to any dynamic pools managed by this server. The first argument contains the client and the transaction identification information. The second argument contains the IPv4 address in the `ciaddr` field. The third argument contains the IPv4 address in the `requested-ip-address` option (if present).

26.14.181 DHCP4_V6_ONLY_PREFERRED_MISSING_IN_ACK

```
v6-only-preferred option missing in 0.0.0.0 reply to query: %1
```

An DHCPACK for the 0.0.0.0 address was generated for a client requesting the v6-only-preferred (108) option but the option is not in the response as expected: the erroneous response is dropped, the request query is displayed.

26.14.182 DHCP4_V6_ONLY_PREFERRED_MISSING_IN_OFFER

```
v6-only-preferred option missing in 0.0.0.0 offer to query: %1
```

An DHCP OFFER for the 0.0.0.0 address was generated for a client requesting the v6-only-preferred (108) option but the option is not in the response as expected: the erroneous response is dropped, the discover query is displayed.

26.15 DHCP6

26.15.1 DHCP6_ADDITIONAL_CLASS_EVAL_ERROR

```
%1: Expression '%2' evaluated to %3
```

This error message indicates that a problem was encountered while evaluating the expression of an additional client class. A description of the problem is printed.

26.15.2 DHCP6_ADDITIONAL_CLASS_EVAL_RESULT

```
%1: Expression '%2' evaluated to %3
```

Logged at debug log level 50. This debug message indicates that the expression of an additional client class has been successfully evaluated. The client class name and the result value of the evaluation are printed.

26.15.3 DHCP6_ADDITIONAL_CLASS_NO_TEST

```
additional class %1 has no test expression, adding it to client's classes unconditionally
```

Logged at debug log level 40. This debug message informs that a class was listed for additional evaluation but its definition does not include a test expression to evaluate. The class is unconditionally added to the query.

26.15.4 DHCP6_ADDITIONAL_CLASS_UNDEFINED

```
additional class %1 has no definition
```

Logged at debug log level 40. This debug message informs that a class is listed for additional evaluation but has no definition. The class is ignored.

26.15.5 DHCP6_ADDR_REG_INFORM_CLIENT_CHANGE

```
received an ADDR-REG-INFORM for %1 from client '%2' but the address was registered by  
↪ another client '%3'
```

This information message is issued when a lease for another client already exists for an address being registered. The address, the new client and previous client identifiers are printed.

26.15.6 DHCP6_ADDR_REG_INFORM_FAIL

```
error on ADDR-REG-INFORM from client %1: %2
```

This information message is issued when the processing of an ADDR-REG-INFORM message failed. The address of the client, usually also the address to register, and the description of the problem are printed.

26.15.7 DHCP6_ADD_GLOBAL_STATUS_CODE

```
%1: adding Status Code to DHCPv6 packet: %2
```

Logged at debug log level 50. This message is logged when the server is adding the top-level Status Code option. The first argument includes the client and the transaction identification information. The second argument includes the details of the status code.

26.15.8 DHCP6_ADD_STATUS_CODE_FOR_IA

```
%1: adding Status Code to IA with iaid=%2: %3
```

Logged at debug log level 50. This message is logged when the server is adding the Status Code option to an IA. The first argument includes the client and the transaction identification information. The second argument specifies the IAID. The third argument includes the details of the status code.

26.15.9 DHCP6_ALREADY_RUNNING

```
%1 already running? %2
```

This is an error message that occurs when the DHCPv6 server encounters a pre-existing PID file which contains the PID of a running process. This most likely indicates an attempt to start a second instance of the server using the same configuration file. It is possible, though unlikely that the PID file is a remnant left behind by a server crash or power failure and the PID it contains refers to a process other than the server. In such an event, it would be necessary to manually remove the PID file. The first argument is the DHCPv6 process name, the second contains the PID and PID file.

26.15.10 DHCP6_BUFFER_RECEIVED

```
received buffer from %1:%2 to %3:%4 over interface %5
```

Logged at debug log level 40. This debug message is logged when the server has received a packet over the socket. When the message is logged the contents of the received packet hasn't been parsed yet. The only available information is the interface and the source and destination addresses/ports.

26.15.11 DHCP6_BUFFER_UNPACK

```
parsing buffer received from %1 to %2 over interface %3
```

Logged at debug log level 50. This debug message is issued when the server starts parsing the received buffer holding the DHCPv6 message. The arguments specify the source and destination addresses as well as the interface over which the buffer has been received.

26.15.12 DHCP6_BUFFER_WAIT_SIGNAL

```
signal received while waiting for next packet
```

Logged at debug log level 50. This debug message is issued when the server was waiting for the packet, but the wait has been interrupted by the signal received by the process. The signal will be handled before the server starts waiting for next packets.

26.15.13 DHCP6_CB_ON_DEMAND_FETCH_UPDATES_FAIL

```
error on demand attempt to fetch configuration updates from the configuration_
↪ backend(s): %1
```

This error message is issued when the server attempted to fetch configuration updates from the database and this on demand attempt failed. The sole argument which is returned to the config-backend-pull command caller too contains the reason for failure.

26.15.14 DHCP6_CB_PERIODIC_FETCH_UPDATES_FAIL

```
error on periodic attempt to fetch configuration updates from the configuration_
↪ backend(s): %1
```

This error message is issued when the server attempted to fetch configuration updates from the database and this periodic attempt failed. The server will re-try according to the configured value of the config-fetch-wait-time parameter. The sole argument contains the reason for failure.

26.15.15 DHCP6_CB_PERIODIC_FETCH_UPDATES_RETRIES_EXHAUSTED

```
maximum number of configuration fetch attempts: 10, has been exhausted without success
```

This error indicates that the server has made a number of unsuccessful periodic attempts to fetch configuration updates from a configuration backend. The server will continue to operate but won't make any further attempts to fetch configuration updates. The administrator must fix the configuration in the database and reload (or restart) the server.

26.15.16 DHCP6_CLASSES_ASSIGNED

```
%1: client packet has been assigned on %2 message to the following classes: %3
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified classes. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument specifies the DHCPv6 message type. The third argument includes all classes to which the packet has been assigned.

26.15.17 DHCP6_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION

```
%1: client packet has been assigned to the following classes: %2
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified classes. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument includes all classes to which the packet has been assigned.

26.15.18 DHCP6_CLASS_ASSIGNED

```
%1: client packet has been assigned to the following class: %2
```

Logged at debug log level 40. This debug message informs that incoming packet has been assigned to specified class. This is a normal behavior and indicates successful operation. The first argument specifies the client and transaction identification information. The second argument includes the new class to which the packet has been assigned.

26.15.19 DHCP6_CLASS_UNCONFIGURED

```
%1: client packet belongs to an unconfigured class: %2
```

Logged at debug log level 40. This debug message informs that incoming packet belongs to a class which cannot be found in the configuration. Either a hook written before the classification was added to Kea is used, or class naming is inconsistent.

26.15.20 DHCP6_CLIENT_FQDN_SCRUBBED_EMPTY

```
%1: sanitizing client's FQDN option '%2' yielded an empty string
```

Logged at debug log level 50. This debug message is issued when the result of sanitizing the FQDN option(39) sent by the client is an empty string. When this occurs the server will ignore the FQDN option. The arguments include the client and the FQDN option it sent.

26.15.21 DHCP6_CONFIG_COMPLETE

```
DHCPv6 server has completed configuration: %1
```

This is an informational message announcing the successful processing of a new configuration. it is output during server startup, and when an updated configuration is committed by the administrator. Additional information may be provided.

26.15.22 DHCP6_CONFIG_LOAD_FAIL

```
configuration error using file: %1, reason: %2
```

This error message indicates that the DHCPv6 configuration has failed. If this is an initial configuration (during server's startup) the server will fail to start. If this is a dynamic reconfiguration attempt the server will continue to use an old configuration.

26.15.23 DHCP6_CONFIG_PACKET_QUEUE

```
DHCPv6 packet queue info after configuration: %1
```

This informational message is emitted during DHCPv6 server configuration, immediately after configuring the DHCPv6 packet queue. The information shown depends upon the packet queue type selected.

26.15.24 DHCP6_CONFIG_RECEIVED

```
received configuration: %1
```

Logged at debug log level 10. A debug message listing the configuration received by the DHCPv6 server. The source of that configuration depends on used configuration backend.

26.15.25 DHCP6_CONFIG_START

```
DHCPv6 server is processing the following configuration: %1
```

Logged at debug log level 10. This is a debug message that is issued every time the server receives a configuration. That happens start up and also when a server configuration change is committed by the administrator.

26.15.26 DHCP6_CONFIG_SYNTAX_WARNING

```
configuration syntax warning: %1
```

This warning message indicates that the DHCPv6 configuration had a minor syntax error. The error was displayed and the configuration parsing resumed.

26.15.27 DHCP6_CONFIG_UNRECOVERABLE_ERROR

```
DHCPv6 server new configuration failed with an error which cannot be recovered
```

This fatal error message is issued when a new configuration raised an error which cannot be recovered. A correct configuration must be applied as soon as possible as the server is no longer working. The configuration can be fixed in several ways. If the control channel is open, config-set with a valid configuration can be used. Alternatively, the original config file on disk could be fixed and SIGHUP signal could be sent (or the config-reload command issued). Finally, the server could be restarted completely.

26.15.28 DHCP6_CONFIG_UNSUPPORTED_OBJECT

```
DHCPv6 server configuration includes an unsupported object: %1
```

This error message is issued when the configuration includes an unsupported object (i.e. a top level element).

26.15.29 DHCP6_DATA_DIRECTORY_DEPRECATED

'data-directory' has been deprecated and should no longer be used.

This warning message is emitted when the configuration parsing detects that the *data-directory* parameter has been specified. This parameter is no longer supported. This supported path, determined at compile time, may be overridden at runtime by setting the environment variable 'KEA_DHCP6_DATA_DIR'.

26.15.30 DHCP6_DB_RECONNECT_DISABLED

database reconnect is disabled: retries left: %1, reconnect wait time: %2, manager ID: %3, timer: %4

This is an informational message indicating that connectivity to either the lease or host database or both and that automatic reconnect is not enabled.

26.15.31 DHCP6_DB_RECONNECT_FAILED

maximum number of database reconnect attempts: %1, has been exhausted without success, manager ID: %2, timer: %3

This error indicates that the server failed to reconnect to the lease and/or host database(s) after making the maximum configured number of reconnect attempts. This might cause the server to shut down as specified in the configuration. Loss of connectivity is typically a network or database server issue.

26.15.32 DHCP6_DB_RECONNECT_LOST_CONNECTION

database connection lost: manager ID: %1, timer: %2.

This info message indicates that the connection has been lost and the dhcp service might have been disabled, as specified in the configuration, in order to try to recover the connection.

26.15.33 DHCP6_DB_RECONNECT_NO_DB_CTL

unexpected error in database reconnect

This is an error message indicating a programmatic error that should not occur. It prohibits the server from attempting to reconnect to its databases if connectivity is lost, and the server exits. This error should be reported.

26.15.34 DHCP6_DB_RECONNECT_SUCCEEDED

database connection recovered: manager ID: %1, timer: %2.

This info message indicates that the connection has been recovered and the dhcp service has been restored.

26.15.35 DHCP6_DDNS_CREATE_ADD_NAME_CHANGE_REQUEST

%1: created name change request: %2

Logged at debug log level 50. This debug message is logged when the new NameChangeRequest has been created to perform the DNS Update, which adds new RRs.

26.15.36 DHCP6_DDNS_FQDN_GENERATED

```
%1: generated FQDN for the client: %2
```

Logged at debug log level 55. This debug message is logged when the server generated FQDN (name) for the client which message is processed. The names may be generated by the server when required by the server's policy or when the client doesn't provide any specific FQDN in its message to the server. The first argument includes the client and transaction identification information. The second argument includes the generated FQDN.

26.15.37 DHCP6_DDNS_GENERATED_FQDN_UPDATE_FAIL

```
%1: failed to update the lease using address %2, after generating FQDN for a client,↵
↪reason: %3
```

This message indicates the failure when trying to update the lease and/or options in the server's response with the hostname generated by the server from the acquired address. The first argument includes the client and the transaction identification information. The second argument is a leased address. The third argument includes the reason for the failure.

26.15.38 DHCP6_DDNS_GENERATE_FQDN

```
%1: client did not send a FQDN option; FQDN will be
```

Logged at debug log level 50. generated for the client. This debug message is issued when the server did not receive a FQDN option from the client and client name replacement is enabled. This provides a means to create DNS entries for unsophisticated clients.

26.15.39 DHCP6_DDNS_RECEIVE_FQDN

```
%1: received DHCPv6 Client FQDN option: %2
```

Logged at debug log level 50. This debug message is logged when server has found the DHCPv6 Client FQDN Option sent by a client and started processing it. The first argument includes the client and transaction identification information. The second argument includes the received FQDN.

26.15.40 DHCP6_DDNS_REMOVE_OLD_LEASE_FQDN

```
%1: FQDN for a lease: %2 has changed. New values: hostname = %3, reverse mapping = %4,↵
↪forward mapping = %5
```

Logged at debug log level 50. This debug message is logged during lease renewal when an old lease that is no longer being offered has a different FQDN than the renewing lease. Thus the old DNS entries need to be removed. The first argument includes the client and the transaction identification information. The second argument holds the details about the lease for which the FQDN information and/or mappings have changed. The remaining arguments hold the new FQDN information and flags for mappings.

26.15.41 DHCP6_DDNS_REQUEST_SEND_FAILED

```
failed sending a request to kea-dhcp-ddns, error: %1, ncr: %2
```

This error message indicates that IPv6 DHCP server failed to send a DDNS update request to the DHCP-DDNS server. This is most likely a configuration or networking error.

26.15.42 DHCP6_DDNS_RESPONSE_FQDN_DATA

%1: including FQDN option in the server's response: %2

Logged at debug log level 50. This debug message is issued when the server is adding the Client FQDN option in its response to the client. The first argument includes the client and transaction identification information. The second argument includes the details of the FQDN option being included. Note that the name carried in the FQDN option may be modified by the server when the lease is acquired for the client.

26.15.43 DHCP6_DECLINE_FAIL

%1: error on decline lease for address %2: %3

This error message indicates that the software failed to decline a lease from the lease database due to an error during a database operation. The first argument includes the client and the transaction identification information. The second argument holds the IPv6 address which decline was attempted. The last one contains the reason for failure.

26.15.44 DHCP6_DECLINE_FAIL_DUID_MISMATCH

Client %1 sent DECLINE for address %2, but it belongs to client with DUID %3

This informational message is printed when a client attempts to decline a lease, but that lease belongs to a different client. The decline request will be rejected.

26.15.45 DHCP6_DECLINE_FAIL_IAID_MISMATCH

Client %1 sent DECLINE for address %2, but used a wrong IAID (%3), instead of expected %4

This informational message is printed when a client attempts to decline a lease. The server has a lease for this address, it belongs to this client, but the recorded IAID does not match what client has sent. This means the server will reject this Decline.

26.15.46 DHCP6_DECLINE_FAIL_LEASE_WITHOUT_DUID

Client %1 sent DECLINE for address %2, but the associated lease has no DUID

This error condition likely indicates database corruption, as every IPv6 lease is supposed to have a DUID, even if it is an empty one.

26.15.47 DHCP6_DECLINE_FAIL_NO_LEASE

Client %1 sent DECLINE for address %2, but there's no lease for it

This informational message is printed when a client tried to decline an address, but the server has no lease for said address. This means that the server's and client's perception of the leases are different. The likely causes of this could be: a confused (e.g. skewed clock) or broken client (e.g. client moved to a different location and didn't notice) or possibly an attack (a rogue client is trying to decline random addresses). The server will inform the client that his decline request was rejected and client should be able to recover from that.

26.15.48 DHCP6_DECLINE_LEASE

Client %1 sent DECLINE for address %2 and the server marked it as declined. The lease will be recovered in %3 seconds.

This informational message indicates that the client leased an address, but discovered that it is being used by some other device and reported this to the server by sending a Decline message. The server marked the lease as declined. This likely indicates a misconfiguration in the network. Either the server is configured with an incorrect pool or there are devices that have statically assigned addresses that are supposed to be assigned by the DHCP server. Both client (will request a different address) and server (will recover the lease after decline-probation-time elapses) will recover automatically. However, if the underlying problem is not solved, the conditions leading to this message may reappear.

26.15.49 DHCP6_DECLINE_PROCESS_IA

Processing of IA (IAID: %1) from client %2 started.

Logged at debug log level 50. This debug message is printed when the server starts processing an IA_NA option received in Decline message. It's expected that the option will contain an address that is being declined. Specific information will be printed in a separate message.

26.15.50 DHCP6_DEVELOPMENT_VERSION

This software is a development branch of Kea. It is not recommended for production use.

This warning message is displayed when the version is a development (vs stable) one: the second number of the version is odd.

26.15.51 DHCP6_DHCP4O6_PACKET_RECEIVED

received DHCPv4o6 packet from DHCPv4 server (type %1) for %2 port %3 on interface %4

Logged at debug log level 40. This debug message is printed when the server is receiving a DHCPv4o6 from the DHCPv4 server over inter-process communication.

26.15.52 DHCP6_DHCP4O6_RECEIVE_FAIL

failed to receive DHCPv4o6: %1

Logged at debug log level 50. This debug message indicates the inter-process communication with the DHCPv4 server failed. The reason for the error is included in the message.

26.15.53 DHCP6_DHCP4O6_RECEIVING

receiving DHCPv4o6 packet from DHCPv4 server

Logged at debug log level 50. This debug message is printed when the server is receiving a DHCPv4o6 from the DHCPv4 server over inter-process communication socket.

26.15.54 DHCP6_DHCP4O6_RESPONSE_DATA

%1: responding with packet %2 (type %3), packet details: %4

Logged at debug log level 55. A debug message including the detailed data about the packet being sent to the client. The first argument contains the client and the transaction identification information. The second and third argument contains the packet name and type respectively. The fourth argument contains detailed packet information.

26.15.55 DHCP6_DHCP4O6_SEND_FAIL

```
%1: failed to send DHCPv4o6 packet: %2
```

This error is output if the IPv6 DHCP server fails to send an assembled DHCPv4o6 message to a client. The reason for the error is included in the message.

26.15.56 DHCP6_DYNAMIC_RECONFIGURATION

```
initiate server reconfiguration using file: %1, after receiving SIGHUP signal or config-  
↪reload command
```

This is the info message logged when the DHCPv6 server starts reconfiguration as a result of receiving SIGHUP signal or config-reload command.

26.15.57 DHCP6_DYNAMIC_RECONFIGURATION_FAIL

```
dynamic server reconfiguration failed with file: %1
```

This is a fatal error message logged when the dynamic reconfiguration of the DHCP server failed.

26.15.58 DHCP6_DYNAMIC_RECONFIGURATION_SUCCESS

```
dynamic server reconfiguration succeeded with file: %1
```

This is info message logged when the dynamic reconfiguration of the DHCP server succeeded.

26.15.59 DHCP6_FLEX_ID

```
%1: flexible identifier generated for incoming packet: %2
```

Logged at debug log level 40. This debug message is printed when host reservation type is set to flexible identifier and the expression specified in its configuration generated (was evaluated to) an identifier for incoming packet. This debug message is mainly intended as a debugging assistance for flexible identifier.

26.15.60 DHCP6_HOOK_ADDR6_REGISTER_DROP

```
%1: ADDR-REG-INFORM for %2 is dropped, because a callout set the next step to DROP
```

Logged at debug log level 40. This debug message is printed when a callout installed on the `addr6_register` hook point sets the next step to DROP. For this particular hook point, the value setting instructs the server to cancel the address registration and drop the packet.

26.15.61 DHCP6_HOOK_ADDR6_REGISTER_SKIP

```
%1: lease %2 operation for %3 is skipped, because a callout set the next step to SKIP
```

Logged at debug log level 40. This debug message is printed when a callout installed on the `addr6_register` hook point sets the next step to SKIP. For this particular hook point, the value setting instructs the server to skip the lease add or update operation for the registered address so not maintaining the registration state.

26.15.62 DHCP6_HOOK_BUFFER_RCVD_DROP

```
received buffer from %1 to %2 over interface %3 was dropped because a callout set the
↳ drop flag
```

Logged at debug log level 15. This debug message is printed when a callout installed on buffer6_receive hook point set the drop flag. For this particular hook point, the setting of the flag by a callout instructs the server to drop the packet. The arguments specify the source and destination address as well as the name of the interface over which the buffer has been received.

26.15.63 DHCP6_HOOK_BUFFER_RCVD_SKIP

```
received buffer from %1 to %2 over interface %3 is not parsed because a callout set the
↳ next step to SKIP
```

Logged at debug log level 50. This debug message is printed when a callout installed on buffer6_receive hook point set the next step status to skip. For this particular hook point, this value set by a callout instructs the server to not parse the buffer because it was already parsed by the hook. The arguments specify the source and destination address as well as the name of the interface over which the buffer has been received.

26.15.64 DHCP6_HOOK_BUFFER_SEND_SKIP

```
%1: prepared DHCPv6 response was dropped because a callout set the next step to SKIP
```

Logged at debug log level 40. This debug message is printed when a callout installed on buffer6_send hook point set the next step to SKIP value. For this particular hook point, the SKIP setting a callout instructs the server to drop the packet. Server completed all the processing (e.g. may have assigned, updated or released leases), but the response will not be sent to the client. The argument includes the client and transaction identification information.

26.15.65 DHCP6_HOOK_DDNS_UPDATE

```
A hook has updated the DDNS parameters: hostname %1=>%2, forward update %3=>%4, reverse
↳ update %5=>%6
```

Logged at debug log level 15. This message indicates that there was a hook called on ddns6_update hook point and that hook updated the DDNS update parameters: hostname, or whether to conduct forward (A record) or reverse (PTR record) DDNS updates.

26.15.66 DHCP6_HOOK_DECLINE_DROP

```
During Decline processing (client=%1, interface=%2, addr=%3) hook callout set next step
↳ to DROP, dropping packet.
```

Logged at debug log level 15. This message indicates that the server received DECLINE message, it was verified to be correct and matching server's lease information. The server called hooks for the lease6_decline hook point and one of the callouts set next step status to DROP. The server will now abort processing of the packet as if it was never received. The lease will continue to be assigned to this client.

26.15.67 DHCP6_HOOK_DECLINE_SKIP

```
During Decline processing (client=%1, interface=%2, addr=%3) hook callout set status to
↳ SKIP, skipping decline.
```


Logged at debug log level 50. This message indicates that the server received DECLINE message, it was verified to be correct and matching server's lease information. The server called hooks for the lease6_decline hook point and one of the callouts set next step status to SKIP. The server will skip the operation of moving the lease to the declined state and will continue processing the packet. In particular, it will send a REPLY message as if the decline actually took place.

26.15.68 DHCP6_HOOK_LEASE6_RELEASE_NA_SKIP

```
%1: DHCPv6 address lease was not released because a callout set the next step to SKIP
```

Logged at debug log level 40. This debug message is printed when a callout installed on the lease6_release hook point set the next step to SKIP. For this particular hook point, this setting by a callout instructs the server to not release a lease. If a client requested the release of multiples leases (by sending multiple IA options), the server will retain this particular lease and proceed with other releases as usual. The argument holds the client and transaction identification information.

26.15.69 DHCP6_HOOK_LEASE6_RELEASE_PD_SKIP

```
%1: prefix lease was not released because a callout set the next step to SKIP
```

Logged at debug log level 40. This debug message is printed when a callout installed on lease6_release hook point set the next step to SKIP value. For this particular hook point, that setting by a callout instructs the server to not release a lease. If client requested release of multiples leases (by sending multiple IA options), the server will retains this particular lease and will proceed with other renewals as usual. The argument holds the client and transaction identification information.

26.15.70 DHCP6_HOOK_LEASES6_COMMITTED_DROP

```
%1: packet is dropped, because a callout set the next step to DROP
```

Logged at debug log level 15. This debug message is printed when a callout installed on the leases6_committed hook point sets the next step to DROP.

26.15.71 DHCP6_HOOK_LEASES6_COMMITTED_PARK

```
%1: packet is parked, because a callout set the next step to PARK
```

Logged at debug log level 40. This debug message is printed when a callout installed on the leases6_committed hook point sets the next step to PARK.

26.15.72 DHCP6_HOOK_LEASES6_PARKING_LOT_FULL

```
The parked-packet-limit %1, has been reached, dropping query: %2
```

Logged at debug log level 15. This debug message occurs when the parking lot used to hold client queries while the hook library work for them completes has reached or exceeded the limit set by the parked-packet-limit global parameter. This can occur when kea-dhcp6 is using hook libraries (e.g. HA) that implement the "leases6-committed" callout and client queries are arriving faster than those callouts can fulfill them.

26.15.73 DHCP6_HOOK_PACKET_RCVD_SKIP

```
%1: packet is dropped, because a callout set the next step to SKIP
```

Logged at debug log level 40. This debug message is printed when a callout installed on the pkt6_receive hook point sets the next step to SKIP. For this particular hook point, the value setting instructs the server to drop the packet.

26.15.74 DHCP6_HOOK_PACKET_SEND_DROP

%1: prepared DHCPv6 response was not sent because a callout set the next step to DROP

Logged at debug log level 15. This debug message is printed when a callout installed on the pkt6_send hook point set the next step to DROP. For this particular hook point, the setting of the value by a callout instructs the server to drop the packet. This effectively means that the client will not get any response, even though the server processed client's request and acted on it (e.g. possibly allocated a lease). The argument specifies the client and transaction identification information.

26.15.75 DHCP6_HOOK_PACKET_SEND_SKIP

%1: prepared DHCPv6 response is not built because a callout set the next step to SKIP

Logged at debug log level 40. This debug message is printed when a callout installed on the pkt6_send hook point set the next step to SKIP. For this particular hook point, the setting of the value by a callout instructs the server to not build the wire data (pack) because it was already done by the book. The argument specifies the client and transaction identification information.

26.15.76 DHCP6_HOOK_SUBNET6_SELECT_DROP

%1: packet was dropped because a callout set the drop flag

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet6_select hook point set the drop flag. For this particular hook point, the setting of the flag instructs the server to drop the received packet. The argument holds the client and transaction identification information.

26.15.77 DHCP6_HOOK_SUBNET6_SELECT_PARK

%1: packet was parked

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet6_select hook point set the park flag. The argument holds the client and transaction identification information.

26.15.78 DHCP6_HOOK_SUBNET6_SELECT_SKIP

%1: no subnet was selected because a callout set the next step to SKIP

Logged at debug log level 40. This debug message is printed when a callout installed on the subnet6_select hook point set the next step to SKIP value. For this particular hook point, the setting of this value instructs the server not to choose a subnet, an action that severely limits further processing; the server will be only able to offer global options - no addresses or prefixes will be assigned. The argument holds the client and transaction identification information.

26.15.79 DHCP6_INIT_FAIL

failed to initialize Kea server: %1

The server has failed to establish communication with the rest of Kea, failed to read JSON configuration file or encountered any other critical issue that prevents it from starting up properly. Attached error message provides more details about the issue.

26.15.80 DHCP6_LEASE_ADVERT

```
%1: lease for address %2 and iaid=%3 will be advertised
```

Logged at debug log level 50. This informational message indicates that the server will advertise an address to the client in the ADVERTISE message. The client will request allocation of this address with the REQUEST message sent in the next message exchange. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated address and IAID.

26.15.81 DHCP6_LEASE_ADVERT_FAIL

```
%1: failed to advertise an address lease for iaid=%2
```

Logged at debug log level 50. This message indicates that in response to a received SOLICIT, the server failed to advertise a non-temporary lease for a given client. There may be many reasons for such failure. Each failure is logged in a separate log entry. The first argument holds the client and transaction identification information. The second argument holds the IAID.

26.15.82 DHCP6_LEASE_ALLOC

```
%1: lease for address %2 and iaid=%3 has been allocated for %4 seconds
```

Logged at debug log level 50. This informational message indicates that in response to a client's REQUEST message, the server successfully granted a non-temporary address lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated address, IAID and validity lifetime.

26.15.83 DHCP6_LEASE_ALLOC_FAIL

```
%1: failed to grant an address lease for iaid=%2
```

Logged at debug log level 50. This message indicates that in response to a received REQUEST, the server failed to grant a non-temporary address lease for the client. There may be many reasons for such failure. Each failure is logged in a separate log entry. The first argument holds the client and transaction identification information. The second argument holds the IAID.

26.15.84 DHCP6_LEASE_DATA

```
%1: detailed lease information for iaid=%2: %3
```

Logged at debug log level 55. This debug message is used to print the detailed information about the allocated lease or a lease which will be advertised to the client. The first argument holds the client and the transaction identification information. The second argument holds the IAID. The third argument holds the detailed lease information.

26.15.85 DHCP6_LEASE_NA_WITHOUT_DUID

```
%1: address lease for address %2 does not have a DUID
```

This error message indicates a database consistency problem. The lease database has an entry indicating that the given address is in use, but the lease does not contain any client identification. This is most likely due to a software error: please raise a bug report. As a temporary workaround, manually remove the lease entry from the database. The first argument includes the client and transaction identification information. The second argument holds the address to be released.

26.15.86 DHCP6_LEASE_PD_WITHOUT_DUID

```
%1: lease for prefix %2/%3 does not have a DUID
```

This error message indicates a database consistency failure. The lease database has an entry indicating that the given prefix is in use, but the lease does not contain any client identification. This is most likely due to a software error: please raise a bug report. As a temporary workaround, manually remove the lease entry from the database. The first argument includes client and transaction identification information. The second and third argument hold the prefix and the prefix length.

26.15.87 DHCP6_LEASE_QUERY_ERROR_GETTING_RELAY_INFO

```
failed to get relay information for lease: %1, reason: %2
```

This is warning message that indicates the server was unable to use the relay information stored in the lease's user-context to construct the lq-relay-data option for the DHCPV6_LEASEQUERY_REPLY. The server will still send the reply with to the requester but without the lq-relay-data option. The most likely cause for this would be either a corrupted lease file or a programmatic error and it should be reported. The first argument is the lease detail, the second argument is the specific error.

26.15.88 DHCP6_LEASE_QUERY_PACKET_PACK

```
%1: preparing on-wire format of the packet to be sent
```

This debug message is issued when the server starts preparing the on-wire format of the packet to be sent back to the client. The argument specifies the client and the transaction identification information.

26.15.89 DHCP6_LEASE_QUERY_PACKET_PACK_FAILED

```
%1: preparing on-wire-format of the packet to be sent failed %2
```

This error message is issued when preparing an on-wire format of the packet has failed. The first argument identifies the client and the

26.15.90 DHCP6_LEASE_QUERY_PACKET_UNPACK_FAILED

```
failed to parse query from %1 to %2, received over interface %3, reason: %4
```

Logged at debug log level 40. This debug message is issued when the received DHCPV6_LEASEQUERY is malformed and can't be parsed by the buffer6_receive callout. The query will be dropped by the server. The first three arguments specify source IP address, destination IP address and the interface. The last argument provides a reason for failure.

26.15.91 DHCP6_LEASE_QUERY_PREFIX_LENGTH_LIST

```
the list of prefix lengths to use when searching will be: %1
```

Logged at debug log level 40. This debug message is emitted after a (re)configuration event to display the list of delegated prefix lengths that will be used when searching for a delegated prefix to which the query address belongs. The argument is the list of prefix lengths in the order they will be used during searches.

26.15.92 DHCP6_LEASE_QUERY_PROCESS_FAILED

```
processing failed for lease query: %1, reason: %2
```

Logged at debug log level 40. This error message is issued when the server encountered an error processing a DHCPV6_LEASEQUERY. The first argument provides query details, the second an explanation of the error.

26.15.93 DHCP6_LEASE_QUERY_RECEIVED

```
received query: %1
```

Logged at debug log level 40. This debug message is printed when the DHCPV6_LEASEQUERY query has been received.

26.15.94 DHCP6_LEASE_QUERY_REPLY_SEND_FAILED

```
unable to send response: %1, iface: %2, address %3:%4 error: %5
```

This error message is issued when the server was unable to send a lease query reply back to the requester. The first argument provides query details, followed by the output interface, IP address and port, and finally the error itself.

26.15.95 DHCP6_LEASE_QUERY_REPLY_SENT

```
response: %1, sent to %2:%3
```

Logged at debug log level 40. This debug message is printed when a response to a DHCPV6_LEASEQUERY has been sent to a requester. The first argument provides response details, the second and third arguments are the IP address and port to which the response was sent.

26.15.96 DHCP6_LEASE_RENEW

```
%1: lease for address %2 and iaid=%3 has been allocated
```

This informational message indicates that in response to a client's REQUEST message, the server successfully renewed a non-temporary address lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated address and IAID.

26.15.97 DHCP6_LEASE_REUSE

```
%1: lease for address %2 and iaid=%3 has been reused for %4 seconds
```

This informational message indicates that in response to a client's message, the server successfully reused a non-temporary address lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated address, IAID and validity lifetime.

26.15.98 DHCP6_MULTI_THREADING_INFO

```
enabled: %1, number of threads: %2, queue size: %3
```

This is a message listing some information about the multi-threading parameters with which the server is running.

26.15.99 DHCP6_NOT_RUNNING

IPv6 DHCP server is not running

A warning message is issued when an attempt is made to shut down the IPv6 DHCP server but it is not running.

26.15.100 DHCP6_NO_INTERFACES

failed to detect any network interfaces

During startup the IPv6 DHCP server failed to detect any network interfaces and is therefore shutting down.

26.15.101 DHCP6_OPEN_SOCKET

opening service sockets on port %1

Logged at debug log level 0. A debug message issued during startup, this indicates that the IPv6 DHCP server is about to open sockets on the specified port.

26.15.102 DHCP6_OPEN_SOCKETS_FAILED

maximum number of open service sockets attempts: %1, has been exhausted without success

This error indicates that the server failed to bind service sockets after making the maximum configured number of reconnect attempts. This might cause the server to shut down as specified in the configuration.

26.15.103 DHCP6_OPEN_SOCKETS_NO_RECONNECT_CTL

unexpected error in bind service sockets.

This is an error message indicating a programmatic error that should not occur. It prohibits the server from attempting to bind to its service sockets if they are unavailable, and the server exits. This error should be reported.

26.15.104 DHCP6_PACKET_DROP_DHCP_DISABLED

%1: DHCP service is globally disabled

Logged at debug log level 15. This debug message is issued when a packet is dropped because the DHCP service has been temporarily disabled. This affects all received DHCP packets. The service may be enabled by the "dhcp-enable" control command or automatically after a specified amount of time since receiving "dhcp-disable" command.

26.15.105 DHCP6_PACKET_DROP_DROP_CLASS

dropped as member of the special class 'DROP': %1 %2

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified into the special class 'DROP' and dropped. The packet details are displayed.

26.15.106 DHCP6_PACKET_DROP_DROP_CLASS2

dropped as member of the special class 'DROP' after host reservation lookup: %1 %2

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified after host reservation lookup into the special class 'DROP' and dropped. The packet details are displayed.

26.15.107 DHCP6_PACKET_DROP_DROP_CLASS_EARLY

```
dropped as member of the special class 'DROP' after early global host reservations.↵  
↵lookup: %1 %2
```

Logged at debug log level 15. This debug message is emitted when an incoming packet was classified after early global host reservations lookup into the special class 'DROP' and dropped. The packet details are displayed.

26.15.108 DHCP6_PACKET_DROP_DUPLICATE

```
dropped as sent by the same client than a packet being processed by another thread:↵  
↵dropped %1 %2 by thread %3 as duplicate of %4 %5 processed by thread %6
```

Logged at debug log level 15. Currently multi-threading processing avoids races between packets sent by the same client by dropping new packets until processing is finished. Packet details and thread identifiers are included for both packets in this warning message.

26.15.109 DHCP6_PACKET_DROP_PARSE_FAIL

```
%1: failed to parse packet from %2 to %3, received over interface %4, reason: %5, %6
```

Logged at debug log level 15. The DHCPv6 server has received a packet that it is unable to interpret. The reason why the packet is invalid is included in the message.

26.15.110 DHCP6_PACKET_DROP_SERVERID_MISMATCH

```
%1: dropping packet with server identifier: %2, server is using: %3
```

Logged at debug log level 15. A debug message noting that server has received message with server identifier option that not matching server identifier that server is using.

26.15.111 DHCP6_PACKET_DROP_UNICAST

```
%1: dropping unicast %2 packet as this packet should be sent to multicast
```

Logged at debug log level 15. This debug message is issued when the server drops the unicast packet, because packets of this type must be sent to multicast. The first argument specifies the client and transaction identification information, the second argument specifies packet type.

26.15.112 DHCP6_PACKET_OPTIONS_SKIPPED

```
%1: An error unpacking an option, caused subsequent options to be skipped: %2
```

Logged at debug log level 50. A debug message issued when an option failed to unpack correctly, making it impossible to unpack the remaining options in the packet. The server will still attempt to service the packet.

26.15.113 DHCP6_PACKET_PROCESS_EXCEPTION

```
%1: exception occurred during packet processing
```

This error message indicates that a non-standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation.

26.15.114 DHCP6_PACKET_PROCESS_EXCEPTION_MAIN

exception occurred during packet processing

This error message indicates that a non-standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation. This error message may appear in main server processing loop.

26.15.115 DHCP6_PACKET_PROCESS_FAIL

%1: processing of %2 message received from %3 failed: %4

Logged at debug log level 40. This is a general catch-all message indicating that the processing of the specified packet type from the indicated address failed. The reason is given in the message. The server will not send a response but will instead ignore the packet.

26.15.116 DHCP6_PACKET_PROCESS_STD_EXCEPTION

%1: exception occurred during packet processing: %2

This error message indicates that a standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation.

26.15.117 DHCP6_PACKET_PROCESS_STD_EXCEPTION_MAIN

exception occurred during packet processing: %1

This error message indicates that a standard exception was raised during packet processing that was not caught by other, more specific exception handlers. This packet will be dropped and the server will continue operation. This error message may appear in main server processing loop.

26.15.118 DHCP6_PACKET_QUEUE_FULL

multi-threading packet queue is full

Logged at debug log level 40. A debug message noting that the multi-threading packet queue is full so the oldest packet of the queue was dropped to make room for the received one.

26.15.119 DHCP6_PACKET_RECEIVED

%1: %2 (type %3) received from %4 to %5 on interface %6

An INFO message noting that the server has received the specified type of packet on the specified interface. The first argument specifies the client and transaction identification information. The second and third argument specify the name of the DHCPv6 message and its numeric type respectively. The remaining arguments specify the source address, destination IP address and the name of the interface on which the message has been received.

26.15.120 DHCP6_PACKET_RECEIVE_FAIL

error on attempt to receive packet: %1

The IPv6 DHCP server tried to receive a packet but an error occurred during this attempt. The reason for the error is included in the message.

26.15.121 DHCP6_PACKET_SEND

```
%1: trying to send packet %2 (type %3) from [%4]:%5 to [%6]:%7 on interface %8
```

An INFO message noting that the server is attempting to send the specified type of packet. The arguments specify the client identification information (HW address and client identifier), DHCP message name and type, source IPv6 address and port, destination IPv6 address and port and the interface name.

26.15.122 DHCP6_PACKET_SEND_FAIL

```
%1: failed to send DHCPv6 packet: %2
```

This error is output if the IPv6 DHCP server fails to send an assembled DHCP message to a client. The reason for the error is included in the message.

26.15.123 DHCP6_PACK_FAIL

```
%1: failed to assemble response correctly: %2
```

This error is output if the server failed to assemble the data to be returned to the client into a valid packet. The reason is most likely to be to a programming error: please raise a bug report.

26.15.124 DHCP6_PARSER_COMMIT_EXCEPTION

```
parser failed to commit changes
```

On receipt of message containing details to a change of the IPv6 DHCP server configuration, a set of parsers were successfully created, but one of them failed to commit its changes due to a low-level system exception being raised. Additional messages may be output indicating the reason.

26.15.125 DHCP6_PARSER_COMMIT_FAIL

```
parser failed to commit changes: %1
```

On receipt of message containing details to a change of the IPv6 DHCP server configuration, a set of parsers were successfully created, but one of them failed to commit its changes. The reason for the failure is given in the message.

26.15.126 DHCP6_PARSER_EXCEPTION

```
failed to create or run parser for configuration element %1
```

On receipt of message containing details to a change of its configuration, the IPv6 DHCP server failed to create a parser to decode the contents of the named configuration element, or the creation succeeded but the parsing actions and committal of changes failed. The message has been output in response to a non-Kea exception being raised. Additional messages may give further information. The most likely cause of this is that the specification file for the server (which details the allowable contents of the configuration) is not correct for this version of Kea. This may be the result of an interrupted installation of an update to Kea.

26.15.127 DHCP6_PARSER_FAIL

```
failed to create or run parser for configuration element %1: %2
```

On receipt of message containing details to a change of its configuration, the IPv6 DHCP server failed to create a parser to decode the contents of the named configuration element, or the creation succeeded but the parsing actions and committal of changes failed. The reason for the failure is given in the message.

26.15.128 DHCP6_PD_LEASE_ADVERT

```
%1: lease for prefix %2/%3 and iaid=%4 will be advertised
```

Logged at debug log level 50. This informational message indicates that the server will advertise a prefix to the client in the ADVERTISE message. The client will request allocation of this prefix with the REQUEST message sent in the next message exchange. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated prefix, prefix length and IAID.

26.15.129 DHCP6_PD_LEASE_ADVERT_FAIL

```
%1: failed to advertise a prefix lease for iaid=%2
```

Logged at debug log level 50. This message indicates that in response to a received SOLICIT, the server failed to advertise a prefix lease for a given client. There may be many reasons for such failure. Each failure is logged in a separate log entry. The first argument holds the client and transaction identification information. The second argument holds the IAID.

26.15.130 DHCP6_PD_LEASE_ALLOC

```
%1: lease for prefix %2/%3 and iaid=%4 has been allocated for %5 seconds
```

Logged at debug log level 50. This informational message indicates that in response to a client's REQUEST message, the server successfully granted a prefix lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated prefix, prefix length, IAID and validity lifetime.

26.15.131 DHCP6_PD_LEASE_ALLOC_FAIL

```
%1: failed to grant a prefix lease for iaid=%2
```

Logged at debug log level 50. This message indicates that in response to a received REQUEST, the server failed to grant a prefix lease for the client. There may be many reasons for such failure. Each failure is logged in a separate log entry. The first argument holds the client and transaction identification information. The second argument holds the IAID.

26.15.132 DHCP6_PD_LEASE_RENEW

```
%1: lease for prefix %2/%3 and iaid=%4 has been allocated
```

This informational message indicates that in response to a client's REQUEST message, the server successfully renewed a prefix lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated prefix, prefix length and IAID.

26.15.133 DHCP6_PD_LEASE_REUSE

```
%1: lease for prefix %2/%3 and iaid=%4 has been reused for %5 seconds
```

This informational message indicates that in response to a client's message, the server successfully reused a prefix lease. This is a normal behavior and indicates successful operation. The first argument includes the client and transaction identification information. The remaining arguments hold the allocated prefix, prefix length, IAID and validity lifetime.

26.15.134 DHCP6_PROCESS_IA_NA_EXTEND

```
%1: extending lease lifetime for IA_NA option with iaid=%2
```

Logged at debug log level 50. This message is logged when the server is starting to extend the lifetime of the address lease associated with the particular IAID. The first argument includes the client and transaction identification information. The second argument contains the IAID.

26.15.135 DHCP6_PROCESS_IA_NA_RELEASE

```
%1: releasing lease for IA_NA option with iaid=%2
```

Logged at debug log level 50. This message is logged when the server is trying to release the client's as a result of receiving the RELEASE message. The first argument includes the client and transaction identification information. The second argument contains the IAID.

26.15.136 DHCP6_PROCESS_IA_NA_REQUEST

```
%1: server is processing IA_NA option with iaid=%2 and hint=%3
```

Logged at debug log level 50. This is a debug message that indicates the processing of a received IA_NA option. The first argument contains the client and the transaction identification information. The second argument holds the IAID of the IA_NA option. The third argument may hold the hint for the server about the address that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.15.137 DHCP6_PROCESS_IA_NA_SOLICIT

```
%1: server is processing IA_NA option with iaid=%2 and hint=%3
```

Logged at debug log level 50. This is a debug message that indicates the processing of a received IA_NA option. The first argument contains the client and the transaction identification information. The second argument holds the IAID of the IA_NA option. The third argument may hold the hint for the server about the address that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.15.138 DHCP6_PROCESS_IA_PD_EXTEND

```
%1: extending lease lifetime for IA_PD option with iaid=%2
```

Logged at debug log level 50. This message is logged when the server is starting to extend the lifetime of the prefix lease associated with the particular IAID. The first argument includes the client and transaction identification information. The second argument contains the IAID.

26.15.139 DHCP6_PROCESS_IA_PD_REQUEST

```
%1: server is processing IA_PD option with iaid=%2 and hint=%3
```

Logged at debug log level 50. This is a debug message that indicates a processing of received IA_PD option. The first argument contains the client and the transaction identification information. The second argument holds the IAID of the IA_PD option. The third argument may hold the hint for the server about the prefix that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.15.140 DHCP6_PROCESS_IA_PD_SOLICIT

```
%1: server is processing IA_PD option with iaid=%2 and hint=%3
```

Logged at debug log level 50. This is a debug message that indicates a processing of received IA_PD option. The first argument contains the client and the transaction identification information. The second argument holds the IAID of the IA_PD option. The third argument may hold the hint for the server about the prefix that the client would like to have allocated. If there is no hint, the argument should provide the text indicating that the hint hasn't been sent.

26.15.141 DHCP6_QUERY_DATA

```
%1, packet details: %2
```

Logged at debug log level 55. A debug message printing the details of the received packet. The first argument includes the client and the transaction identification information.

26.15.142 DHCP6_QUERY_LABEL

```
received query: %1
```

This information message indicates that a query was received. It displays the client and the transaction identification information.

26.15.143 DHCP6_RAPID_COMMIT

```
%1: Rapid Commit option received, following 2-way exchange
```

Logged at debug log level 50. This debug message is issued when the server found a Rapid Commit option in the client's message and 2-way exchanges are supported by the server for the subnet on which the client is connected. The argument specifies the client and transaction identification information.

26.15.144 DHCP6_RECLAIM_EXPIRED_LEASES_FAIL

```
failed to reclaim expired leases: %1
```

This error message indicates that the reclaim expired leases operation failed and provides the cause of failure.

26.15.145 DHCP6_RECLAIM_EXPIRED_LEASES_SKIPPED

```
dhcp6 service is currently disabled. Try again in %1 seconds.
```

Logged at debug log level 40. This debug message is emitted when lease reclamation was scheduled to begin but skipped because DHCPv6 service was disabled. Reclamation will continue to be scheduled according to the configured value of reclaim-timer-wait-time.

26.15.146 DHCP6_REGISTERED_LEASE_ADD_FAIL

```
error in registered lease add for %1
```

This error message indicates that the registered lease add failed and provides the address being registered.

26.15.147 DHCP6_REGISTERED_LEASE_UPDATE_FAIL

```
error in registered lease update for %1: %2
```

This error message indicates that the registered lease update failed and provides the registered address and the cause of failure.

26.15.148 DHCP6_RELEASE_NA

```
%1: binding for address %2 and iaid=%3 was released properly
```

This informational message indicates that an address was released properly. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second and third argument hold the released IPv6 address and IAID respectively.

26.15.149 DHCP6_RELEASE_NA_DELETED

```
%1: binding for address %2 and iaid=%3 was deleted on release
```

This informational message indicates that an address was deleted on release. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second and third argument hold the released IPv6 address and IAID respectively.

26.15.150 DHCP6_RELEASE_NA_EXPIRED

```
%1: binding for address %2 and iaid=%3 expired on release
```

This informational message indicates that an address expired on release. It is a normal operation during client shutdown. The first argument includes the client and transaction identification information. The second and third argument hold the released IPv6 address and IAID respectively.

26.15.151 DHCP6_RELEASE_NA_FAIL

```
%1: failed to remove address lease for address %2 and iaid=%3
```

This error message indicates that the software failed to remove an address lease from the lease database. It probably due to an error during a database operation: resolution will most likely require administrator intervention (e.g. check if DHCP process has sufficient privileges to update the database). It may also be triggered if a lease was manually removed from the database during RELEASE message processing. The first argument holds the client and transaction identification information. The second and third argument hold the released address and IAID respectively.

26.15.152 DHCP6_RELEASE_NA_FAIL_WRONG_DUID

```
%1: client tried to release address %2, but it belongs to another client using duid=%3
```

This warning message indicates that a client tried to release an address that belongs to a different client. This should not happen in normal circumstances and may indicate a misconfiguration of the client. However, since the client releasing the address will stop using it anyway, there is a good chance that the situation will correct itself.

26.15.153 DHCP6_RELEASE_NA_FAIL_WRONG_IAID

```
%1: client tried to release address %2, but it used wrong IAID (expected %3, but got %4)
```

This warning message indicates that client tried to release an address that does belong to it, but the address was expected to be in a different IA (identity association) container. This probably means that the client's support for multiple addresses is flawed.

26.15.154 DHCP6_RELEASE_PD

```
%1: prefix %2/%3 for iaid=%4 was released properly
```

This informational message indicates that a prefix was released properly. It is a normal operation during client shutdown. The first argument holds the client and transaction identification information. The second and third argument hold the prefix and its length. The fourth argument holds IAID.

26.15.155 DHCP6_RELEASE_PD_DELETED

```
%1: prefix %2/%3 for iaid=%4 was deleted on release
```

This informational message indicates that a prefix was deleted on release. It is a normal operation during client shutdown. The first argument holds the client and transaction identification information. The second and third argument hold the prefix and its length. The fourth argument holds IAID.

26.15.156 DHCP6_RELEASE_PD_EXPIRED

```
%1: prefix %2/%3 for iaid=%4 expired on release
```

This informational message indicates that a prefix expired on release. It is a normal operation during client shutdown. The first argument holds the client and transaction identification information. The second and third argument hold the prefix and its length. The fourth argument holds IAID.

26.15.157 DHCP6_RELEASE_PD_FAIL

```
%1: failed to release prefix %2/%3 for iaid=%4
```

This error message indicates that the software failed to remove a prefix lease from the lease database. It probably due to an error during a database operation: resolution will most likely require administrator intervention (e.g. check if DHCP process has sufficient privileges to update the database). It may also be triggered if a lease was manually removed from the database during RELEASE message processing. The first argument hold the client and transaction identification information. The second and third argument define the prefix and its length. The fourth argument holds the IAID.

26.15.158 DHCP6_RELEASE_PD_FAIL_WRONG_DUID

```
%1: client tried to release prefix %2/%3, but it belongs to another client (duid=%4)
```

This warning message indicates that client tried to release a prefix that belongs to a different client. This should not happen in normal circumstances and may indicate a misconfiguration of the client. However, since the client releasing the prefix will stop using it anyway, there is a good chance that the situation will correct itself. The first argument includes the client and the transaction identification information. The second and third argument include the prefix and prefix length. The last argument holds the DUID of the client holding the lease.

26.15.159 DHCP6_RELEASE_PD_FAIL_WRONG_IAID

```
%1: client tried to release prefix %2/%3, but it used wrong IAID (expected %4, but got  
->%5)
```

This warning message indicates that client tried to release a prefix that does belong to it, but the address was expected to be in a different IA (identity association) container. This probably means that the client's support for multiple prefixes is flawed. The first argument includes the client and transaction identification information. The second and third argument identify the prefix. The fourth and fifth argument hold the expected IAID and IAID found respectively.

26.15.160 DHCP6_REQUIRED_OPTIONS_CHECK_FAIL

```
%1: %2 message received from %3 failed the following check: %4
```

Logged at debug log level 40. This message indicates that received DHCPv6 packet is invalid. This may be due to a number of reasons, e.g. the mandatory client-id option is missing, the server-id forbidden in that particular type of message is present, there is more than one instance of client-id or server-id present, etc. The exact reason for rejecting the packet is included in the message.

26.15.161 DHCP6_RESERVATIONS_LOOKUP_FIRST_ENABLED

```
Multi-threading is enabled and host reservations lookup is always performed first.
```

This is a message informing that host reservations lookup is performed before lease lookup when multi-threading is enabled overwriting configured value.

26.15.162 DHCP6_RESPONSE_DATA

```
%1: responding with packet %2 (type %3), packet details: %4
```

Logged at debug log level 55. A debug message including the detailed data about the packet being sent to the client. The first argument contains the client and the transaction identification information. The second and third argument contains the packet name and type respectively. The fourth argument contains detailed packet information.

26.15.163 DHCP6_SECURITY_CHECKS_DISABLED

```
Invoked with command line option -X, Security checks are disabled!!
```

This warning is emitted when internal security checks normally performed by kea-dhcp6 have been disabled via command line option '-X'. This means the server is not enforcing restrictions on resource paths or permissions. This mode of operation may expose your environment to security vulnerabilities and should only be used after careful consideration.

26.15.164 DHCP6_SERVER_FAILED

```
server failed: %1
```

The IPv6 DHCP server has encountered a fatal error and is terminating. The reason for the failure is included in the message.

26.15.165 DHCP6_SHUTDOWN

```
server shutdown
```

Logged at debug log level 40. The IPv6 DHCP server has terminated normally.

26.15.166 DHCP6_SHUTDOWN_REQUEST

```
shutdown of server requested
```

Logged at debug log level 40. This debug message indicates that a shutdown of the IPv6 server has been requested via a call to the 'shutdown' method of the core Dhcpv6Srv object.

26.15.167 DHCP6_SRV_CONSTRUCT_ERROR

```
error creating Dhcpv6Srv object, reason: %1
```

This error message indicates that during startup, the construction of a core component within the IPv6 DHCP server (the Dhcpv6 server object) has failed. As a result, the server will exit. The reason for the failure is given within the message.

26.15.168 DHCP6_SRV_D2STOP_ERROR

```
error stopping IO with DHCP_DDNS during shutdown: %1
```

This error message indicates that during shutdown, an error occurred while stopping IO between the DHCPv6 server and the DHCP_DDNS server. This is probably due to a programmatic error is not likely to impact either server upon restart. The reason for the failure is given within the message.

26.15.169 DHCP6_SRV_UNLOAD_LIBRARIES_ERROR

```
error unloading hooks libraries during shutdown: %1
```

This error message indicates that during shutdown, unloading hooks libraries failed to close them. If the list of libraries is empty it is a programmatic error in the server code. If it is not empty it could be a programmatic error in one of the hooks libraries which could lead to a crash during finalization.

26.15.170 DHCP6_STARTED

```
Kea DHCPv6 server version %1 started
```

This informational message indicates that the IPv6 DHCP server has processed all configuration information and is ready to process DHCPv6 packets. The version is also printed.

26.15.171 DHCP6_STARTING

```
Kea DHCPv6 server version %1 (%2) starting
```

This informational message indicates that the IPv6 DHCP server has processed any command-line switches and is starting. The version is also printed.

26.15.172 DHCP6_START_INFO

```
pid: %1, server port: %2, client port: %3, verbose: %4
```

Logged at debug log level 0. This is a debug message issued during the IPv6 DHCP server startup. It lists some information about the parameters with which the server is running.

26.15.173 DHCP6_SUBNET_DATA

```
%1: the selected subnet details: %2
```

Logged at debug log level 55. This debug message includes the details of the subnet selected for the client. The first argument includes the client and the transaction identification information. The second argument includes the subnet details.

26.15.174 DHCP6_SUBNET_DYNAMICALLY_CHANGED

```
%1: changed selected subnet %2 to subnet %3 from shared network %4 for client assignments
```

Logged at debug log level 45. This debug message indicates that the server is using another subnet than initially selected for client assignments. This newly selected subnet belongs to the same shared network as the original subnet. Some reasons why the new subnet was selected include: address pool exhaustion in the original subnet or the fact that the new subnet includes some static reservations for this client.

26.15.175 DHCP6_SUBNET_SELECTED

```
%1: the subnet with ID %2 was selected for client assignments
```

Logged at debug log level 45. This is a debug message noting the selection of a subnet to be used for address and option assignment. Subnet selection is one of the early steps in the processing of incoming client message. The first argument includes the client and the transaction identification information. The second argument holds the selected subnet id.

26.15.176 DHCP6_SUBNET_SELECTION_FAILED

```
%1: failed to select subnet for the client
```

Logged at debug log level 50. This debug message indicates that the server failed to select the subnet for the client which has sent a message to the server. The cause is likely due to a misconfiguration of the server. The packet processing will continue, but the response will only contain generic configuration and no addresses or prefixes. The argument includes the client and the transaction identification information.

26.15.177 DHCP6_UNKNOWN_MSG_RECEIVED

```
%1: received unknown message (type %2) on interface %3
```

Logged at debug log level 40. This debug message is printed when server receives a message of unknown type. That could either mean missing functionality or invalid or broken relay or client. The list of formally defined message types is available here: <http://www.iana.org/assignments/dhcpv6-parameters>.

26.15.178 DHCP6_USING_SERVERID

server is using server-id %1 and stores in the file %2

This info message is logged when the server reads its server-id from a file or generates it. This message is a notification to the administrator what server-id will be used and where it is persisted. Typically, there is no need to modify the server id. However, it is possible to do it in the Kea configuration file. It is important to understand the implications of such modification. The clients will remember previous server-id, and will use it to extend their leases. As a result, they will have to go through a rebinding phase to re-acquire their leases and associate them with a new server id.

26.16 DHCPDRV

26.16.1 DHCPDRV_CFGMGR_ADD_IFACE

listening on interface %1

An info message issued when a new interface is being added to the collection of interfaces on which the server listens to DHCP messages.

26.16.2 DHCPDRV_CFGMGR_ADD_SUBNET4

adding subnet %1

Logged at debug log level 40. A debug message reported when the DHCP configuration manager is adding the specified IPv4 subnet to its database.

26.16.3 DHCPDRV_CFGMGR_ADD_SUBNET6

adding subnet %1

Logged at debug log level 40. A debug message reported when the DHCP configuration manager is adding the specified IPv6 subnet to its database.

26.16.4 DHCPDRV_CFGMGR_ALL_IFACES_ACTIVE

enabling listening on all interfaces

Logged at debug log level 40. A debug message issued when the server is being configured to listen on all interfaces.

26.16.5 DHCPDRV_CFGMGR_CFG_DHCP_DDNS

Setting DHCP-DDNS configuration to: %1

Logged at debug log level 40. A debug message issued when the server's DHCP-DDNS settings are changed.

26.16.6 DHCPDRV_CFGMGR_CONFIG4_MERGED

Configuration backend data has been merged.

This is an informational message emitted when the DHCPv4 server has successfully merged configuration data retrieved from its configuration backends into the current configuration.

26.16.7 DHCPDRV_CFGMGR_CONFIG6_MERGED

Configuration backend data has been merged.

This is an informational message emitted when the DHCPv6 server has successfully merged configuration data retrieved from its configuration backends into the current configuration.

26.16.8 DHCPDRV_CFGMGR_CONFIGURE_SERVERID

server configuration includes specification of a server identifier

This warning message is issued when the server specified configuration of a server identifier. If this new configuration overrides an existing server identifier, this will affect existing bindings of the clients. Clients will use old server identifier when they renew their bindings. The server will not respond to those renews, and the clients will eventually transition to rebinding state. The server should reassign existing bindings and the clients will subsequently use new server identifier. It is recommended to not modify the server identifier, unless there is a good reason for it, to avoid increased number of renewals and a need for rebinding (increase of multicast traffic, which may be received by multiple servers).

26.16.9 DHCPDRV_CFGMGR_DEL_SUBNET4

IPv4 subnet %1 removed

Logged at debug log level 40. This debug message is issued when a subnet is successfully removed from the server configuration. The argument identifies the removed subnet.

26.16.10 DHCPDRV_CFGMGR_DEL_SUBNET6

IPv6 subnet %1 removed

Logged at debug log level 40. This debug message is issued when a subnet is successfully removed from the server configuration. The argument identifies the removed subnet.

26.16.11 DHCPDRV_CFGMGR_FLQ_POPULATE_FREE_ADDRESS_LEASES

populating free address leases for the FLQ allocator in subnet %1; it can take a while!

This informational message is issued when the server begins building a queue of free address leases for the given subnet. It can take a considerable amount of time, depending on the size of the address pools.

26.16.12 DHCPDRV_CFGMGR_FLQ_POPULATE_FREE_ADDRESS_LEASES_DONE

populated %1 free address leases for the FLQ allocator in subnet %2 in %3

This informational message is issued when the server ends building a queue of free address leases for a given subnet. The first argument logs the number of free leases, the second argument logs the subnet, and the third argument logs a duration.

26.16.13 DHCPDRV_CFGMGR_FLQ_POPULATE_FREE_PREFIX_LEASES

populating free prefix leases for the FLQ allocator in subnet %1; it can take a while!

This informational message is issued when the server begins building a queue of free leases for the given subnet. It can take a considerable amount of time, depending on the size of the delegated prefix pools.

26.16.14 DHCPDRV_CFGMGR_FLQ_POPULATE_FREE_PREFIX_LEASES_DONE

```
populated %1 free prefix leases for the FLQ allocator in subnet %2 completed in %3
```

This informational message is issued when the server ends building a queue of free prefix leases for a given subnet. The first argument logs the number of free leases, the second argument logs the subnet, and the third argument logs a duration.

26.16.15 DHCPDRV_CFGMGR_IPV4_RESERVATIONS_NON_UNIQUE_IGNORED

```
ignoring "ip-reservations-unique" setting because at least one of the host database_
↳backends does not support non-unique IP reservations in a subnet
```

This warning message is issued when the server failed to use the new setting of the ip-reservations-unique global parameter configured via the configuration backend. Some host database backends used apparently do not support specifying several reservations for the same IP address in a subnet. The administrator should either stop using the backend that does not support this setting or set the value of the ip-reservations-unique to true to resolve the configuration issue.

26.16.16 DHCPDRV_CFGMGR_IPV6_RESERVATIONS_NON_UNIQUE_IGNORED

```
ignoring "ip-reservations-unique" setting because at least one of the host database_
↳backends does not support non unique IP reservations in a subnet
```

This warning message is issued when the server failed to use the new setting of the ip-reservations-unique global parameter configured via the configuration backend. Some host database backends used apparently do not support specifying several reservations for the same IP address or delegated prefix in a subnet. The administrator should either stop using the backend that does not support this setting or set the value of the ip-reservations-unique to true to resolve the configuration issue.

26.16.17 DHCPDRV_CFGMGR_IP_RESERVATIONS_UNIQUE_DUPLICATES_DETECTED

```
the "ip-reservations-unique" flag is set to true and multiple reservations for the IP_
↳address %1 in subnet %2 are not allowed causing error: %3
```

This warning message is issued when the DHCP server is configured to not allow multiple reservations for the same IP address. However, the host database backend contains multiple reservations for the IP address logged as the first argument, in the subnet logged as second argument, causing problems with lease allocation logged as third argument.

26.16.18 DHCPDRV_CFGMGR_IP_RESERVATIONS_UNIQUE_DUPLICATES_POSSIBLE

```
setting "ip-reservations-unique" from false to true poses a risk that some host backends_
↳may still contain multiple reservations for the same IP address
```

This warning message is issued when the DHCP server is configured to not allow multiple reservations for the same IP address. However, the host database backends may still contain multiple reservations for the same IP addresses causing problems with lease allocation for certain addresses. Please ensure that all such duplicates are removed.

26.16.19 DHCPDRV_CFGMGR_NEW_SUBNET4

```
a new subnet has been added to configuration: %1
```

This is an informational message reporting that the configuration has been extended to include the specified IPv4 subnet.

26.16.20 DHCPDRV_CFGMGR_NEW_SUBNET6

```
a new subnet has been added to configuration: %1
```

This is an informational message reporting that the configuration has been extended to include the specified subnet.

26.16.21 DHCPDRV_CFGMGR_OPTION_DUPLICATE

```
multiple options with the code: %1 added to the subnet: %2
```

This warning message is issued on an attempt to configure multiple options with the same option code for the particular subnet. Adding multiple options is uncommon for DHCPv6, but it is not prohibited.

26.16.22 DHCPDRV_CFGMGR_RENEW_GTR_REBIND

```
in %1, the value of renew-timer %2 is greater than the value of rebind-timer %3,
↳ ignoring renew-timer
```

A warning message that indicates the configured renew-timer is greater than the configured rebind-timer. The server will ignore the renew timer value and send the rebind timer value only. This is considered a non-fatal configuration error.

26.16.23 DHCPDRV_CFGMGR_SOCKET_RAW_UNSUPPORTED

```
use of raw sockets is unsupported on this OS, UDP sockets will be used
```

This warning message is logged when the user specified that the DHCPv4 server should use the raw sockets to receive the DHCP messages and respond to the clients, but the use of raw sockets is not supported on the particular environment. The raw sockets are useful when the server must respond to the directly connected clients which don't have an address yet. If the raw sockets are not supported by Kea on the particular platform, Kea will fall back to use of the IP/UDP sockets. The responses to the directly connected clients will be broadcast. The responses to relayed clients will be unicast as usual.

26.16.24 DHCPDRV_CFGMGR_SOCKET_TYPE_DEFAULT

```
"dhcp-socket-type" not specified , using default socket type %1
```

This informational message is logged when the administrator hasn't specified the "dhcp-socket-type" parameter in configuration for interfaces. In such case, the default socket type will be used.

26.16.25 DHCPDRV_CFGMGR_SOCKET_TYPE_SELECT

```
using socket type %1
```

This informational message is logged when the DHCPv4 server selects the socket type to be used for all sockets that will be opened on the interfaces. Typically, the socket type is specified by the server administrator. If the socket type hasn't been specified, the raw socket will be selected. If the raw socket has been selected but Kea doesn't support the use of raw sockets on the particular OS, it will use an UDP socket instead.

26.16.26 DHCPDRV_CFGMGR_SUBNET4

```
retrieved subnet %1 for address hint %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv4 subnet when given the address hint specified as the address is within the subnet.

26.16.27 DHCPDRV_CFGMGR_SUBNET4_ADDR

```
selected subnet %1 for packet received by matching address %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv4 subnet for a received packet. This particular subnet was selected, because an IPv4 address was matched which belonged to that subnet.

26.16.28 DHCPDRV_CFGMGR_SUBNET4_IFACE

```
selected subnet %1 for packet received over interface %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv4 subnet for a packet received over the given interface. This particular subnet was selected, because it was specified as being directly reachable over the given interface. (see 'interface' parameter in the subnet4 definition).

26.16.29 DHCPDRV_CFGMGR_SUBNET4_RELAY

```
selected subnet %1, because of matching relay addr %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv4 subnet, because detected relay agent address matches value specified for this subnet.

26.16.30 DHCPDRV_CFGMGR_SUBNET6

```
retrieved subnet %1 for address hint %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv6 subnet when given the address hint specified as the address is within the subnet.

26.16.31 DHCPDRV_CFGMGR_SUBNET6_IFACE

```
selected subnet %1 for packet received over interface %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv6 subnet for a packet received over given interface. This particular subnet was selected, because it was specified as being directly reachable over given interface. (see 'interface' parameter in the subnet6 definition).

26.16.32 DHCPDRV_CFGMGR_SUBNET6_IFACE_ID

```
selected subnet %1 (interface-id match) for incoming packet
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv6 subnet for a received packet. This particular subnet was selected, because value of interface-id option matched what was configured in the server's interface-id option for that selected subnet6. (see 'interface-id' parameter in the subnet6 definition).

26.16.33 DHCPDRV_CFGMGR_SUBNET6_RELAY

```
selected subnet %1, because of matching relay addr %2
```

Logged at debug log level 40. This is a debug message reporting that the DHCP configuration manager has returned the specified IPv6 subnet, because detected relay agent address matches value specified for this subnet.

26.16.34 DHCPDRV_CFGMGR_UNICAST_LINK_LOCAL

```
specified link local address %1 for unicast traffic on interface %2
```

This warning message is logged when user specified a link-local address to receive unicast traffic. The warning message is issued because it is an uncommon use.

26.16.35 DHCPDRV_CFGMGR_UPDATE_SUBNET4

```
updating subnet %1 (result %2)
```

Logged at debug log level 40. A debug message reported when the DHCP configuration manager is updating the specified IPv4 subnet in its current configuration. Subnet ID and result (expected to be true) are displayed.

26.16.36 DHCPDRV_CFGMGR_UPDATE_SUBNET6

```
updating subnet %1 (result %2)
```

Logged at debug log level 40. A debug message reported when the DHCP configuration manager is replacing the specified IPv6 subnet in its current configuration. Subnet ID and result (expected to be true) are displayed.

26.16.37 DHCPDRV_CFGMGR_USE_ADDRESS

```
listening on address %1, on interface %2
```

A message issued when the server is configured to listen on the explicitly specified IP address on the given interface.

26.16.38 DHCPDRV_CFGMGR_USE_ALLOCATOR

```
using the %1 allocator for %2 leases in subnet %3
```

A message issued when the configuration manager starts using a given allocator for a subnet.

26.16.39 DHCPDRV_CFGMGR_USE_UNICAST

```
listening on unicast address %1, on interface %2
```

An info message issued when configuring the DHCP server to listen on the unicast address on the specific interface.

26.16.40 DHCPDRV_CLASS_WITH_ADDITIONAL_AND_LIFETIMES

```
class: %1 has 'only-in-additional-list' true while specifying one or more lease life_
↪time values. Life time values will be ignored.
```

This warning is emitted whenever a class is configured with 'only-in-addition-list' true as well as specifying one or more lease life time parameters (e.g. 'valid-lifetime', 'preferred-lifetime', or 'offer-lifetime'). Additional list classes are evaluated after lease assignment, thus parameters that would otherwise impact lease life times will have no affect.

26.16.41 DHCP_SRV_CLIENT_CLASS_DEPRECATED

The parameter 'client-class' is deprecated. Use 'client-classes' list parameter instead

This warning message is emitted when configuration parsing detects the use of the deprecated 'client-class' parameter. It has been replaced by 'client-classes'. Users should migrate to the new list parameter.

26.16.42 DHCP_SRV_CLOSE_DB

closing currently open %1 database

Logged at debug log level 40. This is a debug message, issued when the DHCP server closes the currently open lease database. It is issued at program shutdown and whenever the database access parameters are changed: in the latter case, the server closes the currently open database, and opens a database using the new parameters.

26.16.43 DHCP_SRV_DDNS_TTL_TOO_LARGE

%1 of lease life time %2 is %3, using maximum of %4 instead.

Logged at debug log level 55. A debug message issued when the DDNS TTL value calculated using the ddns-ttl-percent if specified or (0.33 if not) is larger than the specified value of ddns-ttl-max. Kea will ignore the value and use the specified maximum instead. The message details include the percentage, the lease life time, the calculated TTL, and the value actually used.

26.16.44 DHCP_SRV_DDNS_TTL_TOO_SMALL

%1 of lease life time %2 is %3, using minimum of %4 instead.

Logged at debug log level 55. A debug message issued when the DDNS TTL value calculated using the ddns-ttl-percent if specified or (0.33 if not) is too small. Kea will ignore the value and use the minimum (ddns-ttl-min if specified or 600 seconds if not). The message details include the percentage, the lease life time, the calculated TTL, and the value actually used.

26.16.45 DHCP_SRV_DHCP4O6_RECEIVED_BAD_PACKET

received bad DHCPv4o6 packet: %1

A bad DHCPv4o6 packet was received.

26.16.46 DHCP_SRV_DHCP_DDNS_ERROR_EXCEPTION

error handler for DHCP_DDNS IO generated an expected exception: %1

This is an error message that occurs when an attempt to send a request to kea-dhcp-ddns fails there registered error handler threw an uncaught exception. This is a programmatic error which should not occur. By convention, the error handler should not propagate exceptions. Please report this error.

26.16.47 DHCP_SRV_DHCP_DDNS_HANDLER_NULL

error handler for DHCP_DDNS IO is not set.

This is an error message that occurs when an attempt to send a request to kea-dhcp-ddns fails and there is no registered error handler. This is a programmatic error which should never occur and should be reported.

26.16.48 DHCP_SRV_DHCP_DDNS_NCR_REJECTED

```
NameChangeRequest rejected by the sender: %1, ncr: %2
```

This is an error message indicating that NameChangeSender used to deliver DDNS update requests to kea-dhcp-ddns rejected the request. This most likely cause is the sender's queue has reached maximum capacity. This would imply that requests are being generated faster than they can be delivered.

26.16.49 DHCP_SRV_DHCP_DDNS_NCR_SENT

```
NameChangeRequest sent to kea-dhcp-ddns: %1
```

Logged at debug log level 50. A debug message issued when a NameChangeRequest has been successfully sent to kea-dhcp-ddns.

26.16.50 DHCP_SRV_DHCP_DDNS_SENDER_STARTED

```
NameChangeRequest sender has been started: %1
```

An informational message issued when a communication with kea-dhcp-ddns has been successfully started.

26.16.51 DHCP_SRV_DHCP_DDNS_SENDER_STOPPED

```
NameChangeRequest sender has been stopped.
```

An informational message issued when a communication with kea-dhcp-ddns has been stopped. This normally occurs during reconfiguration and as part of normal shutdown. It may occur if kea-dhcp-ddns communications break down.

26.16.52 DHCP_SRV_DHCP_DDNS_SUSPEND_UPDATES

```
DHCP_DDNS updates are being suspended.
```

This is a warning message indicating the DHCP_DDNS updates have been turned off. This should only occur if IO errors communicating with kea-dhcp-ddns have been experienced. Any such errors should have preceding entries in the log with details. No further attempts to communicate with kea-dhcp-ddns will be made without intervention.

26.16.53 DHCP_SRV_EVAL_ERROR

```
%1: Expression '%2' evaluated to %3
```

This error message indicates that a problem was encountered while evaluating an expression of a client class. A description of the problem is printed.

26.16.54 DHCP_SRV_EVAL_RESULT

```
%1: Expression '%2' evaluated to %3
```

Logged at debug log level 50. This debug message indicates that the expression of a client class has been successfully evaluated. The client class name and the result value of the evaluation are printed.

26.16.55 DHCPDRV_FORENSIC_BACKENDS_REGISTERED

the following forensic backend types are available: %1

This informational message lists all possible forensic backends that could be used in forensic logging.

26.16.56 DHCPDRV_FORENSIC_BACKEND_DEREGISTER

deregistered forensic backend type: %1

Logged at debug log level 40. This debug message is issued when a backend factory was deregistered. It is no longer possible to use forensic backend of this type.

26.16.57 DHCPDRV_FORENSIC_BACKEND_REGISTER

registered forensic backend type: %1

Logged at debug log level 40. This debug message is issued when a backend factory was successfully registered. It is now possible to use forensic backend of this type.

26.16.58 DHCPDRV_HOOK_LEASE4_RECOVER_SKIP

DHCPv4 lease %1 was not recovered from the declined state because a callout set the skip status.

Logged at debug log level 40. This debug message is printed when a callout installed on lease4_recover hook point set the next step status to SKIP. For this particular hook point, this indicates that the server should not recover the lease from declined state. The server will leave the lease as it is, in the declined state. The server will attempt to recover it the next time decline recovery procedure takes place.

26.16.59 DHCPDRV_HOOK_LEASE4_RENEW_SKIP

DHCPv4 lease was not renewed because a callout set the skip flag.

Logged at debug log level 40. This debug message is printed when a callout installed on lease4_renew hook point set the skip flag. For this particular hook point, the setting of the flag by a callout instructs the server to not renew a lease. The server will use existing lease as it is, without extending its lifetime.

26.16.60 DHCPDRV_HOOK_LEASE4_SELECT_SKIP

Lease4 creation was skipped, because of callout skip flag.

Logged at debug log level 40. This debug message is printed when a callout installed on lease4_select hook point sets the skip flag. It means that the server was told that no lease4 should be assigned. The server will not put that lease in its database and the client will get a NAK packet.

26.16.61 DHCPDRV_HOOK_LEASE6_EXTEND_SKIP

DHCPv6 lease lifetime was not extended because a callout set the skip flag for message %1

Logged at debug log level 40. This debug message is printed when a callout installed on lease6_renew or the lease6_rebind hook point set the skip flag. For this particular hook point, the setting of the flag by a callout instructs the server to not extend the lifetime for a lease. If the client requested renewal of multiple leases (by sending multiple IA options), the server will skip the renewal of the one in question and will proceed with other renewals as usual.

26.16.62 DHCPDRV_HOOK_LEASE6_RECOVER_SKIP

```
DHCPv6 lease %1 was not recovered from declined state because a callout set the skip_
↪status.
```

Logged at debug log level 40. This debug message is printed when a callout installed on lease6_recover hook point set the next step status to SKIP. For this particular hook point, this indicates that the server should not recover the lease from declined state. The server will leave the lease as it is, in the declined state. The server will attempt to recover it the next time decline recovery procedure takes place.

26.16.63 DHCPDRV_HOOK_LEASE6_SELECT_SKIP

```
Lease6 (non-temporary) creation was skipped, because of callout skip flag.
```

Logged at debug log level 40. This debug message is printed when a callout installed on lease6_select hook point sets the skip flag. It means that the server was told that no lease6 should be assigned. The server will not put that lease in its database and the client will get a NoAddrsAvail for that IA_NA option.

26.16.64 DHCPDRV_HOST_MGR_DB_OPEN_CONNECTION_WITH_RETRY_FAILED

```
Failed to connect to database: %1 with error: %2
```

This is an informational message issued when the server failed to connect to the host database. The operation started a retry to connect procedure. The database access string with password redacted is logged, along with the error and details for the reconnect procedure.

26.16.65 DHCPDRV_LEASE4_EXTENDED_INFO_SANITY_FAIL

```
extended info for lease %1 failed checks (%2)
```

This error message is printed when a lease extended info failed to pass sanity checks. The detail of the found problem was displayed and the extended info deleted from the lease user context.

26.16.66 DHCPDRV_LEASE4_EXTENDED_INFO_UPGRADED

```
extended info for lease %1 was upgraded
```

Logged at debug log level 40. This debug message is printed when a lease extended info was upgraded.

26.16.67 DHCPDRV_LEASE6_EXTENDED_INFO_SANITY_FAIL

```
extended info for lease %1 failed checks (%2)
```

This error message is printed when a lease extended info failed to pass sanity checks. The detail of the found problem was displayed and the extended info deleted from the lease user context.

26.16.68 DHCPDRV_LEASE6_EXTENDED_INFO_UPGRADED

```
extended info for lease %1 was upgraded
```

Logged at debug log level 40. This debug message is printed when a lease extended info was upgraded.

26.16.69 DHCPDRV_LEASE_MGR_BACKENDS_REGISTERED

```
the following lease backend types are available: %1
```

This informational message lists all possible lease backends that could be used in lease-database.

26.16.70 DHCPDRV_LEASE_MGR_BACKEND_DEREGISTER

```
deregistered lease backend type: %1
```

Logged at debug log level 40. This debug message is issued when a backend factory was deregistered. It is no longer possible to use lease backend of this type.

26.16.71 DHCPDRV_LEASE_MGR_BACKEND_REGISTER

```
registered lease backend type: %1
```

Logged at debug log level 40. This debug message is issued when a backend factory was successfully registered. It is now possible to use lease backend of this type.

26.16.72 DHCPDRV_LEASE_MGR_CALLBACK_EXCEPTION

```
exception occurred in a lease manager callback for callback type %1, subnet id %2, and ↵  
↵lease %3: %4
```

This warning message is printed when one of the callback functions registered in the lease manager causes an error. The callback functions can serve different purposes and they likely log the detailed error messages. This error message possibly indicates an unhandled error. The first argument indicates a callback type. The second argument prints the subnet id. The third argument prints the lease for which the error has occurred. The last argument prints the error text.

26.16.73 DHCPDRV_LEASE_MGR_CALLBACK_UNKNOWN_EXCEPTION

```
unknown exception occurred in a lease manager callback for callback type %1, subnet id ↵  
↵%2, and lease %3
```

This warning message is printed when one of the callback functions registered in the lease manager causes an unknown error. The callback functions can serve different purposes and they likely log the detailed error messages. This error message possibly indicates an unhandled error. The first argument indicates a callback type. The second argument prints the subnet id. The third argument prints the lease for which the error has occurred. This log message variant contains no error text because it is triggered by an unknown exception.

26.16.74 DHCPDRV_LEASE_MGR_DB_OPEN_CONNECTION_WITH_RETRY_FAILED

```
Failed to connect to database: %1 with error: %2
```

This is an informational message issued when the server failed to connect to the lease database. The operation started a retry to connect procedure. The database access string with password redacted is logged, along with the error and details for the reconnect procedure.

26.16.75 DHCPDRV_LEASE_SANITY_FAIL

The lease %1 with subnet-id %2 failed subnet-id checks (%3).

This warning message is printed when the lease being loaded does not match the configuration. Due to lease-checks value, the lease will be loaded, but it will most likely be unused by Kea, as there is no subnet that matches the IP address associated with the lease.

26.16.76 DHCPDRV_LEASE_SANITY_FAIL_DISCARD

The lease %1 with subnet-id %2 failed subnet-id checks (%3) and was dropped.

This warning message is printed when a lease was loaded, but Kea was told (by setting lease-checks parameter) to discard leases with inconsistent data. The lease was discarded, because either there is no subnet configured with matching subnet-id or the address of the lease does not belong to the subnet.

26.16.77 DHCPDRV_LEASE_SANITY_FIXED

The lease %1 with subnet-id %2 failed subnet-id checks, but was corrected to subnet-id ↪%3.

This informational message is printed when a lease was loaded, but had incorrect subnet-id value. The lease-checks parameter was set to a value that told Kea to try to correct the problem. There is a matching subnet, so Kea updated subnet-id and loaded the lease successfully.

26.16.78 DHCPDRV_MEMFILE_ADD_ADDR4

adding IPv4 lease with address %1

Logged at debug log level 50. A debug message issued when the server is about to add an IPv4 lease with the specified address to the memory file backend database.

26.16.79 DHCPDRV_MEMFILE_ADD_ADDR6

adding IPv6 lease with address %1

Logged at debug log level 50. A debug message issued when the server is about to add an IPv6 lease with the specified address to the memory file backend database.

26.16.80 DHCPDRV_MEMFILE_BEGIN_BUILD_EXTENDED_INFO_TABLES6

building extended info tables with %1 sanity check level, tables %2

Logged at debug log level 40. A debug message issued when the server is building extended info tables. The extended info sanity check level and the fact tables are enabled or disabled are displayed.

26.16.81 DHCPDRV_MEMFILE_BEGIN_EXTRACT_EXTENDED_INFO4

extract extended info with %1 sanity check level%2

Logged at debug log level 40. A debug message issued when the server is extracting extended info. The extended info sanity check level and update in file when requested are displayed.

26.16.82 DHCPDRV_MEMFILE_BUILD_EXTENDED_INFO_TABLES6

```
building extended info tables saw %1 leases, extended info sanity checks modified %2,
↳ leases and %3 leases were entered into tables
```

Extended info tables build was finished. Some statistics are displayed, the updated in database is returned to the command interface.

26.16.83 DHCPDRV_MEMFILE_BUILD_EXTENDED_INFO_TABLES6_ERROR

```
building extended info tables got an exception on the lease for %1: %2
```

An error message issued when the server is building extended info tables and receives an exception processing a lease.

26.16.84 DHCPDRV_MEMFILE_COMMIT

```
committing to memory file database
```

Logged at debug log level 50. The code has issued a commit call. For the memory file database this is a no-op.

26.16.85 DHCPDRV_MEMFILE_CONVERTING_LEASE_FILES

```
running LFC now to convert lease files to the current schema: %1.%2
```

A warning message issued when the server has detected lease files that need to be either upgraded or downgraded to match the server's schema, and that the server is automatically running the LFC process to perform the conversion. This should only occur the first time the server is launched following a Kea installation upgrade (or downgrade).

26.16.86 DHCPDRV_MEMFILE_DB

```
opening memory file lease database: %1
```

This informational message is logged when a DHCP server (either V4 or V6) is about to open a memory file lease database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.16.87 DHCPDRV_MEMFILE_DELETE_ADDR4

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the memory file database for the specified address.

26.16.88 DHCPDRV_MEMFILE_DELETE_ADDR6

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the memory file database for the specified address.

26.16.89 DHCPDRV_MEMFILE_DELETE_EXPIRED_RECLAIMED4

```
deleting reclaimed IPv4 leases that expired more than %1 seconds ago
```

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv4 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.16.90 DHCPDRV_MEMFILE_DELETE_EXPIRED_RECLAIMED6

```
deleting reclaimed IPv6 leases that expired more than %1 seconds ago
```

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv6 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.16.91 DHCPDRV_MEMFILE_DELETE_EXPIRED_RECLAIMED_START

```
starting deletion of %1 expired-reclaimed leases
```

Logged at debug log level 50. A debug message issued when the server has found expired-reclaimed leases to be removed. The number of leases to be removed is logged in the message.

26.16.92 DHCPDRV_MEMFILE_EXTRACT_EXTENDED_INFO4

```
extracting extended info saw %1 leases, extended info sanity checks modified %2 /  
↪updated %3 leases and %4 leases have relay or remote id
```

Logged at debug log level 40. Extended info extraction was finished. Some statistics are displayed, the updated in database is returned to the command interface.

26.16.93 DHCPDRV_MEMFILE_EXTRACT_EXTENDED_INFO4_ERROR

```
extracting extended info got an exception on the lease for %1: %2
```

Logged at debug log level 40. A debug message issued when the server is extracting extended info and receives an exception processing a lease.

26.16.94 DHCPDRV_MEMFILE_FAILED_TO_OPEN

```
Could not open lease file: %1
```

This error is issued when the lease file could not be opened. The argument contains the details. The most likely cause is that the *name* parameter for memfile back end contains a path other than the supported path. The argument will contain the details. The path component may simply be omitted. To override the default supported path, set the environment variable `KEA_DHCP_DATA_DIR` prior to starting the server.

26.16.95 DHCPDRV_MEMFILE_GET4

```
obtaining all IPv4 leases
```

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases from the memory file database.

26.16.96 DHCPDRV_MEMFILE_GET6

obtaining all IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases from the memory file database.

26.16.97 DHCPDRV_MEMFILE_GET6_DUID

obtaining IPv6 leases for DUID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain IPv6 leases from the memory file database for the DUID.

26.16.98 DHCPDRV_MEMFILE_GET_ADDR4

obtaining IPv4 lease for address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the memory file database for the specified address.

26.16.99 DHCPDRV_MEMFILE_GET_ADDR6

obtaining IPv6 lease for address %1 and lease type %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the memory file database for the specified address.

26.16.100 DHCPDRV_MEMFILE_GET_CLIENTID

obtaining IPv4 leases for client ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the memory file database for a client with the specified client identification.

26.16.101 DHCPDRV_MEMFILE_GET_EXPIRED4

obtaining maximum %1 of expired IPv4 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv4 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.16.102 DHCPDRV_MEMFILE_GET_EXPIRED6

obtaining maximum %1 of expired IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv6 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.16.103 DHCPDRV_MEMFILE_GET_HOSTNAME4

obtaining IPv4 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the memory file database for a client with the specified hostname.

26.16.104 DHCPDRV_MEMFILE_GET_HOSTNAME6

obtaining IPv6 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the memory file database for a client with the specified hostname.

26.16.105 DHCPDRV_MEMFILE_GET_HWADDR

obtaining IPv4 leases for hardware address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the memory file database for a client with the specified hardware address.

26.16.106 DHCPDRV_MEMFILE_GET_IAID_DUID

obtaining IPv6 leases for IAID %1 and DUID %2 and lease type %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the memory file database for a client with the specified IAID (Identity Association ID) and DUID (DHCP Unique Identifier).

26.16.107 DHCPDRV_MEMFILE_GET_IAID_SUBID_DUID

obtaining IPv6 leases for IAID %1, Subnet ID %2, DUID %3 and lease type %4

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the memory file database for a client with the specified IAID (Identity Association ID), Subnet ID and DUID (DHCP Unique Identifier).

26.16.108 DHCPDRV_MEMFILE_GET_PAGE4

obtaining at most %1 IPv4 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.16.109 DHCPDRV_MEMFILE_GET_PAGE6

obtaining at most %1 IPv6 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.16.110 DHCPDRV_MEMFILE_GET_RELAYID4

obtaining at most %1 IPv4 leases starting from address %2 with relay id %3 and cltt↵
↵between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a relay id and client transaction time between start and end dates.

26.16.111 DHCPDRV_MEMFILE_GET_RELAYID6

obtaining at most %1 IPv6 leases starting from address %2 with relay id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a relay id.

26.16.112 DHCPDRV_MEMFILE_GET_REMOTEID4

obtaining at most %1 IPv4 leases starting from address %2 with remote id %3 and cltt↵
↵between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a remote id and client transaction time between start and end dates.

26.16.113 DHCPDRV_MEMFILE_GET_REMOTEID6

obtaining at most %1 IPv6 leases starting from address %2 with remote id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a remote id.

26.16.114 DHCPDRV_MEMFILE_GET_SUBID4

obtaining IPv4 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases for a given subnet identifier from the memory file database.

26.16.115 DHCPDRV_MEMFILE_GET_SUBID6

obtaining IPv6 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases for a given subnet identifier from the memory file database.

26.16.116 DHCPDRV_MEMFILE_GET_SUBID_CLIENTID

obtaining IPv4 lease for subnet ID %1 and client ID %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the memory file database for a client with the specified subnet ID and client ID.

26.16.117 DHCPDRV_MEMFILE_GET_SUBID_HWADDR

obtaining IPv4 lease for subnet ID %1 and hardware address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the memory file database for a client with the specified subnet ID and hardware address.

26.16.118 DHCPDRV_MEMFILE_GET_SUBID_PAGE6

obtaining at most %1 IPv6 leases starting from address %2 for subnet ID %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases from the memory file database beginning with the specified address for a given subnet identifier.

26.16.119 DHCPDRV_MEMFILE_LEASE_FILE_LOAD

loading leases from file %1

An info message issued when the server is about to start reading DHCP leases from the lease file. All leases currently held in the memory will be replaced by those read from the file.

26.16.120 DHCPDRV_MEMFILE_LEASE_LOAD

loading lease %1

Logged at debug log level 55. A debug message issued when DHCP lease is being loaded from the file to memory.

26.16.121 DHCPDRV_MEMFILE_LEASE_LOAD_ROW_ERROR

discarding row %1, error: %2

An error message issued if the DHCP lease being loaded from the given row of the lease file fails. The log message should contain the specific reason the row was discarded. The server continues loading the remaining data. This may indicate a corrupt lease file.

26.16.122 DHCPDRV_MEMFILE_LFC_EXECUTE

executing Lease File Cleanup using: %1

An informational message issued when the memfile lease database backend starts a new process to perform Lease File Cleanup.

26.16.123 DHCPDRV_MEMFILE_LFC_LEASE_FILE_RENAME_FAIL

failed to rename the current lease file %1 to %2, reason: %3

An error message logged when the memfile lease database backend fails to move the current lease file to a new file on which the cleanup should be performed. This effectively means that the lease file cleanup does not take place.

26.16.124 DHCPDRV_MEMFILE_LFC_LEASE_FILE_REOPEN_FAIL

```
failed to reopen lease file %1 after preparing input file for lease file cleanup,
reason: %2, new leases will not persist!
```

An error message logged when the memfile lease database backend failed to re-open or re-create the lease file after renaming the lease file for lease file cleanup. The server continues to operate but leases do not persist to disk.

26.16.125 DHCPDRV_MEMFILE_LFC_SETUP

```
setting up the Lease File Cleanup interval to %1 sec
```

An informational message logged when the memfile lease database backend configures the LFC to be executed periodically. The argument holds the interval in seconds in which the LFC will be executed.

26.16.126 DHCPDRV_MEMFILE_LFC_SPAWN_FAIL

```
lease file cleanup failed to run because kea-lfc process couldn't be spawned
```

This error message is logged when the Kea server fails to run kea-lfc, the program that cleans up the lease file. The server will try again the next time a lease file cleanup is scheduled. Although this message should not appear and the reason why it did investigated, the occasional failure to start the lease file cleanup will not impact operations. Should the failure persist however, the size of the lease file will increase without bound.

26.16.127 DHCPDRV_MEMFILE_LFC_START

```
starting Lease File Cleanup
```

An informational message issued when the Memfile lease database backend starts the periodic Lease File Cleanup.

26.16.128 DHCPDRV_MEMFILE_LFC_UNREGISTER_TIMER_FAILED

```
failed to unregister timer 'memfile-lfc': %1
```

Logged at debug log level 40. This debug message is logged when Memfile backend fails to unregister timer used for lease file cleanup scheduling. There are several reasons why this could occur, although the most likely cause is that the system is being shut down and some other component has unregistered the timer. The message includes the reason for this error.

26.16.129 DHCPDRV_MEMFILE_NEEDS_DOWNGRADING

```
version of lease file: %1 schema is later than version %2
```

A warning message issued when the schema of the lease file loaded by the server is newer than the memfile schema of the server. The server converts the lease data from newer schemas to its schema as it is read, therefore the lease information in use by the server will be correct. Note though, that any data stored in newer schema fields will be dropped. What remains is for the file itself to be rewritten using the current schema.

26.16.130 DHCPDRV_MEMFILE_NEEDS_UPGRADING

```
version of lease file: %1 schema is earlier than version %2
```

A warning message issued when the schema of the lease file loaded by the server pre-dates the memfile schema of the server. Note that the server converts the lease data from older schemas to the current schema as it is read, therefore the lease information in use by the server will be correct. What remains is for the file itself to be rewritten using the current schema.

26.16.131 DHCPDRV_MEMFILE_NO_STORAGE

running in non-persistent mode, leases will be lost after restart

A warning message issued when writes of leases to disk have been disabled in the configuration. This mode is useful for some kinds of performance testing but should not be enabled in normal circumstances. Non-persistence mode is enabled when 'persist4=no persist6=no' parameters are specified in the database access string.

26.16.132 DHCPDRV_MEMFILE_READ_HWADDR_FAIL

failed to read hardware address from lease file: %1

A warning message issued when read attempt of the hardware address stored in a disk file failed. The parameter should provide the exact nature of the failure. The database read will continue, but that particular lease will no longer have hardware address associated with it.

26.16.133 DHCPDRV_MEMFILE_ROLLBACK

rolling back memory file database

Logged at debug log level 50. The code has issued a rollback call. For the memory file database this is a no-op.

26.16.134 DHCPDRV_MEMFILE_UPDATE_ADDR4

updating IPv4 lease for address %1

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv4 lease from the memory file database for the specified address.

26.16.135 DHCPDRV_MEMFILE_UPDATE_ADDR6

updating IPv6 lease for address %1

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv6 lease from the memory file database for the specified address.

26.16.136 DHCPDRV_MEMFILE_WIPE_LEASES4

removing all IPv4 leases from subnet %1

This informational message is printed when removal of all leases from specified IPv4 subnet is commencing. This is a result of receiving administrative command.

26.16.137 DHCPDRV_MEMFILE_WIPE_LEASES4_FINISHED

removing all IPv4 leases from subnet %1 finished, removed %2 leases

This informational message is printed when removal of all leases from a specified IPv4 subnet has finished. The number of removed leases is printed.

26.16.138 DHCPDRV_MEMFILE_WIPE_LEASES6

```
removing all IPv6 leases from subnet %1
```

This informational message is printed when removal of all leases from specified IPv6 subnet is commencing. This is a result of receiving administrative command.

26.16.139 DHCPDRV_MEMFILE_WIPE_LEASES6_FINISHED

```
removing all IPv6 leases from subnet %1 finished, removed %2 leases
```

This informational message is printed when removal of all leases from a specified IPv6 subnet has finished. The number of removed leases is printed.

26.16.140 DHCPDRV_MT_DISABLED_QUEUE_CONTROL

```
disabling dhcp queue control when multi-threading is enabled.
```

This warning message is issued when dhcp queue control is disabled automatically if multi-threading is enabled. These two options are incompatible and can not both be enabled at the same time.

26.16.141 DHCPDRV_MULTIPLE_RAW_SOCKETS_PER_IFACE

```
current configuration will result in opening multiple broadcast capable sockets on some  
↪ interfaces and some DHCP messages may be duplicated
```

A warning message issued when the current configuration indicates that multiple sockets, capable of receiving broadcast traffic, will be opened on some of the interfaces. It must be noted that this may lead to receiving and processing the same DHCP message multiple times, as it will be received by each socket individually.

26.16.142 DHCPDRV_NOTYPE_DB

```
no 'type' keyword to determine database backend: %1
```

This is an error message, logged when an attempt has been made to access a database backend, but where no 'type' keyword has been included in the access string. The access string (less any passwords) is included in the message.

26.16.143 DHCPDRV_NO_SOCKETS_OPEN

```
no interface configured to listen to DHCP traffic
```

This warning message is issued when the current server configuration specifies no interfaces that the server should listen on, or when the specified interfaces are not configured to receive the traffic.

26.16.144 DHCPDRV_ONLY_IF_REQUIRED_DEPRECATED

```
The parameter 'only-if-required' is deprecated. Use 'only-in-additional-list' instead
```

This warning message is emitted when configuration parsing detects the use of the deprecated 'only-if-required' parameter. It has been replaced by 'only-in-additional-list'. Users should migrate to the new parameter.

26.16.145 DHCPDRV_OPEN_SOCKET_FAIL

```
failed to open socket: %1
```

A warning message issued when IfaceMgr fails to open and bind a socket. The reason for the failure is appended as an argument of the log message.

26.16.146 DHCPDRV_QUEUE_NCR

```
%1: Name change request to %2 DNS entry queued: %3
```

Logged at debug log level 55. A debug message which is logged when the NameChangeRequest to add or remove a DNS entries for a particular lease has been queued. The first argument includes the client identification information. The second argument indicates whether the DNS entry is to be added or removed. The third argument carries the details of the NameChangeRequest.

26.16.147 DHCPDRV_QUEUE_NCR_FAILED

```
%1: queuing %2 name change request failed for lease %3: %4
```

This error message is logged when sending a NameChangeRequest to DHCP DDNS failed. The first argument includes the client identification information. The second argument indicates whether the DNS entry is to be added or removed. The third argument specifies the leased address. The last argument provides the reason for failure.

26.16.148 DHCPDRV_QUEUE_NCR_SKIP

```
%1: skip queuing name change request for lease: %2
```

Logged at debug log level 50. This debug message is issued when the server decides to not queue the name change request because the lease doesn't include the FQDN, the forward and reverse update is disabled for this lease or the DNS updates are disabled in the configuration. The first argument includes the client identification information. The second argument includes the leased address.

26.16.149 DHCPDRV_REQUIRE_CLIENT_CLASSES_DEPRECATED

```
The parameter 'require-client-classes' is deprecated. Use 'evaluate-additional-classes'  instead
```

This warning message is emitted when configuration parsing detects the use of the deprecated 'require-client-classes' parameter. It has been replaced by 'evaluate-additional-classes'. Users should migrate to the new parameter.

26.16.150 DHCPDRV_SUBNET4O6_SELECT_FAILED

```
Failed to select any subnet for the DHCPv4o6 packet
```

Logged at debug log level 40. A debug message issued when the server was unable to select any subnet for the DHCPv4o6 packet.

26.16.151 DHCPDRV_SUBNET4_SELECT_BY_ADDRESS_NO_MATCH

```
No subnet matches address: %1
```

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified address.

26.16.152 DHCPDRV_SUBNET4_SELECT_BY_INTERFACE_NO_MATCH

No subnet matches interface: %1

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified interface name.

26.16.153 DHCPDRV_SUBNET4_SELECT_BY_RELAY_ADDRESS_NO_MATCH

No subnet matches relay address: %1

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified relay address.

26.16.154 DHCPDRV_SUBNET4_SELECT_NO_RAI_OPTIONS

No RAI options found to use for subnet selection.

Logged at debug log level 40. A debug message issued by the server when the client query does not include RAI options suitable for use with subnet selection.

26.16.155 DHCPDRV_SUBNET4_SELECT_NO_RELAY_ADDRESS

Relay address (giaddr) in client packet is empty.

Logged at debug log level 40. A debug message issued when no relay address was specified to use for subnet selection.

26.16.156 DHCPDRV_SUBNET4_SELECT_NO_USABLE_ADDRESS

No subnet selected because no suitable address to use for subnet selection was found.

Logged at debug log level 40. A debug message issued when the server was find a suitable address to use for subnet selection.

26.16.157 DHCPDRV_SUBNET6_SELECT_BY_ADDRESS_NO_MATCH

No subnet matches address: %1

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified address.

26.16.158 DHCPDRV_SUBNET6_SELECT_BY_INTERFACE_ID_NO_MATCH

No subnet matches interface id: %1

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified interface id.

26.16.159 DHCPDRV_SUBNET6_SELECT_BY_INTERFACE_NO_MATCH

No subnet matches interface: %1

Logged at debug log level 40. A debug message issued when the server was unable to select a subnet using the specified interface name.

26.16.160 DHCPDRV_TEMPLATE_EVAL_ERROR

```
%1: Expression '%2' evaluated to %3
```

This error message indicates that a problem was encountered while evaluating an expression of a template client class. A description of the problem is printed.

26.16.161 DHCPDRV_TEMPLATE_EVAL_RESULT

```
%1: Expression '%2' evaluated to %3
```

Logged at debug log level 50. This debug message indicates that the expression of a template client class has been successfully evaluated. The client class name and the result value of the evaluation are printed.

26.16.162 DHCPDRV_TIMERMGR_CALLBACK_FAILED

```
running handler for timer %1 caused exception: %2
```

This error message is emitted when the timer elapsed and the operation associated with this timer has thrown an exception. The timer name and the reason for exception is logged.

26.16.163 DHCPDRV_TIMERMGR_REGISTER_TIMER

```
registering timer: %1, using interval: %2 ms
```

Logged at debug log level 40. A debug message issued when the new interval timer is registered in the Timer Manager. This timer will have a callback function associated with it, and this function will be executed according to the interval specified. The unique name of the timer and the interval at which the callback function will be executed is included in the message.

26.16.164 DHCPDRV_TIMERMGR_RUN_TIMER_OPERATION

```
running operation for timer: %1
```

Logged at debug log level 50. A debug message issued when the Timer Manager is about to run a periodic operation associated with the given timer. An example of such operation is a periodic cleanup of expired leases. The name of the timer is included in the message.

26.16.165 DHCPDRV_TIMERMGR_START_TIMER

```
starting timer: %1
```

Logged at debug log level 40. A debug message issued when the registered interval timer is being started. If this operation is successful the timer will periodically execute the operation associated with it. The name of the started timer is included in the message.

26.16.166 DHCPDRV_TIMERMGR_STOP_TIMER

```
stopping timer: %1
```

Logged at debug log level 40. A debug message issued when the registered interval timer is being stopped. The timer remains registered and can be restarted if necessary. The name of the timer is included in the message.

26.16.167 DHCPDRV_TIMERMGR_UNREGISTER_ALL_TIMERS

```
unregistering all timers
```

Logged at debug log level 40. A debug message issued when all registered interval timers are being unregistered from the Timer Manager.

26.16.168 DHCPDRV_TIMERMGR_UNREGISTER_TIMER

```
unregistering timer: %1
```

Logged at debug log level 40. A debug message issued when one of the registered interval timers is unregistered from the Timer Manager. The name of the timer is included in the message.

26.16.169 DHCPDRV_UNKNOWN_DB

```
unknown database type: %1
```

The database access string specified a database type (given in the message) that is unknown to the software. This is a configuration error.

26.17 DHCP

26.17.1 DHCP_DDNS_ADD_FAILED

```
DHCP_DDNS Request ID %1: Transaction outcome %2
```

This is an error message issued after DHCP_DDNS attempts to submit DNS mapping entry additions have failed. The precise reason for the failure should be documented in preceding log entries.

26.17.2 DHCP_DDNS_ADD_SUCCEEDED

```
DHCP_DDNS Request ID %1: successfully added the DNS mapping addition for this request: %2
```

This is an informational message issued after DHCP_DDNS has submitted DNS mapping additions which were received and accepted by an appropriate DNS server.

26.17.3 DHCP_DDNS_AT_MAX_TRANSACTIONS

```
application has %1 queued requests but has reached maximum number of %2 concurrent_
↳ transactions
```

Logged at debug log level 55. This is a debug message that indicates that the application has DHCP_DDNS requests in the queue but is working as many concurrent requests as allowed.

26.17.4 DHCP_DDNS_CLEARED_FOR_SHUTDOWN

```
application has met shutdown criteria for shutdown type: %1
```

Logged at debug log level 0. This is a debug message issued when the application has been instructed to shutdown and has met the required criteria to exit.

26.17.5 DHCP_DDNS_CONFIGURE

```
configuration %1 received: %2
```

Logged at debug log level 40. This is a debug message issued when the DHCP-DDNS application configure method has been invoked.

26.17.6 DHCP_DDNS_CONFIGURED_CALLOUT_DROP

```
configuration was rejected because a callout set the next step to 'drop': %1
```

This error message indicates that the DHCP-DDNS had failed configuration attempt because the next step of the configured callout was set to 'drop' by a hook library. The error message provided by the hook library is displayed.

26.17.7 DHCP_DDNS_CONFIG_CHECK_FAIL

```
DHCP-DDNS server configuration check failed: %1
```

This error message indicates that the DHCP-DDNS had failed configuration check. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.17.8 DHCP_DDNS_CONFIG_FAIL

```
DHCP-DDNS server configuration failed: %1
```

This error message indicates that the DHCP-DDNS had failed configuration attempt. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.17.9 DHCP_DDNS_CONFIG_SYNTAX_WARNING

```
DHCP-DDNS server configuration syntax warning: %1
```

This warning message indicates that the DHCP-DDNS configuration had a minor syntax error. The error was displayed and the configuration parsing resumed.

26.17.10 DHCP_DDNS_FAILED

```
application experienced a fatal error: %1
```

This is a debug message issued when the DHCP-DDNS application encounters an unrecoverable error from within the event loop.

26.17.11 DHCP_DDNS_FORWARD_ADD_BAD_DNSCLIENT_STATUS

```
DHCP_DDNS Request ID %1: received an unknown DNSClient status: %2, while adding a  
↪forward address mapping for FQDN %3 to DNS server %4
```

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was adding a forward address mapping. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.12 DHCP_DDNS_FORWARD_ADD_BUILD_FAILURE

DNS Request ID %1: update message to add a forward DNS entry could not be constructed.
↳ for this request: %2, reason: %3

This is an error message issued when an error occurs attempting to construct the server bound packet requesting a forward address addition. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue.

26.17.13 DHCP_DDNS_FORWARD_ADD_IO_ERROR

DHCP_DDNS Request ID %1: encountered an IO error sending a forward mapping add for FQDN
↳ %2 to DNS server %3

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a forward address add. The application will retry against the same server or others as appropriate.

26.17.14 DHCP_DDNS_FORWARD_ADD_REJECTED

DNS Request ID %1: Server, %2, rejected a DNS update request to add the address mapping.
↳ for FQDN, %3, with an RCODE: %4

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.15 DHCP_DDNS_FORWARD_ADD_RESP_CORRUPT

DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while
↳ adding forward address mapping for FQDN, %3

This is an error message issued when the response received by DHCP_DDNS, to a update request to add a forward address mapping, is mangled or malformed. The application will retry against the same server or others as appropriate.

26.17.16 DHCP_DDNS_FORWARD_ADD_TIMEOUT

DHCP_DDNS Request ID %1: timed out waiting for a response to forward mapping add for
↳ FQDN %2 to DNS server %3

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a forward address add. The application will retry against the same server or others as appropriate.

26.17.17 DHCP_DDNS_FORWARD_REMOVE_ADDRS_BAD_DNSCLIENT_STATUS

DHCP_DDNS Request ID %1: received an unknown DNSClient status: %2, while removing a
↳ forward address mapping for FQDN %3 to DNS server %4

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was removing a forward address mapping. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.18 DHCP_DDNS_FORWARD_REMOVE_ADDRS_BUILD_FAILURE

DNS Request ID %1: update message to remove a forward DNS Address entry could not be constructed for this request: %2, reason: %3

This is an error message issued when an error occurs attempting to construct the server bound packet requesting a forward address (A or AAAA) removal. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue. */sar/*

26.17.19 DHCP_DDNS_FORWARD_REMOVE_ADDRS_IO_ERROR

DHCP_DDNS Request ID %1: encountered an IO error sending a forward mapping address removal for FQDN %2 to DNS server %3

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a forward address remove. The application will retry against the same server or others as appropriate.

26.17.20 DHCP_DDNS_FORWARD_REMOVE_ADDRS_REJECTED

DNS Request ID %1: Server, %2, rejected a DNS update request to remove the forward address mapping for FQDN, %3, with an RCODE: %4

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.21 DHCP_DDNS_FORWARD_REMOVE_ADDRS_RESP_CORRUPT

DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while removing forward address mapping for FQDN, %3

This is an error message issued when the response received by DHCP_DDNS, to a update request to remove a forward address mapping, is mangled or malformed. The application will retry against the same server or others as appropriate.

26.17.22 DHCP_DDNS_FORWARD_REMOVE_ADDRS_TIMEOUT

DHCP_DDNS Request ID %1: timed out waiting for a response to forward mapping address removal for FQDN %2 to DNS server %3

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a forward mapping address removal. The application will retry against the same server or others as appropriate.

26.17.23 DHCP_DDNS_FORWARD_REMOVE_RRS_BAD_DNSCLIENT_STATUS

DHCP_DDNS Request ID %1: received an unknown DNSclient status: %2, while removing forward RRs for FQDN %3 to DNS server %4

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was removing forward RRs. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.24 DHCP_DDNS_FORWARD_REMOVE_RRS_BUILD_FAILURE

```
DNS Request ID %1: update message to remove forward DNS RR entries could not be
↳constructed for this request: %2, reason: %3
```

This is an error message issued when an error occurs attempting to construct the server bound packet requesting forward RR (DHCID RR) removal. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue.

26.17.25 DHCP_DDNS_FORWARD_REMOVE_RRS_IO_ERROR

```
DHCP_DDNS Request ID %1: encountered an IO error sending a forward RR removal for FQDN
↳%2 to DNS server %3
```

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a forward RR remove. The application will retry against the same server.

26.17.26 DHCP_DDNS_FORWARD_REMOVE_RRS_REJECTED

```
DNS Request ID %1: Server, %2, rejected a DNS update request to remove forward RR
↳entries for FQDN, %3, with an RCODE: %4
```

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.27 DHCP_DDNS_FORWARD_REMOVE_RRS_RESP_CORRUPT

```
DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while
↳removing forward RRs for FQDN, %3
```

This is an error message issued when the response received by DHCP_DDNS, to a update request to remove forward RRs mapping, is mangled or malformed. The application will retry against the same server or others as appropriate. */sar/*

26.17.28 DHCP_DDNS_FORWARD_REMOVE_RRS_TIMEOUT

```
DHCP_DDNS Request ID %1: timed out waiting for response to forward RR removal for FQDN
↳%2 to DNS server %3
```

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a forward RR removal. The application will retry against the same server or others as appropriate.

26.17.29 DHCP_DDNS_FORWARD_REPLACE_BAD_DNSCLIENT_STATUS

```
DHCP_DDNS Request ID %1: received an unknown DNSclient status: %2, while replacing
↳forward address mapping for FQDN %3 to DNS server %4
```

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was replacing a forward address mapping. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.30 DHCP_DDNS_FORWARD_REPLACE_BUILD_FAILURE

DNS Request ID %1: update message to replace a forward DNS entry could not be constructed from this request: %2, reason: %3

This is an error message issued when an error occurs attempting to construct the server bound packet requesting a forward address replacement. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue.

26.17.31 DHCP_DDNS_FORWARD_REPLACE_IO_ERROR

DHCP_DDNS Request ID %1: encountered an IO error sending a forward mapping replace for FQDN %2 to DNS server %3

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a forward mapping replace. The application will retry against the same server or others as appropriate.

26.17.32 DHCP_DDNS_FORWARD_REPLACE_REJECTED

DNS Request ID %1: Server, %2, rejected a DNS update request to replace the address mapping for FQDN, %3, with an RCODE: %4

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.33 DHCP_DDNS_FORWARD_REPLACE_RESP_CORRUPT

DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while replacing forward address mapping for FQDN, %3

This is an error message issued when the response received by DHCP_DDNS, to a update request to replace a forward address mapping, is mangled or malformed. The application will retry against the same server or others as appropriate.

26.17.34 DHCP_DDNS_FORWARD_REPLACE_TIMEOUT

DHCP_DDNS Request ID %1: timed out waiting for a response to forward mapping replace for FQDN %2 to DNS server %3

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a forward mapping replace. The application will retry against the same server or others as appropriate.

26.17.35 DHCP_DDNS_FWD_REQUEST_IGNORED

Request ID %1: Forward updates are disabled, the forward portion of request will be ignored: %2

Logged at debug log level 55. This is a debug message issued when forward DNS updates are disabled and DHCP_DDNS receives an update request containing a forward DNS update. The forward update will not performed.

26.17.36 DHCP_DDNS_INVALID_NCR

```
application received an invalid DNS update request: %1
```

This is an error message that indicates that an invalid request to update a DNS entry was received by the application. Either the format or the content of the request is incorrect. The request will be ignored.

26.17.37 DHCP_DDNS_INVALID_RESPONSE

```
received response to DNS Update message is malformed: %1
```

Logged at debug log level 50. This is a debug message issued when the DHCP-DDNS application encountered an error while decoding a response to DNS Update message. Typically, this error will be encountered when a response message is malformed.

26.17.38 DHCP_DDNS_LISTENING_ON_ALL_INTERFACES

```
the DHCP-DDNS server has been configured to listen on %1. This is an insecure_
↳ configuration supported for testing purposes only
```

This is a warning message issued when the DHCP-DDNS server is configured to listen at either *0.0.0.0* or *::*. It is possible for a malicious attacker to send bogus NameChangeRequests to it and change entries in the DNS. For this reason, listening on all interfaces should only be used when deploying in containers or for testing purposes. A future version of Kea will disable this ability by default.

26.17.39 DHCP_DDNS_NCR_FLUSH_IO_ERROR

```
DHCP-DDNS Last send before stopping did not complete successfully: %1
```

This is an error message that indicates the DHCP-DDNS client was unable to complete the last send prior to exiting send mode. This is a programmatic error, highly unlikely to occur, and should not impair the application's ability to process requests.

26.17.40 DHCP_DDNS_NCR_LISTEN_CLOSE_ERROR

```
application encountered an error while closing the listener used to receive_
↳ NameChangeRequests : %1
```

This is an error message that indicates the application was unable to close the listener connection used to receive NameChangeRequests. Closure may occur during the course of error recovery or during normal shutdown procedure. In either case the error is unlikely to impair the application's ability to process requests but it should be reported for analysis.

26.17.41 DHCP_DDNS_NCR_RECV_NEXT_ERROR

```
application could not initiate the next read following a request receive.
```

This is an error message indicating that NameChangeRequest listener could not start another read after receiving a request. While possible, this is highly unlikely and is probably a programmatic error. The application should recover on its own.

26.17.42 DHCP_DDNS_NCR_SEND_CLOSE_ERROR

DHCP-DDNS client encountered an error while closing the sender connection used to send `NameChangeRequests`: %1

This is an error message that indicates the DHCP-DDNS client was unable to close the connection used to send `NameChangeRequests`. Closure may occur during the course of error recovery or during normal shutdown procedure. In either case the error is unlikely to impair the client's ability to send requests but it should be reported for analysis.

26.17.43 DHCP_DDNS_NCR_SEND_NEXT_ERROR

DHCP-DDNS client could not initiate the next request send following send completion: %1

This is an error message indicating that `NameChangeRequest` sender could not start another send after completing the send of the previous request. While possible, this is highly unlikely and is probably a programmatic error. The application should recover on its own.

26.17.44 DHCP_DDNS_NCR_UDP_CLEAR_READY_ERROR

NCR UDP watch socket failed to clear: %1

This is an error message that indicates the application was unable to reset the UDP NCR sender ready status after completing a send. This is programmatic error that should be reported. The application may or may not continue to operate correctly.

26.17.45 DHCP_DDNS_NCR_UDP_RECV_CANCELED

UDP socket receive was canceled while listening for DNS Update requests

Logged at debug log level 40. This is a debug message indicating that the listening on a UDP socket for DNS update requests has been canceled. This is a normal part of suspending listening operations.

26.17.46 DHCP_DDNS_NCR_UDP_RECV_ERROR

UDP socket receive error while listening for DNS Update requests: %1

This is an error message indicating that an I/O error occurred while listening over a UDP socket for DNS update requests. This could indicate a network connectivity or system resource issue.

26.17.47 DHCP_DDNS_NCR_UDP_SEND_CANCELED

UDP socket send was canceled while sending a DNS Update request to DHCP_DDNS: %1

This is an informational message indicating that sending requests via UDP socket to DHCP_DDNS has been interrupted. This is a normal part of suspending send operations.

26.17.48 DHCP_DDNS_NCR_UDP_SEND_ERROR

UDP socket send error while sending a DNS Update request: %1

This is an error message indicating that an IO error occurred while sending a DNS update request to DHCP_DDNS over a UDP socket. This could indicate a network connectivity or system resource issue.

26.17.49 DHCP_DDNS_NOT_ON_LOOPBACK

the DHCP-DDNS server has been configured to listen on %1 which is not the local loopback.
 ↳ This is an insecure configuration supported for testing purposes only

This is a warning message issued when the DHCP-DDNS server is configured to listen at an address other than the loopback address (127.0.0.1 or ::1). It is possible for a malicious attacker to send bogus NameChangeRequests to it and change entries in the DNS. For this reason, addresses other than the IPv4 or IPv6 loopback addresses should only be used for testing purposes. A future version of Kea will implement authentication to guard against such attacks.

26.17.50 DHCP_DDNS_NO_ELIGIBLE_JOBS

although there are queued requests, there are pending transactions for each, Queue_↳
 ↳count: %1 Transaction count: %2

Logged at debug log level 55. This is a debug message issued when all of the queued requests represent clients for which there is an update already in progress. This may occur under normal operations but should be temporary situation.

26.17.51 DHCP_DDNS_NO_FWD_MATCH_ERROR

Request ID %1: the configured list of forward DDNS domains does not contain a match for:
 ↳%2 The request has been discarded.

This is an error message that indicates that DHCP_DDNS received a request to update the forward DNS information for the given FQDN but for which there are no configured DDNS domains in the DHCP_DDNS configuration. Either the DHCP_DDNS configuration needs to be updated or the source of the FQDN itself should be investigated.

26.17.52 DHCP_DDNS_NO_MATCH

No DNS servers match FQDN %1

This is warning message issued when there are no domains in the configuration which match the cited fully qualified domain name (FQDN). The DNS Update request for the FQDN cannot be processed.

26.17.53 DHCP_DDNS_NO_REV_MATCH_ERROR

Request ID %1: the configured list of reverse DDNS domains does not contain a match for:
 ↳%2 The request has been discarded.

This is an error message that indicates that DHCP_DDNS received a request to update the reverse DNS information for the given FQDN but for which there are no configured DDNS domains in the DHCP_DDNS configuration. Either the DHCP_DDNS configuration needs to be updated or the source of the FQDN itself should be investigated.

26.17.54 DHCP_DDNS_QUEUE_MGR_QUEUE_FULL

application request queue has reached maximum number of entries %1

This an error message indicating that DHCP-DDNS is receiving DNS update requests faster than they can be processed. This may mean the maximum queue needs to be increased, the DHCP-DDNS clients are simply generating too many requests too quickly, or perhaps upstream DNS servers are experiencing load issues.

26.17.55 DHCP_DDNS_QUEUE_MGR_QUEUE_RECEIVE

```
Request ID %1: received and queued a request.
```

Logged at debug log level 55. This is an informational message indicating that the NameChangeRequest listener used by DHCP-DDNS to receive a request has received a request and queued it for further processing.

26.17.56 DHCP_DDNS_QUEUE_MGR_RECONFIGURING

```
application is reconfiguring the queue manager
```

Logged at debug log level 40. This is an informational message indicating that DHCP_DDNS is reconfiguring the queue manager as part of normal startup or in response to a new configuration.

26.17.57 DHCP_DDNS_QUEUE_MGR_RECOVERING

```
application is attempting to recover from a queue manager IO error
```

This is an informational message indicating that DHCP_DDNS is attempting to restart the queue manager after it suffered an IO error while receiving requests.

26.17.58 DHCP_DDNS_QUEUE_MGR_RECV_ERROR

```
application's queue manager was notified of a request receive error by its listener.
```

This is an error message indicating that the NameChangeRequest listener used by DHCP-DDNS to receive requests encountered an IO error. There should be corresponding log messages from the listener layer with more details. This may indicate a network connectivity or system resource issue.

26.17.59 DHCP_DDNS_QUEUE_MGR_RESUME_ERROR

```
application could not restart the queue manager, reason: %1
```

This is an error message indicating that DHCP_DDNS's Queue Manager could not be restarted after stopping due to a full receive queue. This means that the application cannot receive requests. This is most likely due to DHCP_DDNS configuration parameters referring to resources such as an IP address or port, that is no longer unavailable. DHCP_DDNS will attempt to restart the queue manager if given a new configuration.

26.17.60 DHCP_DDNS_QUEUE_MGR_RESUMING

```
application is resuming listening for requests now that the request queue size has↵  
↵reached %1 of a maximum %2 allowed
```

This is an informational message indicating that DHCP_DDNS, which had stopped accepting new requests, has processed enough entries from the receive queue to resume accepting requests.

26.17.61 DHCP_DDNS_QUEUE_MGR_STARTED

```
application's queue manager has begun listening for requests.
```

Logged at debug log level 0. This is a debug message indicating that DHCP_DDNS's Queue Manager has successfully started and is now listening for NameChangeRequests.

26.17.62 DHCP_DDNS_QUEUE_MGR_START_ERROR

```
application could not start the queue manager, reason: %1
```

This is an error message indicating that DHCP_DDNS's Queue Manager could not be started. This means that the application cannot receive requests. This is most likely due to DHCP_DDNS configuration parameters referring to resources such as an IP address or port, that are unavailable. DHCP_DDNS will attempt to restart the queue manager if given a new configuration.

26.17.63 DHCP_DDNS_QUEUE_MGR_STOPPED

```
application's queue manager has stopped listening for requests.
```

Logged at debug log level 40. This is a debug message indicating that DHCP_DDNS's Queue Manager has stopped listening for NameChangeRequests. This may be because of normal event such as reconfiguration or as a result of an error. There should be log messages preceding this one to indicate why it has stopped.

26.17.64 DHCP_DDNS_QUEUE_MGR_STOPPING

```
application is stopping the queue manager for %1
```

Logged at debug log level 0. This is an informational message indicating that DHCP_DDNS is stopping the queue manager either to reconfigure it or as part of application shutdown.

26.17.65 DHCP_DDNS_QUEUE_MGR_STOP_ERROR

```
application encountered an error stopping the queue manager: %1
```

This is an error message indicating that DHCP_DDNS encountered an error while trying to stop the queue manager. This error is unlikely to occur or to impair the application's ability to function but it should be reported for analysis.

26.17.66 DHCP_DDNS_QUEUE_MGR_UNEXPECTED_HANDLER_ERROR

```
application's queue manager request receive handler experienced an unexpected exception  
↪ %1:
```

This is an error message indicating that an unexpected error occurred within the DHCP_DDNS's Queue Manager request receive completion handler. This is most likely a programmatic issue that should be reported. The application may recover on its own.

26.17.67 DHCP_DDNS_QUEUE_MGR_UNEXPECTED_STOP

```
application's queue manager receive was
```

aborted unexpectedly while queue manager state is: %1 This is an error message indicating that DHCP_DDNS's Queue Manager request receive was unexpected interrupted. Normally, the read is receive is only interrupted as a normal part of stopping the queue manager. This is most likely a programmatic issue that should be reported.

26.17.68 DHCP_DDNS_REMOVE_FAILED

```
DHCP_DDNS Request ID %1: Transaction outcome: %2
```

This is an error message issued after DHCP_DDNS attempts to submit DNS mapping entry removals have failed. The precise reason for the failure should be documented in preceding log entries.

26.17.69 DHCP_DDNS_REMOVE_SUCCEEDED

DHCP_DDNS Request ID %1: successfully removed the DNS mapping addition for this request:
↪ %2

This is an informational message issued after DHCP_DDNS has submitted DNS mapping removals which were received and accepted by an appropriate DNS server.

26.17.70 DHCP_DDNS_REQUEST_DROPPED

Request ID %1: Request contains no enabled update requests and will be dropped: %2

Logged at debug log level 55. This is a debug message issued when DHCP_DDNS receives a request which does not contain updates in a direction that is enabled. In other words, if only forward updates are enabled and request is received that asks only for reverse updates then the request is dropped.

26.17.71 DHCP_DDNS_REVERSE_REMOVE_BAD_DNSCLIENT_STATUS

DHCP_DDNS Request ID %1: received an unknown DNSClient status: %2, while removing
↪ reverse address mapping for FQDN %3 to DNS server %4

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was removing a reverse address mapping. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.72 DHCP_DDNS_REVERSE_REMOVE_BUILD_FAILURE

DNS Request ID %1: update message to remove a reverse DNS entry could not be constructed.
↪ from this request: %2, reason: %3

This is an error message issued when an error occurs attempting to construct the server bound packet requesting a reverse PTR removal. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue.

26.17.73 DHCP_DDNS_REVERSE_REMOVE_IO_ERROR

DHCP_DDNS Request ID %1: encountered an IO error sending a reverse mapping remove for
↪ FQDN %2 to DNS server %3

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a reverse mapping remove. The application will retry against the same server or others as appropriate.

26.17.74 DHCP_DDNS_REVERSE_REMOVE_REJECTED

DNS Request ID %1: Server, %2, rejected a DNS update request to remove the reverse
↪ mapping for FQDN, %3, with an RCODE: %4

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.75 DHCP_DDNS_REVERSE_REMOVE_RESP_CORRUPT

DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while removing reverse address mapping for FQDN, %3

This is an error message issued when the response received by DHCP_DDNS, to a update request to remove a reverse address, is mangled or malformed. The application will retry against the same server or others as appropriate.

26.17.76 DHCP_DDNS_REVERSE_REMOVE_TIMEOUT

DHCP_DDNS Request ID %1: timed out waiting for a response to reverse mapping remove for FQDN %2 to DNS server %3

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a reverse mapping remove. The application will retry against the same server or others as appropriate.

26.17.77 DHCP_DDNS_REVERSE_REPLACE_BAD_DNSCLIENT_STATUS

DHCP_DDNS Request ID %1: received an unknown DNSClient status: %2, while replacing reverse address mapping for FQDN %3 to DNS server %4

This is an error message issued when DNSClient returns an unrecognized status while DHCP_DDNS was replacing a reverse address mapping. The request will be aborted. This is most likely a programmatic issue and should be reported.

26.17.78 DHCP_DDNS_REVERSE_REPLACE_BUILD_FAILURE

DNS Request ID %1: update message to replace a reverse DNS entry could not be constructed from this request: %2, reason: %3

This is an error message issued when an error occurs attempting to construct the server bound packet requesting a reverse PTR replacement. This is due to invalid data contained in the NameChangeRequest. The request will be aborted. This is most likely a configuration issue.

26.17.79 DHCP_DDNS_REVERSE_REPLACE_IO_ERROR

DHCP_DDNS Request ID %1: encountered an IO error sending a reverse mapping replacement for FQDN %2 to DNS server %3

This is an error message issued when a communication error occurs while DHCP_DDNS is carrying out a reverse mapping replacement. The application will retry against the same server or others as appropriate.

26.17.80 DHCP_DDNS_REVERSE_REPLACE_REJECTED

DNS Request ID %1: Server, %2, rejected a DNS update request to replace the reverse mapping for FQDN, %3, with an RCODE: %4

This is an error message issued when an update was rejected by the DNS server it was sent to for the reason given by the RCODE. The rcode values are defined in RFC 2136.

26.17.81 DHCP_DDNS_REVERSE_REPLACE_RESP_CORRUPT

DHCP_DDNS Request ID %1: received a corrupt response from the DNS server, %2, while
↪ replacing reverse address mapping for FQDN, %3

This is an error message issued when the response received by DHCP_DDNS, to a update request to replace a reverse address, is mangled or malformed. The application will retry against the same server or others as appropriate.

26.17.82 DHCP_DDNS_REVERSE_REPLACE_TIMEOUT

DHCP_DDNS Request ID %1: timed out waiting for a response to reverse mapping replacement
↪ for FQDN %2 to DNS server %3

This is an error message issued when no response is received from the DNS server before exceeding dns-server-timeout while DHCP_DDNS is carrying out a reverse mapping replacement. The application will retry against the same server or others as appropriate.

26.17.83 DHCP_DDNS_REV_REQUEST_IGNORED

Request ID %1: Reverse updates are disabled, the reverse portion of request will be
↪ ignored: %2

Logged at debug log level 55. This is a debug message issued when reverse DNS updates are disabled and DHCP_DDNS receives an update request containing a reverse DNS update. The reverse update will not performed.

26.17.84 DHCP_DDNS_RUN_EXIT

application is exiting the event loop

Logged at debug log level 0. This is a debug message issued when the DHCP-DDNS server exits its event lo

26.17.85 DHCP_DDNS_SECURITY_CHECKS_DISABLED

Invoked with command line option -X, Security checks are disabled!!

This warning is emitted when internal security checks normally performed by kea-dhcp-ddns have been disabled via command line option '-X'. This means the server is not enforcing restrictions on resource paths or permissions. This mode of operation may expose your environment to security vulnerabilities and should only be used after consideration.

26.17.86 DHCP_DDNS_SHUTDOWN_COMMAND

application received shutdown command with args: %1

Logged at debug log level 0. This is a debug message issued when the application has been instructed to shut down by the controller.

26.17.87 DHCP_DDNS_STARTED

Kea DHCP-DDNS server version %1 started

This informational message indicates that the DHCP-DDNS server has processed all configuration information and is ready to begin processing. The version is also printed.

26.17.88 DHCP_DDNS_STARTING_TRANSACTION

Request ID %1:

Logged at debug log level 50. This is a debug message issued when DHCP-DDNS has begun a transaction for a given request.

26.17.89 DHCP_DDNS_STATE_MODEL_UNEXPECTED_ERROR

Request ID %1: application encountered an unexpected error while carrying out a `NameChangeRequest`: %2

This is error message issued when the application fails to process a `NameChangeRequest` correctly. Some or all of the DNS updates requested as part of this update did not succeed. This is a programmatic error and should be reported.

26.17.90 DHCP_DDNS_TRANS_SEND_ERROR

Request ID %1: application encountered an unexpected error while attempting to send a `DNS update`: %2

This is error message issued when the application is able to construct an update message but the attempt to send it suffered an unexpected error. This is most likely a programmatic error, rather than a communications issue. Some or all of the DNS updates requested as part of this request did not succeed.

26.17.91 DHCP_DDNS_UDP_SENDER_WATCH_SOCKET_CLOSE_ERROR

watch socket failed to close: %1

This is an error message that indicates the application was unable to close the inbound or outbound side of a NCR sender's watch socket. While technically possible the error is highly unlikely to occur and should not impair the application's ability to process requests.

26.17.92 DHCP_DDNS_UNCAUGHT_NCR_RECV_HANDLER_ERROR

unexpected exception thrown from the application receive completion handler: %1

This is an error message that indicates that an exception was thrown but not caught in the application's request receive completion handler. This is a programmatic error that needs to be reported. Dependent upon the nature of the error the application may or may not continue operating normally.

26.17.93 DHCP_DDNS_UNCAUGHT_NCR_SEND_HANDLER_ERROR

unexpected exception thrown from the DHCP-DDNS client send completion handler: %1

This is an error message that indicates that an exception was thrown but not caught in the application's send completion handler. This is a programmatic error that needs to be reported. Dependent upon the nature of the error the client may or may not continue operating normally.

26.17.94 DHCP_DDNS_UPDATE_REQUEST_SENT

Request ID %1: %2 to server: %3

Logged at debug log level 50. This is a debug message issued when DHCP_DDNS sends a DNS request to a DNS server.

26.17.95 DHCP_DDNS_UPDATE_RESPONSE_RECEIVED

```
Request ID %1: to server: %2 status: %3
```

Logged at debug log level 50. This is a debug message issued when DHCP_DDNS receives sends a DNS update response from a DNS server.

26.18 EVAL

26.18.1 EVAL_DEBUG_AND

```
%1: Popping %2 and %3 pushing %4
```

Logged at debug log level 55. This debug message indicates that two values are popped from the value stack. Then are then combined via logical and and the result is pushed onto the value stack.

26.18.2 EVAL_DEBUG_BRANCH

```
Branching to %1
```

Logged at debug log level 55. This debug message indicates that an unconditional branch is performed to the displayed target.

26.18.3 EVAL_DEBUG_CONCAT

```
%1: Popping %2 and %3 pushing %4
```

Logged at debug log level 55. This debug message indicates that the two strings are being popped off of the stack. They are then concatenated and the resulting string is pushed onto the stack. The strings are displayed in hex.

26.18.4 EVAL_DEBUG_EQUAL

```
%1: Popping %2 and %3 pushing result %4
```

Logged at debug log level 55. This debug message indicates that the two strings are being popped off of the value stack and the result of comparing them is being pushed onto the value stack. The strings are displayed in hex.

26.18.5 EVAL_DEBUG_HEXSTRING

```
%1: Pushing hex string %2
```

Logged at debug log level 55. This debug message indicates that the given binary string is being pushed onto the value stack. The string is displayed in hex.

26.18.6 EVAL_DEBUG_IFELSE_FALSE

```
%1: Popping %2 (false) and %3, leaving %4
```

Logged at debug log level 55. This debug message indicates that the condition is false so the iftrue branch value is removed and the ifelse branch value is left on the value stack.

26.18.7 EVAL_DEBUG_IFELSE_TRUE

%1: Popping %2 (true) and %3, leaving %4

Logged at debug log level 55. This debug message indicates that the condition is true so the ifelse branch value is removed and the iftrue branch value is left on the value stack.

26.18.8 EVAL_DEBUG_INT16TOTEXT

%1: Pushing Int16 %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents a 16 bit integer.

26.18.9 EVAL_DEBUG_INT32TOTEXT

%1: Pushing Int32 %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents a 32 bit integer.

26.18.10 EVAL_DEBUG_INT8TOTEXT

%1: Pushing Int8 %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents an 8 bit integer.

26.18.11 EVAL_DEBUG_IPADDRESS

%1: Pushing IPAddress %2

Logged at debug log level 55. This debug message indicates that the given binary string is being pushed onto the value stack. This represents either an IPv4 or IPv6 address. The string is displayed in hex.

26.18.12 EVAL_DEBUG_IPADDRESSTOTEXT

%1: Pushing IPAddress %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents either an IPv4 or IPv6 address.

26.18.13 EVAL_DEBUG_LCASE

%1: Popping string %2 and pushing converted value to lower case %3

Logged at debug log level 55. This debug message indicates that the given string representation is being converted to lower case and pushed onto the value stack.

26.18.14 EVAL_DEBUG_MATCH

```
Matching '%1' on %2, result %3
```

Logged at debug log level 55. This debug message indicates that the given regular expression was matched with the popped value. The result was pushed onto the value stack.

26.18.15 EVAL_DEBUG_MATCH_ERROR

```
Matching '%1' on %2 raised an error: %3
```

This error message indicates that an error occurred while evaluating the given regular expression against the popped value.

26.18.16 EVAL_DEBUG_MEMBER

```
%1: Checking membership of '%2', pushing result %3
```

Logged at debug log level 55. This debug message indicates that the membership of the packet for the client class was checked.

26.18.17 EVAL_DEBUG_NOT

```
%1: Popping %2 pushing %3
```

Logged at debug log level 55. This debug message indicates that the first value is popped from the value stack, negated and then pushed onto the value stack. The string is displayed in text.

26.18.18 EVAL_DEBUG_OPTION

```
%1: Pushing option %2 with value %3
```

Logged at debug log level 55. This debug message indicates that the given string representing the value of the requested option is being pushed onto the value stack. The string may be the text or binary value of the string based on the representation type requested (.text or .hex) or "true" or "false" if the requested type is .exists. The option code may be for either an option or a sub-option as requested in the classification statement.

26.18.19 EVAL_DEBUG_OR

```
%1: Popping %2 and %3 pushing %4
```

Logged at debug log level 55. This debug message indicates that two values are popped from the value stack. Then are then combined via logical or and the result is pushed onto the value stack. The string is displayed in text.

26.18.20 EVAL_DEBUG_PKT

```
%1: Pushing PKT meta data %2 with value %3
```

Logged at debug log level 55. This debug message indicates that the given binary string representing the value of the requested meta data is being pushed onto the value stack. The string is displayed in hex at the exception of interface name.

26.18.21 EVAL_DEBUG_PKT4

%1: Pushing PKT4 field %2 with value %3

Logged at debug log level 55. This debug message indicates that the given binary string representing the value of the requested field is being pushed onto the value stack. The string is displayed in hex.

26.18.22 EVAL_DEBUG_PKT6

%1: Pushing PKT6 field %2 with value %3

Logged at debug log level 55. This debug message indicates that the given binary string representing the value of the requested field is being pushed onto the value stack. The string is displayed in hex.

26.18.23 EVAL_DEBUG_POP_AND_BRANCH_FALSE

Value is false: branching to %1

Logged at debug log level 55. This debug message indicates that a branch on false condition is performed to the displayed target.

26.18.24 EVAL_DEBUG_POP_OR_BRANCH_FALSE

Value is false: keeping it and branching to %1

Logged at debug log level 55. This debug message indicates that a branch on false condition is performed to the displayed target.

26.18.25 EVAL_DEBUG_POP_OR_BRANCH_TRUE

Value is true: keeping it and branching to %1

Logged at debug log level 55. This debug message indicates that a branch on true condition is performed to the displayed target.

26.18.26 EVAL_DEBUG_RELAY6

%1: Pushing PKT6 relay field %2 nest %3 with value %4

Logged at debug log level 55. This debug message indicates that the given binary string representing the value of the requested field is being pushed onto the value stack. The string is displayed in hex.

26.18.27 EVAL_DEBUG_RELAY6_RANGE

%1: Pushing PKT6 relay field %2 nest %3 with value %4

Logged at debug log level 55. This debug message is generated if the nest field is out of range. The empty string will always be the value pushed onto the stack.

26.18.28 EVAL_DEBUG_SPLIT

%1: Popping field %2, delimiters %3, string %4, pushing result %5

Logged at debug log level 55. This debug message indicates that three values are being popped from the stack and a result is being pushed onto the stack. The values being popped are the field, delimiter and string. The result is the extracted field which is pushed onto the stack. The strings are displayed in hex.

26.18.29 EVAL_DEBUG_SPLIT_DELIM_EMPTY

%1: Popping field %2, delimiters %3, string %4, pushing result %5

Logged at debug log level 55. This debug message indicates that the delimiter popped from the stack was empty and so the result will be the entire string. The field, delimiter and string are still popped from the stack and the result is still pushed.

26.18.30 EVAL_DEBUG_SPLIT_EMPTY

%1: Popping field %2, delimiters %3, string %4, pushing result %5

Logged at debug log level 55. This debug message indicates that the string popped from the stack was empty and so the result will also be empty. The field, delimiter and string are still popped from the stack and the result is still pushed.

26.18.31 EVAL_DEBUG_SPLIT_FIELD_OUT_OF_RANGE

%1: Popping field %2, delimiters %3, string %4, pushing result %5

Logged at debug log level 55. This debug message indicates that the field is either less than one or larger than the number of fields in the string popped from the stack. The result will be empty. The field, delimiter and string are still popped from the stack and the result is still pushed.

26.18.32 EVAL_DEBUG_STRING

%1: Pushing text string %2

Logged at debug log level 55. This debug message indicates that the given text string is being pushed onto the value stack. The string is displayed in text.

26.18.33 EVAL_DEBUG_SUBSTRING

%1: Popping length %2, start %3, string %4 pushing result %5

Logged at debug log level 55. This debug message indicates that three values are being popped from the value stack and a result is being pushed onto the value stack. The values being popped are the starting point and length of a substring to extract from the given string. The resulting string is pushed onto the stack. The strings are displayed in hex.

26.18.34 EVAL_DEBUG_SUBSTRING_EMPTY

%1: Popping length %2, start %3, string %4 pushing result %5

Logged at debug log level 55. This debug message indicates that the string popped from the stack was empty and so the result will also be empty. The start, length and string are still popped from the stack and the result is still pushed.

26.18.35 EVAL_DEBUG_SUBSTRING_RANGE

%1: Popping length %2, start %3, string %4 pushing result %5

Logged at debug log level 55. This debug message indicates that the value of start is outside of the string and an empty result will be pushed onto the stack. The start, length and string are still popped from the stack and the result is still pushed. The strings are displayed in hex.

26.18.36 EVAL_DEBUG_SUB_OPTION

%1: Pushing option %2 sub-option %3 with value %4

This debug message indicates that the given string representing the value of the requested sub-option of the requested parent option is being pushed onto the value stack. The string may be the text or binary value of the string based on the representation type requested (.text or .hex) or "true" or "false" if the requested type is .exists. The codes are the parent option and the sub-option codes as requested in the classification statement.

26.18.37 EVAL_DEBUG_SUB_OPTION_NO_OPTION

%1: Requested option %2 sub-option %3, but the parent option is not present, pushing  result %4

This debug message indicates that the parent option was not found. The codes are the parent option and the sub-option codes as requested in the classification statement.

26.18.38 EVAL_DEBUG_TOHEXSTRING

%1: Popping binary value %2 and separator %3, pushing result %4

Logged at debug log level 55. This debug message indicates that two values are being popped from the value stack and a result is being pushed onto the value stack. The values being popped are the binary value to convert and the separator. The binary value is converted to its hexadecimal string representation and pushed onto the stack. The binary value is displayed in hex.

26.18.39 EVAL_DEBUG_UCASE

%1: Popping string %2 and pushing converted value to upper case %3

Logged at debug log level 55. This debug message indicates that the given string representation is being converted to upper case and pushed onto the value stack.

26.18.40 EVAL_DEBUG_UINT16TOTEXT

%1: Pushing UInt16 %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents a 16 bit unsigned integer.

26.18.41 EVAL_DEBUG_UINT32TOTEXT

%1: Pushing UInt32 %2

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents a 32 bit unsigned integer.

26.18.42 EVAL_DEBUG_UINT8TOTEXT

```
%1: Pushing UInt8 %2
```

Logged at debug log level 55. This debug message indicates that the given address string representation is being pushed onto the value stack. This represents an 8 bit unsigned integer.

26.18.43 EVAL_DEBUG_VENDOR_CLASS_DATA

```
%1: Data %2 (out of %3 received) in vendor class found, pushing result '%4'
```

Logged at debug log level 55. This debug message indicates that vendor class option was found and passed enterprise-id checks and has sufficient number of data chunks. The total number of chunks and value pushed are reported as debugging aid.

26.18.44 EVAL_DEBUG_VENDOR_CLASS_DATA_NOT_FOUND

```
%1: Requested data index %2, but option with enterprise-id %3 has only %4 data tuple(s),  
↪ pushing result '%5'
```

Logged at debug log level 55. This debug message indicates that vendor class option was found and passed enterprise-id checks, but does not have sufficient number of data chunks. Note that the index starts at 0, so there has to be at least (index + 1) data chunks.

26.18.45 EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID

```
%1: Pushing enterprise-id %2 as result 0x%3
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor class option was found and its enterprise-id is being reported.

26.18.46 EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID_MISMATCH

```
%1: Was looking for %2, option had %3, pushing result '%4'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor class option was found, but has different enterprise-id than specified in the expression.

26.18.47 EVAL_DEBUG_VENDOR_CLASS_EXISTS

```
%1: Option with enterprise-id %2 found, pushing result '%3'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor class option was found.

26.18.48 EVAL_DEBUG_VENDOR_CLASS_NO_OPTION

```
%1: Option with code %2 missing, pushing result '%3'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor class option was not found.

26.18.49 EVAL_DEBUG_VENDOR_ENTERPRISE_ID

```
%1: Pushing enterprise-id %2 as result 0x%3
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor option was found and its enterprise-id is being reported.

26.18.50 EVAL_DEBUG_VENDOR_ENTERPRISE_ID_MISMATCH

```
%1: Was looking for %2, option had %3, pushing result '%4'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor option was found, but has different enterprise-id than specified in the expression.

26.18.51 EVAL_DEBUG_VENDOR_EXISTS

```
%1: Option with enterprise-id %2 found, pushing result '%3'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor option was found.

26.18.52 EVAL_DEBUG_VENDOR_NO_OPTION

```
%1: Option with code %2 missing, pushing result '%3'
```

Logged at debug log level 55. This debug message indicates that the expression has been evaluated and vendor option was not found.

26.19 FLEX

26.19.1 FLEX_ID_EXPRESSION_EMPTY

```
Specified identifier-expression is empty
```

This warning message is printed when the flex-id library is being loaded, but the expression used to generate the identifier is empty. The library will load, but will not generate any identifiers. Please make sure that the identifier-expression parameter is specified.

26.19.2 FLEX_ID_EXPRESSION_EVALUATED

```
Expression evaluated for packet to "%1" (size: %2)
```

Logged at debug log level 40. This debug message is printed every time a packet evaluation is successful. This means that the identifier expression has been generated. Note that depending on the expression and content of the incoming packet, the expression may be evaluated to an empty string.

26.19.3 FLEX_ID_EXPRESSION_EVALUATED_NP

```
Expression evaluated for packet to 0x%1 (size: %2)
```

This debug message is printed every time a packet evaluation is successful. This means that the identifier expression has been generated. As it is not printable it is converted to hexadecimal.

26.19.4 FLEX_ID_EXPRESSION_HEX

```
evaluated expression in hexadecimal form "%1"
```

Logged at debug log level 40. This debug message provides a hexadecimal representation of the evaluated expression. This is useful for debugging purposes because further logs use hexadecimal format for logging.

26.19.5 FLEX_ID_EXPRESSION_INVALID_JSON_TYPE

```
The identifier-expression is %1, but expected JSON string
```

This error message is printed when the flex-id library is being loaded, but the expression used to generate the identifier is malformed. It has a different JSON type (e.g. is a map) rather than expected string.

26.19.6 FLEX_ID_EXPRESSION_NOT_DEFINED

```
Expression (identifier-expression) is not defined.
```

This warning message is printed when the flex-id library is loaded, but the expression used to generate the identifier is not specified. The library will load, but will not generate any identifiers. Please make sure that the identifier-expression parameter is specified for your library.

26.19.7 FLEX_ID_EXPRESSION_PARSE_FAILED

```
The identifier-expression is [%1], but fails to parse with error: %2
```

This error message is printed when the flex-id library is being loaded, but the expression used to generate the identifier is malformed. It failed to parse.

26.19.8 FLEX_ID_IGNORE_IAID_APPLIED_ON_NA

```
the ignore-iaid has changed IAID (%1) to 0 for the IA_NA option.
```

Logged at debug log level 40. This DEBUG message is printed when the flex-id library did apply the ignore-iaid flag and changed IAID to 0 because the received packet contains exactly one IA_NA.

26.19.9 FLEX_ID_IGNORE_IAID_APPLIED_ON_PD

```
the ignore-iaid has changed IAID (%1) to 0 for the IA_PD option.
```

Logged at debug log level 40. This DEBUG message is printed when the flex-id library did apply the ignore-iaid flag and changed IAID to 0 because the received packet contains exactly one IA_PD.

26.19.10 FLEX_ID_IGNORE_IAID_ENABLED

```
the ignore-iaid is set. It only has an effect on clients with at most one IA_NA and one ↵  
↵ IA_PD.
```

This WARNING message is printed when the flex-id library is being loaded, and the ignore-iaid parameter is set. This flag will enable the server to drop packets which contain more than one IA_NA and one IA_PD.

26.19.11 FLEX_ID_IGNORE_IAID_JSON_TYPE

the ignore-iaid is %1 but expected boolean value

This error message is printed when the flex-id library is being loaded, but the ignore-iaid parameter is malformed, i.e. it has a different type than expected. It is expected to be a boolean value.

26.19.12 FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_NA

the ignore-iaid was not applied on the packet because it contains more than one IA_NA.

Logged at debug log level 40. This WARNING message is printed when the flex-id library did not apply the ignore-iaid flag because the received packet contains more than one IA_NA.

26.19.13 FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_PD

the ignore-iaid was not applied on the packet because it contains more than one IA_PD.

Logged at debug log level 40. This WARNING message is printed when the flex-id library did not apply the ignore-iaid flag because the received packet contains more than one IA_PD.

26.19.14 FLEX_ID_LOAD_ERROR

An error occurred loading the library %1

This error message will be printed when an error is encountered during loading of the library. Details of the problem are printed as parameter to this message.

26.19.15 FLEX_ID_REPLACE_CLIENT_ID_JSON_TYPE

the replace-client-id is %1 but expected boolean value

This error message is printed when the flex-id library is being loaded, but the replace-client-id parameter is malformed, i.e. it has a different type than expected. It is expected to be a boolean value.

26.19.16 FLEX_ID_RESTORE_CLIENT_ID

restoring original client identifier '%1' in the response

Logged at debug log level 40. This debug message is issued when the original (client supplied) client identifier is restored into the server's response.

26.19.17 FLEX_ID_RESTORE_DUID

restoring original DUID "%1" in the response

Logged at debug log level 40. This debug message is issued when the original (client supplied) client identifier is restored into the server's response.

26.19.18 FLEX_ID_UNLOAD

```
Flex-id library has been unloaded.
```

This informational message signifies that the flexible-identifier library has been unloaded.

26.19.19 FLEX_ID_USED_AS_CLIENT_ID

```
using flexible identifier "%1" as client identifier
```

Logged at debug log level 40. This debug message is issued to indicate that the library is removing client supplied client identifier from the received message and is inserting flexible identifier based client identifier instead. The server will use this new client identifier for processing the packet. The original client identifier will be restored in the pkt4_send callout and sent back to the client.

26.19.20 FLEX_ID_USED_AS_DUID

```
using flexible identifier "%1" as DUID
```

Logged at debug log level 40. This debug message is issued to indicate that the library is removing client supplied DUID from the received message and is inserting flexible identifier based DUID instead. The server will use this new DUID for processing the packet. The original DUID will be restored in the pkt6_send callout and sent back to the client.

26.19.21 FLEX_OPTION_LOAD_ERROR

```
loading Flex Option hooks library failed: %1
```

This error message indicates an error during loading the Flex Option hooks library. The details of the error are provided as argument of the log message.

26.19.22 FLEX_OPTION_PROCESS_ADD

```
Added the option code %1 with value %2
```

Logged at debug log level 40. This debug message is printed when an option was added into the response packet. The option code and the value (between quotes if printable, in hexadecimal if not) are provided.

26.19.23 FLEX_OPTION_PROCESS_CLIENT_CLASS

```
Skip processing of the option code %1 for class '%2'
```

Logged at debug log level 40. This debug message is printed when the processing for an option is skipped because the query does not belongs to the client class. The option code and the client class name are provided.

26.19.24 FLEX_OPTION_PROCESS_ERROR

```
An error occurred processing query %1: %2
```

This error message indicates an error during processing of a query by the Flex Option hooks library. The client identification information from the query and the details of the error are provided as arguments of the log message.

26.19.25 FLEX_OPTION_PROCESS_REMOVE

Removed option code %1

Logged at debug log level 40. This debug message is printed when an option was removed from the response packet. The option code is provided.

26.19.26 FLEX_OPTION_PROCESS_SUB_ADD

Added the sub-option code %1 in option code %2 with value %3

Logged at debug log level 40. This debug message is printed when an sub-option was added into the response packet. The sub-option and container option codes, and the value (between quotes if printable, in hexadecimal if not) are provided.

26.19.27 FLEX_OPTION_PROCESS_SUB_CLIENT_CLASS

Skip processing of the sub-option code %1 in option code %2 for class '%3'

Logged at debug log level 40. This debug message is printed when the processing for a sub-option is skipped because the query does not belongs to the client class. The sub-option and container option codes, and the client class name are provided.

26.19.28 FLEX_OPTION_PROCESS_SUB_REMOVE

Removed sub-option code %1 in option code %2

Logged at debug log level 40. This debug message is printed when a sub-option was removed from the response packet. The sub-option and container option codes are provided.

26.19.29 FLEX_OPTION_PROCESS_SUB_SUPERSEDE

Supersedes the sub-option code %1 in option code %2 with value %3

Logged at debug log level 40. This debug message is printed when a sub-option was superseded into the response packet. The sub-option and container option codes, and the value (between quotes if printable, in hexadecimal if not) are provided.

26.19.30 FLEX_OPTION_PROCESS_SUPERSEDE

Supersedes the option code %1 with value %2

Logged at debug log level 40. This debug message is printed when an option was superseded into the response packet. The option code and the value (between quotes if printable, in hexadecimal if not) are provided.

26.19.31 FLEX_OPTION_PROCESS_VENDOR_ID_MISMATCH

Skip processing of vendor option code %1 with vendor id %2 not matching wanted %3

Logged at debug log level 40. This debug message is printed when a sub-option of a vendor option is processed but vendor ids do not match. The code of the vendor option and the two vendor ids are provided.

26.19.32 FLEX_OPTION_UNLOAD

```
Flex Option hooks library has been unloaded
```

This info message indicates that the Flex Option hooks library has been unloaded.

26.20 FUZZ

26.20.1 FUZZ_DATA_READ

```
read %1 byte(s) from AFL via stdin
```

Logged at debug log level 50. A debug message output to indicate how much data has been received from the fuzzer via stdin

26.20.2 FUZZ_INIT_COMPLETE

```
fuzz initialization complete: interface %1, address %2, port %3, max loops %4
```

An informational message output when the fuzzing initialization function has completed successfully. The parameters listed are those which must be/can be set via environment variables.

26.20.3 FUZZ_INIT_FAIL

```
fuzz initialization failure, reason: %1
```

An error message reported if the fuzzing initialization failed. The reason for the failure is given in the message.

26.20.4 FUZZ_READ_FAIL

```
error reading input from fuzzer: %1
```

This error is reported if the read of data from the fuzzer (which is received over stdin) fails, or if a read returns zero bytes. If this occurs, the thread will sleep for a short period before retrying the read. The message includes the reason for the failure.

26.20.5 FUZZ_SEND

```
sent %1 byte(s) to the socket connected to the Kea interface
```

Logged at debug log level 50. A debug message stating that the sendto() call in the main fuzzing function has successfully completed and reporting the number of bytes sent. This call sends data received from AFL to the port on which Kea is listening.

26.20.6 FUZZ_SEND_ERROR

```
failed to send data to Kea input socket: %1
```

This error will be reported if the sendto() call in the fuzzing thread (which sends data received from AFL to the socket on which Kea is listening) fails. The reason for the failure is given in the message. The fuzzing code will attempt to continue from this, but it may cause the fuzzing process to fail.

26.20.7 FUZZ_SHORT_SEND

```
expected to send %1 bytes to Kea input socket but only sent %2
```

A warning message that is output if the `sendto()` call (used to send data from the fuzzing thread to the main Kea processing) did not send as much data as that read from AFL. This may indicate a problem in the underlying communications between the fuzzing thread and the main Kea processing.

26.20.8 FUZZ_SOCKET_CREATE_FAIL

```
failed to crease socket for use by fuzzing thread: %1
```

An error message output when the fuzzing code has failed to create a socket through which it will copy data received on stdin from the AFL fuzzer to the port on which Kea is listening. The program will most likely hang if this occurs.

26.21 GSS

26.21.1 GSS_TSIG_COMMAND_PROCESSED_FAILED

```
command_processed callout failed: %1.
```

This error message is issued when the callout for the `command_processed` callout point failed. The argument contains a reason for the error.

26.21.2 GSS_TSIG_LOAD_FAILED

```
GSS-TSIG hooks library failed to load: %1.
```

This error message indicates that an error occurred attempting to load the GSS-TSIG hooks library. The argument details the error.

26.21.3 GSS_TSIG_LOAD_OK

```
GSS-TSIG hooks library loaded successfully.
```

This info message indicates that the GSS-TSIG hooks library has been loaded successfully.

26.21.4 GSS_TSIG_MANAGER_STARTED

```
hooks library GSS-TSIG key periodic manager started.
```

Logged at debug log level 40. This debug message is issued when the GSS-TSIG key periodic manager has started.

26.21.5 GSS_TSIG_MANAGER_STOPPED

```
hooks library GSS-TSIG key periodic manager stopped.
```

Logged at debug log level 40. This debug message is issued when the GSS-TSIG key periodic manager has stopped.

26.21.6 GSS_TSIG_MANAGER_STOP_ERROR

```
manager stop error: %1
```

This error message is issued when the GSS-TSIG key periodic manager has stopped but an error is detected. The error message in the argument gives details about the problem.

26.21.7 GSS_TSIG_MANAGER_STOP_GENERAL_ERROR

```
manager stop general error
```

This error message is issued when the GSS-TSIG key periodic manager has stopped but a general error is detected.

26.21.8 GSS_TSIG_NEW_KEY

```
new GSS-TSIG key '%1' was created.
```

Logged at debug log level 40. This info message indicates that the GSS-TSIG hooks library has created a new GSS-TSIG key. The name of the new key is displayed.

26.21.9 GSS_TSIG_NEW_KEY_SETUP_FAILED

```
new GSS-TSIG key '%1' setup failed: %2.
```

This warning message is issued when the setup of a new GSS-TSIG key failed. The name of the new key and the error are displayed.

26.21.10 GSS_TSIG_NEW_KEY_SETUP_SUCCEED

```
new GSS-TSIG key '%1' setup succeed.
```

Logged at debug log level 40. This debug message is issued when the setup of a new GSS-TSIG key successfully finished. The name of the new key is displayed.

26.21.11 GSS_TSIG_OLD_KEY_REMOVED

```
%1 old GSS-TSIG keys were removed
```

Logged at debug log level 40. This debug message is issued when some old keys (older than 2 times the maximum TKEY lifetime) were removed. The number of removed keys is displayed.

26.21.12 GSS_TSIG_UNLOAD_OK

```
GSS-TSIG hooks library unloaded successfully.
```

This info message indicates that the GSS-TSIG hooks library has been unloaded successfully.

26.21.13 GSS_TSIG_VERIFIED

```
GSS-TSIG verify succeeded.
```

Logged at debug log level 40. A debug message issued when GSS-TSIG verification succeeded.

26.21.14 GSS_TSIG_VERIFY_FAILED

```
GSS-TSIG verify failed: %1.
```

This info message indicates that GSS-TSIG verification failed. The argument details the error.

26.22 HA

26.22.1 HA_BUFFER4_RECEIVE_FAILED

```
buffer4_receive callout failed: %1
```

This error message is issued when the callout for the `buffer4_receive` hook point failed. This may occur as a result of an internal server error. The argument contains a reason for the error.

26.22.2 HA_BUFFER4_RECEIVE_NOT_FOR_US

```
%1: dropping query to be processed by another server
```

Logged at debug log level 40. This debug message is issued when the received DHCPv4 query is dropped by this server because it should be served by another server. This is the case when the remote server was designated to process the packet as a result of load balancing or because it is a primary server in the hot standby configuration. The argument provides client identification information retrieved from the query.

26.22.3 HA_BUFFER4_RECEIVE_PACKET_OPTIONS_SKIPPED

```
an error unpacking an option, caused subsequent options to be skipped: %1
```

Logged at debug log level 40. A debug message issued when an option failed to unpack correctly, making it impossible to unpack the remaining options in the DHCPv4 query. The server will still attempt to service the packet. The sole argument provides a reason for unpacking error.

26.22.4 HA_BUFFER4_RECEIVE_UNPACK_FAILED

```
failed to parse query from %1 to %2, received over interface %3, reason: %4
```

Logged at debug log level 40. This debug message is issued when received DHCPv4 query is malformed and can't be parsed by the `buffer4_receive` callout. The query will be dropped by the server. The first three arguments specify source IP address, destination IP address and the interface. The last argument provides a reason for failure.

26.22.5 HA_BUFFER6_RECEIVE_FAILED

```
buffer6_receive callout failed: %1
```

This error message is issued when the callout for the `buffer6_receive` hook point failed. This may occur as a result of an internal server error. The argument contains a reason for the error.

26.22.6 HA_BUFFER6_RECEIVE_NOT_FOR_US

```
%1: dropping query to be processed by another server
```


Logged at debug log level 40. This debug message is issued when the received DHCPv6 query is dropped by this server because it should be served by another server. This is the case when the remote server was designated to process the packet as a result of load balancing or because it is a primary server in the hot standby configuration. The argument provides client identification information retrieved from the query.

26.22.7 HA_BUFFER6_RECEIVE_PACKET_OPTIONS_SKIPPED

```
an error unpacking an option, caused subsequent options to be skipped: %1
```

Logged at debug log level 40. A debug message issued when an option failed to unpack correctly, making it impossible to unpack the remaining options in the DHCPv6 query. The server will still attempt to service the packet. The sole argument provides a reason for unpacking error.

26.22.8 HA_BUFFER6_RECEIVE_UNPACK_FAILED

```
failed to parse query from %1 to %2, received over interface %3, reason: %4
```

Logged at debug log level 40. This debug message is issued when received DHCPv6 query is malformed and can't be parsed by the buffer6_receive callout. The query will be dropped by the server. The first three arguments specify source IP address, destination IP address and the interface. The last argument provides a reason for failure.

26.22.9 HA_COMMAND_PROCESSED_FAILED

```
command_processed callout failed: %1
```

This error message is issued when the callout for the command_processed hook point failed. The argument contains a reason for the error.

26.22.10 HA_COMMUNICATION_INTERRUPTED

```
%1: communication with %2 is interrupted
```

This warning message is issued by the server which discovered that the communication to the active partner has been interrupted for a time period longer than the configured heartbeat-delay time. At this stage the server starts the failover procedure by monitoring the DHCP traffic sent to the partner and checking whether the partner server responds to this traffic. If the max-unacked-clients value is set to 0 such verification is disabled in which case the server will transition to the partner-down state.

26.22.11 HA_COMMUNICATION_INTERRUPTED_CLIENT4

```
%1: new client %2 attempting to get a lease from the partner
```

This informational message is issued when the surviving server observes a DHCP packet sent to the partner with which the communication is interrupted. The client whose packet is observed is not yet considered "unacked" because the secs field value does not exceed the configured threshold specified with max-ack-delay.

26.22.12 HA_COMMUNICATION_INTERRUPTED_CLIENT4_UNACKED

```
%1: partner server failed to respond to %2, %3 clients unacked so far, %4 clients left,  
→ before transitioning to the partner-down state
```

This informational message is issued when the surviving server determines that its partner failed to respond to the DHCP query and that this client is considered to not be served by the partner. The surviving server counts such clients

and if the number of such clients exceeds the max-unacked-clients threshold, the server will transition to the partner-down state. The first argument specifies the relationship name. The second argument contains client identification information. The third argument specifies the number of clients to which the server has failed to respond. The fourth argument specifies the number of additional clients which, if not provisioned, will cause the server to transition to the partner-down state.

26.22.13 HA_COMMUNICATION_INTERRUPTED_CLIENT6

```
%1: new client %2 attempting to get a lease from the partner
```

This informational message is issued when the surviving server observes a DHCP packet sent to the partner with which the communication is interrupted. The client whose packet is observed is not yet considered "unacked" because the elapsed time option value does not exceed the configured threshold specified with max-ack-delay. The sole argument specifies client identification information.

26.22.14 HA_COMMUNICATION_INTERRUPTED_CLIENT6_UNACKED

```
%1: partner server failed to respond to %2, %3 clients unacked so far, %4 clients left.
↳ before transitioning to the partner-down state
```

This informational message is issued when the surviving server determines that its partner failed to respond to the DHCP query and that this client is considered to not be served by the partner. The surviving server counts such clients and if the number of such clients exceeds the max-unacked-clients threshold, the server will transition to the partner-down state. The first argument specifies the relationship name. The second argument contains client identification information. The third argument specifies the number of clients to which the server has failed to respond. The fourth argument specifies the number of additional clients which, if not provisioned, will cause the server to transition to the partner-down state.

26.22.15 HA_CONFIGURATION_FAILED

```
failed to configure High Availability hooks library: %1
```

This error message is issued when there is an error configuring the HA hooks library. The argument provides the detailed error message.

26.22.16 HA_CONFIGURATION_SUCCESSFUL

```
HA hook library has been successfully configured
```

This informational message is issued when the HA hook library configuration parser successfully parses and validates the new configuration.

26.22.17 HA_CONFIG_AUTO_FAILOVER_DISABLED

```
%1: auto-failover disabled
```

This warning message is issued to indicate that the 'auto-failover' parameter was administratively disabled for the specified server. The server will not automatically start serving partner's scope when the partner failure is detected. The server administrator will need to enable this scope manually by sending appropriate ha-scopes command.

26.22.18 HA_CONFIG_DHCP_MT_DISABLED

%1: HA multi-threading has been disabled, it cannot be enabled when Kea global multi-threading is disabled

This informational message is issued when HA configuration has enabled multi-threading while Kea global configuration has multi-threading disabled.

26.22.19 HA_CONFIG_DHCP_MT_DISABLED_AND_KEA_MT_ENABLED

%1: HA multi-threading is disabled while Kea global multi-threading is enabled which most likely cause performance degradation.

This warning message is issued when HA configuration has disabled multi-threading while Kea global configuration has multi-threading enabled. This will likely cause performance degradation.

26.22.20 HA_CONFIG_LEASE_SYNCING_DISABLED

%1: lease database synchronization between HA servers is disabled

This warning message is issued when the lease database synchronization is administratively disabled. This is valid configuration if the leases are replicated between lease databases via some other mechanism, e.g. SQL database replication.

26.22.21 HA_CONFIG_LEASE_SYNCING_DISABLED_REMINDER

%1: bypassing SYNCING state because lease database synchronization is administratively disabled

This informational message is issued as a reminder that lease database synchronization is administratively disabled and therefore the server transitions directly from the "waiting" to "ready" state.

26.22.22 HA_CONFIG_LEASE_UPDATES_AND_SYNCING_DIFFER

%1: unusual configuration where "send-lease-updates": %2 and "sync-leases": %3

This warning message is issued when the configuration values of the send-lease-updates and sync-leases parameters differ. This may be a valid configuration but is unusual. Normally, if the lease database with replication is in use, both values are set to false. If a lease database without replication is in use (e.g. memfile), both values are set to true. Providing different values for those parameters means that an administrator either wants the server to not synchronize leases upon startup but later send lease updates to the partner, or the lease database should be synchronized upon startup, but no lease updates are later sent as a result of leases allocation.

26.22.23 HA_CONFIG_LEASE_UPDATES_DISABLED

%1: lease updates will not be generated

This warning message is issued when the lease updates are administratively disabled. This is valid configuration if the leases are replicated to the partner's database via some other mechanism, e.g. SQL database replication.

26.22.24 HA_CONFIG_LEASE_UPDATES_DISABLED_REMINDER

```
%1: lease updates are administratively disabled and will not be generated while in %2_
↪state
```

This informational message is issued as a reminder that the lease updates are administratively disabled and will not be issued in the HA state to which the server has transitioned. The sole argument specifies the state into which the server has transitioned.

26.22.25 HA_CONFIG_SYSTEM_MT_UNSUPPORTED

```
%1: HA multi-threading has been disabled, auto-detection of thread support reports 0
```

This informational message is issued when HA multi-threading configuration has specified auto-detection for the number of threads to use and the system reports the number of concurrent threads as 0. If you know your system can support multiple threads, then you may override this condition by specifying explicit values for `http-listener-threads` and `http-client-threads`.

26.22.26 HA_CONTINUE_HANDLER_FAILED

```
ha-continue command failed: %1
```

This error message is issued to indicate that the `ha-continue` command handler failed while processing the command. The argument provides the reason for failure.

26.22.27 HA_DEINIT_OK

```
unloading High Availability hooks library successful
```

This informational message indicates that the High Availability hooks library has been unloaded successfully.

26.22.28 HA_DHCP4_START_SERVICE_FAILED

```
failed to start DHCPv4 HA services in dhcp4_srv_configured callout: %1
```

This error message is issued when an attempt to start High Availability services for the DHCPv4 server failed in the `dhcp4_srv_configured` callout. This is internal server error and a bug report should be created.

26.22.29 HA_DHCP6_START_SERVICE_FAILED

```
failed to start DHCPv6 HA services in dhcp6_srv_configured callout: %1
```

This error message is issued when an attempt to start High Availability services for the DHCPv6 server failed in the `dhcp6_srv_configured` callout. This is internal server error and a bug report should be created.

26.22.30 HA_DHCP_DISABLE_COMMUNICATIONS_FAILED

```
%1: failed to send request to disable DHCP service of %2: %3
```

This warning message indicates that there was a problem in communication with a HA peer while sending the `dhcp-disable` command. The first argument specifies the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.31 HA_DHCP_DISABLE_FAILED

```
%1: failed to disable DHCP service of %2: %3
```

This warning message indicates that a peer returned an error status code in response to a dhcp-disable command. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.32 HA_DHCP_ENABLE_COMMUNICATIONS_FAILED

```
%1: failed to send request to enable DHCP service of %2: %3
```

This warning message indicates that there was a problem in communication with a HA peer while sending the dhcp-enable command. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.33 HA_DHCP_ENABLE_FAILED

```
%1: failed to enable DHCP service of %2: %3
```

This warning message indicates that a peer returned an error status code in response to a dhcp-enable command. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.34 HA_HEARTBEAT_COMMUNICATIONS_FAILED

```
%1: failed to send heartbeat to %2: %3
```

This warning message indicates that there was a problem in communication with a HA peer while sending a heartbeat. This is a first sign that the peer may be down. The server will keep trying to send heartbeats until it considers that communication is interrupted.

26.22.35 HA_HEARTBEAT_FAILED

```
%1: heartbeat to %2 failed: %3
```

This warning message indicates that a peer returned an error status code in response to a heartbeat. This is the sign that the peer may not function properly. The server will keep trying to send heartbeats until it considers that communication is interrupted.

26.22.36 HA_HEARTBEAT_HANDLER_FAILED

```
heartbeat command failed: %1
```

This error message is issued to indicate that the heartbeat command handler failed while processing the command. The argument provides the reason for failure.

26.22.37 HA_HIGH_CLOCK_SKEW

```
%1: %2, please synchronize clocks!
```

This warning message is issued when the clock skew between the active servers exceeds 30 seconds. The HA service continues to operate but may not function properly, especially for low lease lifetimes. The administrator should synchronize the clocks, e.g. using NTP. If the clock skew exceeds 60 seconds, the HA service will terminate.

26.22.38 HA_HIGH_CLOCK_SKEW_CAUSED_TERMINATION

```
%1: %2, causing HA service to terminate
```

This warning message is issued when the clock skew between the active servers exceeds 60 seconds. The HA service stops. The servers will continue to respond to the DHCP queries but won't exchange lease updates or send heartbeats. The administrator is required to synchronize the clocks and then restart the servers to resume the HA service.

26.22.39 HA_INIT_OK

```
loading High Availability hooks library successful
```

This informational message indicates that the High Availability hooks library has been loaded successfully. Enjoy!

26.22.40 HA_INVALID_PARTNER_STATE_COMMUNICATION_RECOVERY

```
%1: partner is in the communication-recovery state unexpectedly
```

This warning message is issued when a partner is in the communication-recovery state, and this server is not running in the load balancing mode. The server may only transition to the communication-recovery state when it runs in the load balancing mode. The HA mode of both servers must be the same.

26.22.41 HA_INVALID_PARTNER_STATE_HOT_STANDBY

```
%1: partner is in the hot-standby state unexpectedly
```

This warning message is issued when a partner is in the hot-standby state, and this server is not running in the hot standby mode. The server may only transition to the hot-standby state when it runs in the hot standby mode. The HA mode of both servers must be the same.

26.22.42 HA_INVALID_PARTNER_STATE_LOAD_BALANCING

```
%1: partner is in the load-balancing state unexpectedly
```

This warning message is issued when a partner is in the load-balancing state, and this server is not running in the load balancing mode. The server may only transition to the load-balancing state when it runs in the load balancing mode. The HA mode of both servers must be the same.

26.22.43 HA_LEASE4_EXPIRE_FAILED

```
lease4_expire callout failed: %1
```

This error message is issued when the callout for the lease4_expire hook point failed. This includes unexpected errors like wrong arguments provided to the callout by the DHCP server (unlikely internal server error). The argument contains a reason for the error.

26.22.44 HA_LEASE4_EXPIRE_INVALID_HA_SERVER_NAME

```
%1: invalid ha-server-name value for subnet %2
```

This error message is issued when the reclaimed DHCPv4 lease belongs to a subnet which includes ha-server-name value in the user-context but this value is not a string or is empty. It is a server's misconfiguration. The first argument holds the lease information. The second argument is a subnet prefix.

26.22.45 HA_LEASE4_EXPIRE_RECLAMATION_SKIP

```
%1: skipping reclamation of the lease that belongs to a partner
```

Logged at debug log level 40. This debug message is issued when the server is in the terminated state and skips reclamation of the lease that was probably allocated by another server, or is maintained by the other server while the servers are in the HA terminated state. The argument is the lease address.

26.22.46 HA_LEASE4_SERVER_DECLINE_FAILED

```
lease4_server_decline callout failed: %1
```

This error message is issued when the callout for the lease4_server_decline hook point failed. This includes unexpected errors like wrong arguments provided to the callout by the DHCP server (unlikely internal server error). The argument contains a reason for the error.

26.22.47 HA_LEASE6_EXPIRE_FAILED

```
lease4_expire callout failed: %1
```

This error message is issued when the callout for the lease4_expire hook point failed. This includes unexpected errors like wrong arguments provided to the callout by the DHCP server (unlikely internal server error). The argument contains a reason for the error.

26.22.48 HA_LEASE6_EXPIRE_INVALID_HA_SERVER_NAME

```
%1: invalid ha-server-name value for subnet %2
```

This error message is issued when the reclaimed DHCPv6 lease belongs to a subnet which includes ha-server-name value in the user-context but this value is not a string or is empty. It is a server's misconfiguration. The first argument holds the lease information. The second argument is a subnet prefix.

26.22.49 HA_LEASE6_EXPIRE_RECLAMATION_SKIP

```
%1: skipping reclamation of the lease that belongs to a partner
```

Logged at debug log level 40. This debug message is issued when the server is in the terminated state and skips reclamation of the lease that was probably allocated by another server, or is maintained by the other server while the servers are in the HA terminated state. The argument is the lease address.

26.22.50 HA_LEASES4_COMMITTED_FAILED

```
leases4_committed callout failed: %1
```

This error message is issued when the callout for the leases4_committed hook point failed. This includes unexpected errors like wrong arguments provided to the callout by the DHCP server (unlikely internal server error). The argument contains a reason for the error.

26.22.51 HA_LEASES4_COMMITTED_NOTHING_TO_UPDATE

```
%1: leases4_committed callout was invoked without any leases
```

Logged at debug log level 40. This debug message is issued when the "leases4_committed" callout returns because there are neither new leases nor deleted leases for which updates should be sent. The sole argument specifies the details of the client which sent the packet.

26.22.52 HA_LEASES4_COMMITTED_NO_RELATIONSHIP

```
%1: HA relationship not found: %2
```

This error message is issued when the relationship for the server name provided by the earlier callouts was not found in the HA configuration. This error is highly unlikely and rather indicates some programming error. The first argument is the client identification information. The second argument holds a more detailed error message.

26.22.53 HA_LEASES6_COMMITTED_FAILED

```
leases6_committed callout failed: %1
```

This error message is issued when the callout for the leases6_committed hook point failed. This includes unexpected errors like wrong arguments provided to the callout by the DHCP server (unlikely internal server error). The argument contains a reason for the error.

26.22.54 HA_LEASES6_COMMITTED_NOTHING_TO_UPDATE

```
%1: leases6_committed callout was invoked without any leases
```

Logged at debug log level 40. This debug message is issued when the "leases6_committed" callout returns because there are neither new leases nor deleted leases for which updates should be sent. The sole argument specifies the details of the client which sent the packet.

26.22.55 HA_LEASES6_COMMITTED_NO_RELATIONSHIP

```
%1: HA relationship not found: %2
```

This error message is issued when the relationship for the server name provided by the earlier callouts was not found in the HA configuration. This error is highly unlikely and rather indicates some programming error. The first argument is the client identification information. The second argument holds a more detailed error message.

26.22.56 HA_LEASES_BACKLOG_COMMUNICATIONS_FAILED

```
%1: failed to communicate with %2 while sending lease updates backlog: %3
```

This error message is issued to indicate that there was a communication error with a partner server while sending outstanding lease updates after resuming connection. The third argument contains a reason for the error.

26.22.57 HA_LEASES_BACKLOG_FAILED

```
%1: failed to send lease updates backlog to %2: %3
```

This error message is issued to indicate that sending lease updates backlog to a partner server failed. The lease updates backlog is sent to the partner after resuming temporarily broken communication with the partner. If this operation fails the server will transition to the waiting state to initiate full lease database synchronization.

26.22.58 HA_LEASES_BACKLOG_NOTHING_TO_SEND

%1: no leases in backlog after communication recovery

This informational message is issued when there are no outstanding leases to be sent after communication recovery with a partner. This means that the communication interruption was short enough that no DHCP clients obtained any leases from the server while it was in the communication-recovery state. The server may now transition to the load-balancing state.

26.22.59 HA_LEASES_BACKLOG_START

%1: starting to send %2 outstanding lease updates to %3

This informational message is issued when the server starts to send outstanding lease updates to the partner after resuming communications. The first argument specifies the local server's name. The second argument specifies the number of lease updates to be sent. The name of the partner is specified with the third argument.

26.22.60 HA_LEASES_BACKLOG_SUCCESS

%1: sending lease updates backlog to %2 successful in %3

This informational message is issued when server successfully completes sending lease updates backlog to the partner. The first argument specifies the local server's name. The second argument specifies the name of the remote server. The third argument specifies the duration of this operation.

26.22.61 HA_LEASES_SYNC_APPLIED_LEASES

%1: applied %2 leases received from the partner in the local lease database

This informational message outputs the number of leases received from the partner during the database synchronization and applied in the local database. A typical case when only some leases are applied is when the server has multiple relationships and some of the received leases belong to another relationship. The first argument specifies this server name. The second argument specifies the number of applied leases.

26.22.62 HA_LEASES_SYNC_COMMUNICATIONS_FAILED

%1: failed to communicate with %2 while syncing leases: %3

This error message is issued to indicate that there was a communication error with a partner server while trying to fetch leases from its lease database. The argument contains a reason for the error.

26.22.63 HA_LEASES_SYNC_FAILED

%1: failed to synchronize leases with %2: %3

This error message is issued to indicate that there was a problem while parsing a response from the server from which leases have been fetched for local database synchronization. The third argument contains a reason for the error.

26.22.64 HA_LEASES_SYNC_LEASE_PAGE_RECEIVED

%1: received %2 leases from %3

This informational message is issued during lease database synchronization to indicate that a bulk of leases have been received. The first argument specifies the local server's name. The second argument holds the count of leases received. The third argument specifies the partner server name.

26.22.65 HA_LEASE_SYNC_FAILED

```
%1: synchronization failed for lease: %2, reason: %3
```

This warning message is issued when creating or updating a lease in the local lease database fails. The lease information in the JSON format is provided as a first argument. The third argument provides a reason for the failure.

26.22.66 HA_LEASE_SYNC_STALE_LEASE4_SKIP

```
%1: skipping stale lease %2 in subnet %3
```

Logged at debug log level 40. This debug message is issued during lease database synchronization, when fetched IPv4 lease instance appears to be older than the instance in the local database. The newer instance is left in the database and the fetched lease is dropped. The remote server will still hold the older lease instance until it synchronizes its database with this server. The first argument specifies the local server's name. The second argument specifies leased address. The third argument specifies a subnet to which the lease belongs.

26.22.67 HA_LEASE_SYNC_STALE_LEASE6_SKIP

```
%1: skipping stale lease %2 in subnet %3
```

Logged at debug log level 40. This debug message is issued during lease database synchronization, when fetched IPv6 lease instance appears to be older than the instance in the local database. The newer instance is left in the database and the fetched lease is dropped. The remote server will still hold the older lease instance until it synchronizes its database with this server. The first argument specifies the local server's name. The second argument specifies leased address. The second argument specifies a subnet to which the lease belongs.

26.22.68 HA_LEASE_UPDATES_DISABLED

```
%1: lease updates will not be sent to the partner while in %2 state
```

This informational message is issued to indicate that lease updates will not be sent to the partner while the server is in the current state. The second argument specifies the server's current state name. The lease updates are still sent to the backup servers if they are configured but any possible errors in communication with the backup servers are ignored.

26.22.69 HA_LEASE_UPDATES_ENABLED

```
%1: lease updates will be sent to the partner while in %2 state
```

This informational message is issued to indicate that lease updates will be sent to the partner while the server is in the current state. The second specifies the server's current state name.

26.22.70 HA_LEASE_UPDATE_COMMUNICATIONS_FAILED

```
%1: failed to send lease update %2 to %3: %4
```

This warning message indicates that there was a problem in communication with a HA peer while processing a DHCP client query and sending lease update. The client's DHCP message will be dropped.

26.22.71 HA_LEASE_UPDATE_CONFLICT

```
%1: lease update %2 sent to %3 returned conflict status code: %4
```

This warning message indicates that the partner returned a conflict status code in response to a lease update. The client's DHCP message will be dropped. If the server is configured to track conflicting lease updates, it may eventually transition to the terminated state when the configured threshold is exceeded.

26.22.72 HA_LEASE_UPDATE_CREATE_UPDATE_FAILED_ON_PEER

```
%1: failed to create or update the lease having type %2 for address %3, reason: %4
```

This informational message is issued when one of the leases failed to be created or updated on the HA peer while processing the lease updates sent from this server. This may indicate an issue with communication between the peer and its lease database.

26.22.73 HA_LEASE_UPDATE_DELETE_FAILED_ON_PEER

```
%1: failed to delete the lease having type %2 for address %3, reason: %4
```

This informational message is issued when one of the leases failed to delete on the HA peer while processing lease updates sent from this server. Typically, the lease fails to delete when it doesn't exist in the peer's database.

26.22.74 HA_LEASE_UPDATE_FAILED

```
%1: lease update %2 sent to %3 failed: %4
```

This warning message indicates that a peer returned an error status code in response to a lease update. The client's DHCP message will be dropped.

26.22.75 HA_LEASE_UPDATE_REJECTS_CAUSED_TERMINATION

```
%1: too many rejected lease updates cause the HA service to terminate
```

This error message is issued when the HA service terminates because the number of lease updates for which a conflict status code was returned by the partner exceeds the limit set with max-rejected-lease-updates configuration parameter.

26.22.76 HA_LOAD_BALANCING_DUID_MISSING

```
%1: load balancing failed for the DHCPv6 message (transaction id: %2) because DUID is ↵  
↵missing
```

Logged at debug log level 40. This debug message is issued when the HA hook library was unable to load balance an incoming DHCPv6 query because neither client identifier nor HW address was included in the query. The query will be dropped. The sole argument contains transaction id.

26.22.77 HA_LOAD_BALANCING_IDENTIFIER_MISSING

```
%1: load balancing failed for the DHCPv4 message (transaction id: %2) because HW address ↵  
↵and client identifier are missing
```

Logged at debug log level 40. This debug message is issued when the HA hook library was unable to load balance an incoming DHCPv4 query because neither client identifier nor HW address was included in the query. The query will be dropped. The sole argument contains transaction id.

26.22.78 HA_LOAD_BALANCING_LEASE_DUID_MISSING

```
%1: load balancing failed for the DHCPv6 lease %2 because DUID is missing
```

Logged at debug log level 40. This debug message is issued when the HA hook library was unable to load balance a reclaimed DHCPv6 lease because client identifier was not included found in the lease.

26.22.79 HA_LOAD_BALANCING_LEASE_IDENTIFIER_MISSING

```
%1: load balancing failed for the DHCPv4 lease %2 because HW address and client_
↪ identifier are missing
```

Logged at debug log level 40. This debug message is issued when the HA hook library was unable to load balance a reclaimed DHCPv4 lease because neither client identifier nor HW address was included in the query.

26.22.80 HA_LOCAL_DHCP_DISABLE

```
local DHCP service is disabled while the %1 is in the %2 state
```

This informational message is issued to indicate that the local DHCP service is disabled because the server remains in a state in which the server should not respond to DHCP clients, e.g. the server hasn't synchronized its lease database. The first argument specifies server name. The second argument specifies server's state.

26.22.81 HA_LOCAL_DHCP_ENABLE

```
local DHCP service is enabled while the %1 is in the %2 state
```

This informational message is issued to indicate that the local DHCP service is enabled because the server remains in a state in which it should respond to the DHCP clients. The first argument specifies server name. The second argument specifies server's state.

26.22.82 HA_MAINTENANCE_CANCEL_HANDLER_FAILED

```
ha-maintenance-cancel command failed: %1
```

This error message is issued to indicate that the ha-maintenance-cancel command handler failed while processing the command. The argument provides the reason for failure.

26.22.83 HA_MAINTENANCE_NOTIFY_CANCEL_COMMUNICATIONS_FAILED

```
%1: failed to send ha-maintenance-notify to %2 in attempt to cancel its maintenance: %3
```

This warning message indicates that there was a problem in communication with a HA peer while sending the ha-maintenance-notify command with the cancel flag set to true. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.84 HA_MAINTENANCE_NOTIFY_CANCEL_FAILED

```
%1: error returned while processing ha-maintenance-notify by %2 in attempt to cancel its_
↪ maintenance: %3
```

This warning message indicates that a peer returned an error status code in response to a ha-maintenance-notify command with the cancel flag set to true. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.85 HA_MAINTENANCE_NOTIFY_COMMUNICATIONS_FAILED

```
%1: failed to send ha-maintenance-notify to %2: %3
```

This warning message indicates that there was a problem in communication with a HA peer while sending the ha-maintenance-notify command. The first argument provides the local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.86 HA_MAINTENANCE_NOTIFY_FAILED

```
%1: error returned while processing ha-maintenance-notify by %2: %3
```

This warning message indicates that a peer returned an error status code in response to a ha-maintenance-notify command. The first argument provides the remote server's name. The second argument provides a reason for failure.

26.22.87 HA_MAINTENANCE_NOTIFY_HANDLER_FAILED

```
ha-maintenance-notify command failed: %1
```

This error message is issued to indicate that the ha-maintenance-notify command handler failed while processing the command. The argument provides the reason for failure.

26.22.88 HA_MAINTENANCE_SHUTDOWN_SAFE

```
%1: the server can now be shutdown for maintenance as the partner has taken over the  
↪DHCP traffic
```

This informational message is displayed after the server transitions to the in-maintenance state. This server no longer responds to any DHCP queries and its partner - in partner-in-maintenance state - has taken over the DHCP traffic. When the server in-maintenance state is shut down, the partner moves to the partner-down state immediately.

26.22.89 HA_MAINTENANCE_STARTED

```
%1: the server is now in the partner-in-maintenance state and the partner is in-  
↪maintenance state
```

This informational message is displayed when the server receiving the ha-maintenance-start command transitions to the partner-in-maintenance state. The server does it after sending the ha-maintenance-notify to its partner to put the partner in the in-maintenance state. From now on, the server in the partner-in-maintenance state will be responding to all queries and the partner will respond to no queries. The partner may be safely shut down for maintenance in which case this server will automatically transition from the partner-in-maintenance state to the partner-down state.

26.22.90 HA_MAINTENANCE_STARTED_IN_PARTNER_DOWN

```
%1: the server is now in the partner-down mode as a result of requested maintenance
```

This informational message is displayed when the server receiving the ha-maintenance-start command transitions to the partner-down state because it was unable to communicate with the partner while receiving the command. It is assumed that in such situation the partner is already offline for the maintenance. Note that in this case the normal failover procedure does not take place. The server does not wait for a heartbeat to fail several times, nor it monitors the DHCP traffic for not responded queries. In the maintenance case the server transitions to the partner-down state when it first encounters a communication problem with the partner.

26.22.91 HA_MAINTENANCE_START_HANDLER_FAILED

```
ha-maintenance-start command failed: %1
```

This error message is issued to indicate that the ha-maintenance-start command handler failed while processing the command. The argument provides the reason for failure.

26.22.92 HA_MISSING_CONFIGURATION

```
high-availability parameter not specified for High Availability hooks library
```

This error message is issued to indicate that the configuration for the High Availability hooks library hasn't been specified. The 'high-availability' parameter must be specified for the hooks library to load properly.

26.22.93 HA_PAUSE_CLIENT_LISTENER_FAILED

```
%1: pausing multi-threaded HTTP processing failed: %2
```

This error message is emitted when attempting to pause HA's HTTP client and listener threads. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.22.94 HA_PAUSE_CLIENT_LISTENER_ILLEGAL

```
%1: pausing multi-threaded HTTP processing failed: %2
```

This error message is emitted when attempting to pause HA's HTTP client or listener thread pools from a worker thread. This error indicates that a command run on the listener threads is trying to use a critical section which would result in a dead-lock.

26.22.95 HA_RESET_COMMUNICATIONS_FAILED

```
%1: failed to send ha-reset command to %2: %3
```

This warning message indicates a problem with communication with a HA peer while sending the ha-reset command. The first argument specifies the local server name. The second argument specifies a remote server name. The third argument specifies a reason for failure.

26.22.96 HA_RESET_FAILED

```
%1: failed to reset HA state machine of %2: %3
```

This warning message indicates that a peer returned an error status code in response to the ha-reset command. The first argument specifies a local server name. The second argument specifies a remote server name. The third argument specifies a reason for failure.

26.22.97 HA_RESET_HANDLER_FAILED

```
ha-reset command failed: %1
```

This error message is issued to indicate that the ha-reset command handler failed while processing the command. The argument provides the reason for failure.

26.22.98 HA_RESUME_CLIENT_LISTENER_FAILED

```
%1: resuming multi-threaded HTTP processing failed: %2
```

This error message is emitted when attempting to resume HA's HTTP client and listener threads. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.22.99 HA_SCOPES_HANDLER_FAILED

```
ha-scopes command failed: %1
```

This error message is issued to indicate that the ha-scopes command handler failed while processing the command. The argument provides reason for the failure.

26.22.100 HA_SERVICE_STARTED

```
%1: started high availability service in %2 mode as %3 server
```

This informational message is issued when the HA service is started as a result of server startup or reconfiguration. The first argument specifies a local server name. The second argument provides the HA mode. The third argument specifies the role of this server instance in this configuration.

26.22.101 HA_STATE_MACHINE_CONTINUED

```
%1: state machine is un-paused
```

This informational message is issued when the HA state machine is un-paused. This unlocks the server from the current state. It may transition to any other state if it needs to do so, e.g. 'partner-down' if its partner appears to be offline. The server may also remain in the current state if the HA setup state warrants such behavior.

26.22.102 HA_STATE_MACHINE_PAUSED

```
%1: state machine paused in state %2
```

This informational message is issued when the HA state machine is paused. HA state machine may be paused in certain states specified in the HA hooks library configuration. When the state machine is paused, the server remains in the given state until it is explicitly unpaused (via the ha-continue command). If the state machine is paused, the server operates normally but cannot transition to any other state.

26.22.103 HA_STATE_TRANSITION

```
%1: server transitions from %2 to %3 state, partner state is %4
```

This informational message is issued when the server transitions to a new state as a result of some interaction (or lack of thereof) with its partner. The arguments specify local server name, initial server state, new server state and the partner's state.

26.22.104 HA_STATE_TRANSITION_PASSIVE_BACKUP

```
%1: server transitions from %2 to %3 state
```

This informational message is issued when the server in passive-backup mode transitions to a new state. The arguments specify local server name, initial server state and a new server state.

26.22.105 HA_SUBNET4_SELECT_FAILED

```
subnet4_select callout failed: %1
```

This error message is issued when the callout for the subnet4_select hook point failed. This may occur as a result of an internal server error. The argument contains a reason for the error.

26.22.106 HA_SUBNET4_SELECT_INVALID_HA_SERVER_NAME

```
%1: invalid ha-server-name value for subnet %2
```

This error message is issued when the received DHCPv4 query is dropped by this server because the specified ha-server-name value in the subnet's user-context has non-string type or is empty. It is a server's misconfiguration. The first argument is the client identification information. The second argument is a subnet prefix.

26.22.107 HA_SUBNET4_SELECT_NOT_FOR_US

```
%1: dropping query in relationship %2 to be processed by another server
```

Logged at debug log level 40. This debug message is issued when the received DHCPv4 query is dropped by this server because it should be served by another server. This is the case when a remote primary server is operational. The first argument is the client identification information. The second argument is the relationship name.

26.22.108 HA_SUBNET4_SELECT_NO_RELATIONSHIP_FOR_SUBNET

```
%1: HA relationship not found for %2
```

This error message is issued when the received DHCPv4 query is dropped by this server because the server could not find a relationship matching the specified ha-server-name for a subnet. The server name matches no relationship specified in the HA configuration. A typical reason for it is a typo. The first argument is the client identification information. The second argument is the relationship name.

26.22.109 HA_SUBNET4_SELECT_NO_RELATIONSHIP_SELECTOR_FOR_SUBNET

```
%1: unable to determine HA relationship because selected subnet %2 lacks the ha-server-  
↪name
```

This error message is issued when the received DHCPv4 query is dropped by this server because it was unable to determine the HA relationship to which the received query belongs. If there are multiple relationships, it is required to specify ha-server-name value in the user-context at the subnet or shared network level for each subnet and/or shared network. The server uses them as a relationship selector. If these selectors are unspecified for any of the subnets it is a configuration error. The first argument is the client identification information. The second argument is a subnet prefix.

26.22.110 HA_SUBNET4_SELECT_NO_SUBNET_SELECTED

```
%1: unable to determine HA relationship because no subnet has been selected for the  
↪client
```

Logged at debug log level 40. This debug message is issued when the received DHCPv4 query is dropped by this server because it could not select a subnet for this client. Selecting the subnet is required to find a suitable HA relationship. This message is not emitted when the server has only one relationship. The argument is the client identification information.

26.22.111 HA_SUBNET6_SELECT_FAILED

```
subnet6_select callout failed: %1
```

This error message is issued when the callout for the subnet6_select hook point failed. This may occur as a result of an internal server error. The argument contains a reason for the error.

26.22.112 HA_SUBNET6_SELECT_INVALID_HA_SERVER_NAME

```
%1: invalid ha-server-name value for subnet %2
```

This error message is issued when the received DHCPv6 query is dropped by this server because the specified ha-server-name value in the subnet's user-context has non-string type or is empty. It is a server's misconfiguration. The first argument is the client identification information. The second argument is a subnet prefix.

26.22.113 HA_SUBNET6_SELECT_NOT_FOR_US

```
%1: dropping query in relationship %2 to be processed by another server
```

Logged at debug log level 40. This debug message is issued when the received DHCPv6 query is dropped by this server because it should be served by another server. This is the case when a remote primary server is operational. The first argument is the client identification information. The second argument is the relationship name.

26.22.114 HA_SUBNET6_SELECT_NO_RELATIONSHIP_FOR_SUBNET

```
%1: HA relationship not found for %2
```

This error message is issued when the received DHCPv6 query is dropped by this server because the server could not find a relationship matching the specified ha-server-name for a subnet. The server name matches no relationship specified in the HA configuration. A typical reason for it is a typo. The first argument is the client identification information. The second argument is the relationship name.

26.22.115 HA_SUBNET6_SELECT_NO_RELATIONSHIP_SELECTOR_FOR_SUBNET

```
%1: unable to determine HA relationship because selected subnet %2 lacks the ha-server-  
↪name
```

This error message is issued when the received DHCPv6 query is dropped by this server because it was unable to determine the HA relationship to which the received query belongs. If there are multiple relationships, it is required to specify ha-server-name value in the user-context at the subnet or shared network level for each subnet and/or shared network. The server uses them as a relationship selector. If these selectors are unspecified for any of the subnets it is a configuration error. The first argument is the client identification information. The second argument is a subnet prefix.

26.22.116 HA_SUBNET6_SELECT_NO_SUBNET_SELECTED

```
%1: unable to determine HA relationship because no subnet has been selected for the  
↪client
```

Logged at debug log level 40. This debug message is issued when the received DHCPv6 query is dropped by this server because it could not select a subnet for this client. Selecting the subnet is required to find a suitable HA relationship. This message is not emitted when the server has only one relationship. The argument is the client identification information.

26.22.117 HA_SYNC_COMPLETE_NOTIFY_COMMUNICATIONS_FAILED

```
%1: failed to send ha-sync-complete-notify to %2: %3
```

This warning message indicates that there was a problem in communication with an HA peer while sending the ha-sync-complete-notify command. The first argument provides a local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.118 HA_SYNC_COMPLETE_NOTIFY_FAILED

```
%1: error processing ha-sync-complete-notify command on %2: %3
```

This warning message indicates that a peer returned an error status code in response to the ha-sync-complete-notify command. The first argument provides a local server's name. The second argument provides the remote server's name. The third argument provides a reason for failure.

26.22.119 HA_SYNC_COMPLETE_NOTIFY_HANDLER_FAILED

```
ha-sync-complete-notify command failed: %1
```

This error message is issued to indicate that the ha-sync-complete-notify command handler failed while processing the command. The argument provides the reason for failure.

26.22.120 HA_SYNC_FAILED

```
%1: lease database synchronization with %2 failed: %3
```

This error message is issued to indicate that the lease database synchronization failed. The first argument provides the local server's name. The second argument provides the partner server's name. The third argument provides a reason for the failure.

26.22.121 HA_SYNC_HANDLER_FAILED

```
ha-sync command failed: %1
```

This error message is issued to indicate that the ha-sync command handler failed while processing the command. The argument provides the reason for failure.

26.22.122 HA_SYNC_START

```
%1: starting lease database synchronization with %2
```

This informational message is issued when the server starts lease database synchronization with a partner. The arguments specify the local and remote server names.

26.22.123 HA_SYNC_SUCCESSFUL

```
%1: lease database synchronization with %2 completed successfully in %3
```

This informational message is issued when the server successfully completed lease database synchronization with the partner. The first argument specifies local server name. The second argument specifies the name of the partner server. The third argument specifies the duration of the synchronization.

26.22.124 HA_TERMINATED

```
HA %1: service terminated due to an unrecoverable condition. Check previous error_
↪message(s), address the problem and restart!
```

This error message is issued to indicate that the HA service has been stopped due to an unacceptable condition (e.g. too large of a clock skew). The exact cause should appear in a previous error message. Address the condition reported then restart the servers to resume service.

26.22.125 HA_TERMINATED_PARTNER_DID_NOT_RESTART

```
%1: service is terminating because the terminated partner was not restarted within %2_
↪minutes
```

This warning message is issued to indicate that the HA service is terminating because partner server is in the terminated state and was not restarted within an expected time frame. The terminated servers should be restarted after correcting the problem that caused the termination. They can be restarted sequentially but the duration between the restarts should not be too long. If it is long it may mean that the restart of one of the servers was unintentional (e.g., power outage). If the restarted server remains in the waiting state it cannot serve DHCP clients. Moving to the terminated state at least allows for responding to the DHCP traffic.

26.22.126 HA_TERMINATED_RESTART_PARTNER

```
%1: waiting for the partner in the TERMINATED state to be restarted
```

This informational message is issued when the server has been restarted after correcting the clock skew. The partner server is still in the terminated state. The partner must be restarted before the server can synchronize the database and start normal operation.

26.23 HOOKS

26.23.1 HOOKS_ALL_CALLOUTS_DEREGISTERED

```
hook library at index %1 removed all callouts on hook %2
```

Logged at debug log level 55. A debug message issued when all callouts on the specified hook registered by the library with the given index were removed. This is similar to the HOOKS_CALLOUTS_REMOVED message (and the two are likely to be seen together), but is issued at a lower-level in the hook framework.

26.23.2 HOOKS_CALLOUTS_BEGIN

```
begin all callouts for hook %1
```

Logged at debug log level 45. This debug message is issued when callout manager begins to invoke callouts for the hook. The argument specifies the hook name.

26.23.3 HOOKS_CALLOUTS_COMPLETE

```
completed callouts for hook %1 (total callouts duration: %2)
```

Logged at debug log level 45. This debug message is issued when callout manager has completed execution of all callouts for the particular hook. The arguments specify the hook name and total execution time for all callouts in milliseconds.

26.23.4 HOOKS_CALLOUTS_REMOVED

```
callouts removed from hook %1 for library %2
```

Logged at debug log level 45. This is a debug message issued during library unloading. It notes that one of more callouts registered by that library have been removed from the specified hook. This is similar to the `HOOKS_DEREGISTER_ALL_CALLOUTS` message (and the two are likely to be seen together), but is issued at a higher-level in the hook framework.

26.23.5 HOOKS_CALLOUT_CALLED

```
hooks library with index %1 has called a callout on hook %2 that has address %3 (callout_
↳duration: %4)
```

Logged at debug log level 55. Only output at a high debugging level, this message indicates that a callout on the named hook registered by the library with the given index (in the list of loaded libraries) has been called and returned a success state. The address of the callout is given in the message. The message includes the callout execution time in milliseconds.

26.23.6 HOOKS_CALLOUT_DEREGISTERED

```
hook library at index %1 deregistered a callout on hook %2
```

Logged at debug log level 55. A debug message issued when all instances of a particular callouts on the hook identified in the message that were registered by the library with the given index have been removed.

26.23.7 HOOKS_CALLOUT_ERROR

```
error returned by callout on hook %1 registered by library with index %2 (callout_
↳address %3) (callout duration %4)
```

If a callout returns an error status when called, this error message is issued. It identifies the hook to which the callout is attached, the index of the library (in the list of loaded libraries) that registered it and the address of the callout. The error is otherwise ignored. The error message includes the callout execution time in milliseconds.

26.23.8 HOOKS_CALLOUT_EXCEPTION

```
exception thrown by callout on hook %1 registered by library with index %2 (callout_
↳address %3): %4 (callout duration: %5)
```

If a callout throws an exception when called, this error message is issued. It identifies the hook to which the callout is attached, the index of the library (in the list of loaded libraries) that registered it and the address of the callout. The error is otherwise ignored. The error message includes the callout execution time in milliseconds.

26.23.9 HOOKS_CALLOUT_REGISTRATION

```
hooks library with index %1 registering callout for hook '%2'
```

Logged at debug log level 45. This is a debug message, output when a library (whose index in the list of libraries (being) loaded is given) registers a callout.

26.23.10 HOOKS_CLOSE_ERROR

```
failed to close hook library %1: %2
```

Kea has failed to close the named hook library for the stated reason. Although this is an error, this should not affect the running system other than as a loss of resources. If this error persists, you should restart Kea.

26.23.11 HOOKS_HOOK_LIST_RESET

```
the list of hooks has been reset
```

This is a message indicating that the list of hooks has been reset. While this is usual when running the Kea test suite, it should not be seen when running Kea in a production environment. If this appears, please report a bug through the usual channels.

26.23.12 HOOKS_INCORRECT_VERSION

```
hook library %1 is at version %2, require version %3
```

Kea has detected that the named hook library has been built against a version of Kea that is incompatible with the version of Kea running on your system. It has not loaded the library. This is most likely due to the installation of a new version of Kea without rebuilding the hook library. A rebuild and re-install of the library should fix the problem in most cases.

26.23.13 HOOKS_LIBRARY_CLOSED

```
hooks library %1 successfully closed
```

This information message is issued when a user-supplied hooks library has been successfully closed.

26.23.14 HOOKS_LIBRARY_LOADED

```
hooks library %1 successfully loaded
```

This information message is issued when a user-supplied hooks library has been successfully loaded.

26.23.15 HOOKS_LIBRARY_LOADING

```
loading hooks library %1
```

Logged at debug log level 40. This is a debug message output just before the specified library is loaded. If the action is successfully, it will be followed by the HOOKS_LIBRARY_LOADED informational message.

26.23.16 HOOKS_LIBRARY_MULTI_THREADING_COMPATIBLE

```
hooks library %1 reports its multi-threading compatibility as %2
```

Logged at debug log level 45. A debug message issued when the "multi_threading_compatible" function was called. The returned value (0 means not compatible, others compatible) is displayed.

26.23.17 HOOKS_LIBRARY_MULTI_THREADING_NOT_COMPATIBLE

```
hooks library %1 is not compatible with multi-threading
```

When multi-threading is enabled and the library is not compatible (either because the "multi_threading_compatible" function returned 0 or was not implemented) this error message is issued. The library must be removed from the configuration or the multi-threading disabled.

26.23.18 HOOKS_LIBRARY_UNLOADED

```
hooks library %1 successfully unloaded
```

This information message is issued when a user-supplied hooks library has been successfully unloaded.

26.23.19 HOOKS_LIBRARY_UNLOADING

```
unloading library %1
```

Logged at debug log level 40. This is a debug message called when the specified library is being unloaded. If all is successful, it will be followed by the HOOKS_LIBRARY_UNLOADED informational message.

26.23.20 HOOKS_LIBRARY_VERSION

```
hooks library %1 reports its version as %2
```

Logged at debug log level 45. A debug message issued when the version check on the hooks library has succeeded.

26.23.21 HOOKS_LOAD_ERROR

```
'load' function in hook library %1 returned error %2
```

A "load" function was found in the library named in the message and was called. The function returned a non-zero status (also given in the message) which was interpreted as an error. The library has been unloaded and no callouts from it will be installed.

26.23.22 HOOKS_LOAD_EXCEPTION

```
'load' function in hook library %1 threw an exception
```

A "load" function was found in the library named in the message and was called. The function threw an exception (an error indication) during execution, which is an error condition. The library has been unloaded and no callouts from it will be installed.

26.23.23 HOOKS_LOAD_FRAMEWORK_EXCEPTION

```
'load' function in hook library %1 threw an exception: reason %2
```

A "load" function was found in the library named in the message and was called. Either the hooks framework or the function threw an exception (an error indication) during execution, which is an error condition; the cause of the exception is recorded in the message. The library has been unloaded and no callouts from it will be installed.

26.23.24 HOOKS_LOAD_SUCCESS

```
'load' function in hook library %1 returned success
```

Logged at debug log level 40. This is a debug message issued when the "load" function has been found in a hook library and has been successfully called.

26.23.25 HOOKS_MULTI_THREADING_COMPATIBLE_EXCEPTION

```
'multi_threading_compatible' function in hook library %1 threw an exception
```

This error message is issued if the `multi_threading_compatible()` function in the specified hooks library was called and generated an exception. The library is considered unusable and will not be loaded.

26.23.26 HOOKS_NO_LOAD

```
no 'load' function found in hook library %1
```

Logged at debug log level 40. This is a debug message saying that the specified library was loaded but no function called "load" was found in it. Providing the library contained some "standard" functions (i.e. functions with the names of the hooks for the given server), this is not an issue.

26.23.27 HOOKS_NO_UNLOAD

```
no 'unload' function found in hook library %1
```

Logged at debug log level 40. This is a debug message issued when the library is being unloaded. It merely states that the library did not contain an "unload" function.

26.23.28 HOOKS_NO_VERSION

```
no 'version' function found in hook library %1
```

The shared library named in the message was found and successfully loaded, but Kea did not find a function named "version" in it. This function is required and should return the version of Kea against which the library was built. The value is used to check that the library was built against a compatible version of Kea. The library has not been loaded.

26.23.29 HOOKS_OPEN_ERROR

```
failed to open hook library %1: %2
```

Kea failed to open the specified hook library for the stated reason. The library has not been loaded. Kea will continue to function, but without the services offered by the library.

26.23.30 HOOKS_STD_CALLOUT_REGISTERED

```
hooks library %1 registered standard callout for hook %2 at address %3
```

Logged at debug log level 45. This is a debug message, output when the library loading function has located a standard callout (a callout with the same name as a hook point) and registered it. The address of the callout is indicated.

26.23.31 HOOKS_UNLOAD_ERROR

```
'unload' function in hook library %1 returned error %2
```

During the unloading of a library, an "unload" function was found. It was called, but returned an error (non-zero) status, resulting in the issuing of this message. The unload process continued after this message and the library has been unloaded.

26.23.32 HOOKS_UNLOAD_EXCEPTION

```
'unload' function in hook library %1 threw an exception
```

During the unloading of a library, an "unload" function was found. It was called, but in the process generated an exception (an error indication). The unload process continued after this message and the library has been unloaded.

26.23.33 HOOKS_UNLOAD_FRAMEWORK_EXCEPTION

```
'unload' function in hook library %1 threw an exception, reason %2
```

During the unloading of a library, an "unload" function was found. It was called, but in the process either it or the hooks framework generated an exception (an error indication); the cause of the error is recorded in the message. The unload process continued after this message and the library has been unloaded.

26.23.34 HOOKS_UNLOAD_SUCCESS

```
'unload' function in hook library %1 returned success
```

Logged at debug log level 40. This is a debug message issued when an "unload" function has been found in a hook library during the unload process, called, and returned success.

26.23.35 HOOKS_VERSION_EXCEPTION

```
'version' function in hook library %1 threw an exception
```

This error message is issued if the version() function in the specified hooks library was called and generated an exception. The library is considered unusable and will not be loaded.

26.24 HOSTS

26.24.1 HOSTS_BACKENDS_REGISTERED

```
the following host backend types are available: %1
```

This informational message lists all possible host backends that could be used in hosts-database[s].

26.24.2 HOSTS_BACKEND_DEREGISTER

```
deregistered host backend type: %1
```

Logged at debug log level 40. This debug message is issued when a backend factory was deregistered. It is no longer possible to use host backend of this type.

26.24.3 HOSTS_BACKEND_REGISTER

```
registered host backend type: %1
```

Logged at debug log level 40. This debug message is issued when a backend factory was successfully registered. It is now possible to use host backend of this type.

26.24.4 HOSTS_CFG_ADD_HOST

```
add the host for reservations: %1
```

Logged at debug log level 40. This debug message is issued when new host (with reservations) is added to the server's configuration. The argument describes the host and its reservations in detail.

26.24.5 HOSTS_CFG_CACHE_HOST_DATA_SOURCE

```
get host cache data source: %1
```

This informational message is issued when a host cache data source is detected by the host manager.

26.24.6 HOSTS_CFG_CLOSE_HOST_DATA_SOURCE

```
Closing host data source: %1
```

Logged at debug log level 40. This is a normal message being printed when the server closes host data source connection.

26.24.7 HOSTS_CFG_DEL

```
deleted %1 host(s) having %2 IPv6 reservation(s) for subnet id %3 and address %4
```

Logged at debug log level 40. This debug message is issued when reservations are deleted for the specified subnet and address. The first argument specifies how many hosts have been deleted. The second argument specifies how many reservations have been deleted. The third argument is the subnet identifier. The fourth argument is the IP address.

26.24.8 HOSTS_CFG_DEL4

```
deleted %1 host(s) for subnet id %2 and identifier %3
```

Logged at debug log level 40. This debug message is issued when IPv4 reservations are deleted for the specified subnet and identifier. The first argument specifies how many hosts have been deleted. The second argument is the subnet identifier. The third argument is the identifier.

26.24.9 HOSTS_CFG_DEL6

```
deleted %1 host(s) having %2 IPv6 reservation(s) for subnet id %3 and identifier %4
```

Logged at debug log level 40. This debug message is issued when IPv6 reservations are deleted for the specified subnet and identifier. The first argument specifies how many hosts have been deleted. The second argument specifies how many reservations have been deleted. The third argument is the subnet identifier. The fourth argument is the identifier.

26.24.10 HOSTS_CFG_DEL_ALL_SUBNET4

```
deleted all %1 host(s) for subnet id %2
```

Logged at debug log level 40. This debug message is issued when all IPv4 reservations are deleted for the specified subnet. The first argument specifies how many reservations have been deleted. The second argument is the subnet identifier.

26.24.11 HOSTS_CFG_DEL_ALL_SUBNET6

```
deleted all %1 host(s) having %2 IPv6 reservation(s) for subnet id %3
```

Logged at debug log level 40. This debug message is issued when all IPv6 reservations are deleted for the specified subnet. The first argument specifies how many hosts have been deleted. The second argument specifies how many IPv6 (addresses and prefixes) reservations have been deleted. The third argument is the subnet identifier.

26.24.12 HOSTS_CFG_GET_ALL

```
get all hosts with reservations
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts.

26.24.13 HOSTS_CFG_GET_ALL_ADDRESS4

```
get all hosts with reservations for IPv4 address %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts, holding the reservation for the specific IPv4 address, from the configuration. The argument specifies the IPv4 address used to search the hosts.

26.24.14 HOSTS_CFG_GET_ALL_ADDRESS4_COUNT

```
using address %1, found %2 host(s)
```

Logged at debug log level 45. This debug message logs the number of hosts found using the specified IPv4 address. The arguments specify the IPv4 address used and the number of hosts found respectively.

26.24.15 HOSTS_CFG_GET_ALL_ADDRESS4_HOST

```
using address %1 found host: %2
```

Logged at debug log level 55. This debug message is issued when found host with the reservation for the specified IPv4 address. The arguments specify the IPv4 address and the detailed description of the host found.

26.24.16 HOSTS_CFG_GET_ALL_ADDRESS6

```
get all hosts with reservations for IPv6 address %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts, holding the reservation for the specific IPv6 address, from the configuration. The argument specifies the IPv6 address used to search the hosts.

26.24.17 HOSTS_CFG_GET_ALL_ADDRESS6_COUNT

```
using address %1, found %2 host(s)
```

Logged at debug log level 45. This debug message logs the number of hosts found using the specified IPv6 address. The arguments specify the IPv6 address used and the number of hosts found respectively.

26.24.18 HOSTS_CFG_GET_ALL_ADDRESS6_HOST

```
using address %1 found host: %2
```

Logged at debug log level 55. This debug message is issued when found host with the reservation for the specified IPv6 address. The arguments specify the IPv6 address and the detailed description of the host found.

26.24.19 HOSTS_CFG_GET_ALL_COUNT

```
found %1 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found. The argument specifies the number of hosts found.

26.24.20 HOSTS_CFG_GET_ALL_HOST

```
found host: %1
```

Logged at debug log level 55. This debug message includes the details of the host found. The argument specifies found host details.

26.24.21 HOSTS_CFG_GET_ALL_HOSTNAME

```
get all hosts with reservations for hostname %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts with the specific hostname. The argument specifies hostname.

26.24.22 HOSTS_CFG_GET_ALL_HOSTNAME_COUNT

```
using hostname %1, found %2 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the hostname. The arguments specify hostname and the number of hosts found respectively.

26.24.23 HOSTS_CFG_GET_ALL_HOSTNAME_HOST

```
using hostname %1, found host: %2
```

Logged at debug log level 55. This debug message includes the details of the host found using the hostname. The arguments specify hostname and found host details respectively.

26.24.24 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4

```
get all hosts with reservations for hostname %1 and IPv4 subnet %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts with the specific hostname connected to the specific DHCPv4 subnet. The argument specifies hostname and subnet id.

26.24.25 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_COUNT

```
using hostname %1 and IPv4 subnet %2, found %3 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the hostname and the DHCPv4 subnet id. The arguments specify hostname, subnet id and the number of hosts found respectively.

26.24.26 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_HOST

```
using hostname %1 and IPv4 subnet %2, found host: %3
```

Logged at debug log level 55. This debug message includes the details of the host found using the hostname and the DHCPv4 subnet id. The arguments specify hostname, subnet id and found host details respectively.

26.24.27 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6

```
get all hosts with reservations for hostname %1 and IPv6 subnet %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts with the specific hostname connected to the specific DHCPv6 subnet. The argument specifies hostname and subnet id.

26.24.28 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_COUNT

```
using hostname %1 and IPv6 subnet %2, found %3 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the hostname and the DHCPv6 subnet id. The arguments specify hostname, subnet id and the number of hosts found respectively.

26.24.29 HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_HOST

```
using hostname %1 and IPv6 subnet %2, found host: %3
```

Logged at debug log level 55. This debug message includes the details of the host found using the hostname and the DHCPv6 subnet id. The arguments specify hostname, subnet id and found host details respectively.

26.24.30 HOSTS_CFG_GET_ALL_IDENTIFIER

```
get all hosts with reservations using identifier: %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve reservations for all hosts identified by HW address or DUID. The argument holds both the identifier type and the value.

26.24.31 HOSTS_CFG_GET_ALL_IDENTIFIER_COUNT

```
using identifier %1, found %2 host(s)
```

Logged at debug log level 45. This debug message logs the number of hosts found using the specified identifier. The arguments specify the identifier used and the number of hosts found respectively.

26.24.32 HOSTS_CFG_GET_ALL_IDENTIFIER_HOST

```
using identifier: %1, found host: %2
```

Logged at debug log level 55. This debug message is issued when found host identified by the specific identifier. The arguments specify the identifier and the detailed description of the host found.

26.24.33 HOSTS_CFG_GET_ALL_SUBNET_ID4

```
get all hosts with reservations for IPv4 subnet %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts connected to the specific DHCPv4 subnet. The argument specifies subnet id.

26.24.34 HOSTS_CFG_GET_ALL_SUBNET_ID4_COUNT

```
using IPv4 subnet %1, found %2 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the DHCPv4 subnet id. The arguments specify subnet id and the number of hosts found respectively.

26.24.35 HOSTS_CFG_GET_ALL_SUBNET_ID4_HOST

```
using IPv4 subnet %1, found host: %2
```

Logged at debug log level 55. This debug message includes the details of the host found using the DHCPv4 subnet id. The arguments specify subnet id and found host details respectively.

26.24.36 HOSTS_CFG_GET_ALL_SUBNET_ID6

```
get all hosts with reservations for IPv6 subnet %1
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts connected to the specific DHCPv6 subnet. The argument specifies subnet id.

26.24.37 HOSTS_CFG_GET_ALL_SUBNET_ID6_COUNT

```
using IPv6 subnet %1, found %2 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the DHCPv6 subnet id. The arguments specify subnet id and the number of hosts found respectively.

26.24.38 HOSTS_CFG_GET_ALL_SUBNET_ID6_HOST

```
using IPv6 subnet %1, found host: %2
```

Logged at debug log level 55. This debug message includes the details of the host found using the DHCPv6 subnet id. The arguments specify subnet id and found host details respectively.

26.24.39 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4

```
get all hosts with reservations for subnet id %1 and IPv4 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts having the reservation for the given IPv4 address within the given subnet. The first argument specifies subnet identifier. The second argument specifies the IPv4 address for which the reservation is to be returned.

26.24.40 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_COUNT

```
using IPv4 subnet %1 and IPv4 address %2, found %3 host(s)
```

Logged at debug log level 45. This debug message logs the number of hosts found having the reservation for the specified IPv4 address within the specified subnet. The first argument specifies the subnet identifier. The second argument specifies the reserved IPv4 address. The third argument specifies the number of hosts found.

26.24.41 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_HOST

```
using IPv4 subnet %1 and IPv4 address %2, found host: %3
```

Logged at debug log level 55. This debug message is issued when found host having the reservation for the specified IPv4 address in the specified subnet. The first argument specifies the subnet identifier. The second argument specifies the reserved IPv4 address. The third argument specifies host details.

26.24.42 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6

```
get all hosts with reservations for subnet id %1 and IPv6 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve all hosts connected to the specific subnet and having the specific IPv6 address reserved. The arguments specify subnet id and IPv6 address respectively.

26.24.43 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_COUNT

```
using subnet id %1 and address %2, found %3 host(s)
```

Logged at debug log level 45. This debug message include the details of the host found using the subnet id and address. The arguments specify subnet id, address and the number of hosts found respectively.

26.24.44 HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_HOST

```
using subnet id %1 and address %2, found host: %3
```

Logged at debug log level 55. This debug message includes the details of the host found using the subnet id and address. The arguments specify subnet id, address and the number of hosts found respectively. found host details respectively.

26.24.45 HOSTS_CFG_GET_ONE_PREFIX

```
get one host with reservation for prefix %1/%2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host having a reservation for a specified prefix. The arguments specify a prefix and prefix length.

26.24.46 HOSTS_CFG_GET_ONE_PREFIX_HOST

```
using prefix %1/%2, found host: %3
```

Logged at debug log level 55. This debug message includes the details of the host found using the specific prefix/prefix length. The arguments specify prefix, prefix length and host details respectively.

26.24.47 HOSTS_CFG_GET_ONE_PREFIX_NULL

```
host not found using prefix %1/%2
```

Logged at debug log level 55. This debug message is issued when no host was found for a specified prefix and prefix length.

26.24.48 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4

```
get one host with reservation for subnet id %1 and IPv4 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host connected to the specific subnet and having the specific IPv4 address reserved. The arguments specify subnet id and IPv4 address respectively.

26.24.49 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_HOST

```
using subnet id %1 and address %2, found host: %3
```

Logged at debug log level 45. This debug message logs the details of the host found using the subnet id and IPv4 address.

26.24.50 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_NULL

```
host not found using subnet id %1 and address %2
```

Logged at debug log level 45. This debug message is issued when no host was found for the specified subnet id and IPv4 address.

26.24.51 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6

```
get one host with reservation for subnet id %1 and having IPv6 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host connected to the specific subnet and having the specific IPv6 address reserved. The arguments specify subnet id and IPv6 address respectively.

26.24.52 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_HOST

```
using subnet id %1 and address %2, found host: %3
```

Logged at debug log level 45. This debug message logs the details of the host found using the subnet id and IPv6 address.

26.24.53 HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_NULL

```
host not found using subnet id %1 and address %2
```

Logged at debug log level 45. This debug message is issued when no host was found using the specified subnet id and IPv6 address.

26.24.54 HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER

```
get one host with %1 reservation for subnet id %2, identified by %3
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host holding IPv4 or IPv6 reservations, which is connected to a specific subnet and is identified by a specific unique identifier. The first argument identifies if the IPv4 or IPv6 reservation is desired.

26.24.55 HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_HOST

```
using subnet id %1 and identifier %2, found host: %3
```

Logged at debug log level 45. This debug message includes the details of a host found using a subnet id and specific host identifier.

26.24.56 HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_NULL

```
host not found using subnet id %1 and identifier %2
```

Logged at debug log level 45. This debug message is issued when no host was found using the specified subnet id and host identifier.

26.24.57 HOSTS_CFG_UPDATE_ADD

```
add the host for reservations: %1
```

Logged at debug log level 40. This debug message is issued when a new host (with reservations) is added to the server's configuration during an update. The argument describes the host and its reservations in detail.

26.24.58 HOSTS_CFG_UPDATE_DEL4

```
deleted %1 host(s) for subnet id %2 and identifier %3
```

Logged at debug log level 40. This debug message is issued when IPv4 reservations are deleted for the specified subnet and identifier during an update. The first argument specifies how many hosts have been deleted. The second argument is the subnet identifier. The third argument is the identifier.

26.24.59 HOSTS_CFG_UPDATE_DEL6

```
deleted %1 host(s) having %2 IPv6 reservation(s) for subnet id %3 and identifier %4
```

Logged at debug log level 40. This debug message is issued when IPv6 reservations are deleted for the specified subnet and identifier during an update. The first argument specifies how many hosts have been deleted. The second argument specifies how many reservations have been deleted. The third argument is the subnet identifier. The fourth argument is the identifier.

26.24.60 HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_ADDRESS4

```
trying alternate sources for host using subnet id %1 and address %2
```

Logged at debug log level 40. This debug message is issued when the Host Manager doesn't find the host connected to the specific subnet and having the reservation for the specific IPv4 address, and it is starting to search for this host in alternate host data sources.

26.24.61 HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER

```
get one host with IPv4 reservation for subnet id %1, identified by %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host holding IPv4 reservation, which is connected to a specific subnet and is identified by a specific unique identifier.

26.24.62 HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_HOST

```
using subnet id %1 and identifier %2, found in %3 host: %4
```

Logged at debug log level 45. This debug message includes the details of a host returned by an alternate hosts data source using a subnet id and specific host identifier.

26.24.63 HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_NULL

```
host not found using subnet id %1 and identifier %2
```

Logged at debug log level 45. This debug message is issued when no host was found using the specified subnet id and host identifier.

26.24.64 HOSTS_MGR_ALTERNATE_GET6_PREFIX

```
trying alternate sources for host using prefix %1/%2
```

Logged at debug log level 40. This debug message is issued when the Host Manager doesn't find the host connected to the specific subnet and having the reservation for the specified prefix, and it is starting to search for this host in alternate host data sources.

26.24.65 HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_ADDRESS6

```
trying alternate sources for host using subnet id %1 and IPv6 address %2
```

Logged at debug log level 40. This debug message is issued when the Host Manager doesn't find the host connected to the specific subnet and having the reservation for the specified IPv6 address, and it is starting to search for this host in alternate host data sources.

26.24.66 HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER

```
get one host with IPv6 reservation for subnet id %1, identified by %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host holding IPv4 reservation, which is connected to a specific subnet and is identified by a specific unique identifier.

26.24.67 HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_HOST

```
using subnet id %1 and identifier %2, found in %3 host: %4
```

Logged at debug log level 45. This debug message includes the details of a host returned by an alternate host data source using a subnet id and specific host identifier.

26.24.68 HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_NULL

```
host not found using subnet id %1 and identifier %2
```

Logged at debug log level 45. This debug message is issued when no host was found using the specified subnet id and host identifier.

26.24.69 HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS4

```
trying alternate sources for hosts using subnet id %1 and address %2
```

Logged at debug log level 40. This debug message is issued when the Host Manager is starting to search for hosts in alternate host data sources by subnet ID and IPv4 address.

26.24.70 HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS6

```
trying alternate sources for hosts using subnet id %1 and address %2
```

Logged at debug log level 40. This debug message is issued when the Host Manager is starting to search for hosts in alternate host data sources by subnet ID and IPv6 address.

26.24.71 HOSTS_MGR_NON_UNIQUE_IP_UNSUPPORTED

```
host data source %1 does not support the mode in which IP reservations are non-unique
```

This warning message is issued when an administrator attempted to configure the server to allow multiple host reservations for the same IP address or prefix. Some host database backends may not support this mode of operation. In this case the administrator should stop using these backends or fall back to the default setting which requires that IP addresses are unique within a subnet. This setting is guaranteed to work for MySQL and Postgres host backends.

26.25 HOST

26.25.1 HOST_CACHE_ADD

```
add host: %1
```

Logged at debug log level 45. This debug message logs the details of the added host cache entry.

26.25.2 HOST_CACHE_ADD_DUPLICATE

```
duplicate host: %1
```

Logged at debug log level 45. The add operation failed because the entry conflicts with an already existing one. The details of the add operation argument are logged.

26.25.3 HOST_CACHE_COMMAND_CLEAR

```
cache-clear command successful
```

The cache-clear command has been successful.

26.25.4 HOST_CACHE_COMMAND_CLEAR_FAILED

```
cache-clear command failed (reason: %1)
```

The cache-clear command has failed. The reason is logged.

26.25.5 HOST_CACHE_COMMAND_FLUSH

```
cache-flush command successful
```

The cache-flush command has been successful.

26.25.6 HOST_CACHE_COMMAND_FLUSH_FAILED

```
cache-flush command failed (parameters: %1, reason: %2)
```

The cache-flush command has failed. Both the reason as well as the parameters passed are logged.

26.25.7 HOST_CACHE_COMMAND_GET

```
cache-get command successful (returned: %1)
```

The cache-get command has been successful. The number of returned entries is logged.

26.25.8 HOST_CACHE_COMMAND_GET_BY_ID

```
cache-get-by-id command successful (returned: %1)
```

The cache-get-by-id command has been successful. The number of returned entries is logged.

26.25.9 HOST_CACHE_COMMAND_GET_BY_ID_FAILED

```
cache-get-by-id command failed (reason: %1)
```

The cache-get-by-id command has failed. The reason is logged.

26.25.10 HOST_CACHE_COMMAND_GET_FAILED

```
cache-get command failed (reason: %1)
```

The cache-get command has failed. The reason is logged.

26.25.11 HOST_CACHE_COMMAND_INSERT

```
cache-insert command successful (inserted: %1, overwritten: %2)
```

The cache-insert command has been successful. The number of inserted entries and the number of entries overwritten by more recent entries are logged.

26.25.12 HOST_CACHE_COMMAND_INSERT_FAILED

```
cache-insert command failed (parameters: %1, reason: %2)
```

The cache-insert command has failed. Both the reason as well as the parameters passed are logged.

26.25.13 HOST_CACHE_COMMAND_LOAD

```
cache-load command successful (loaded: %1, overwritten: %2)
```

The cache-load command has been successful. The number of loaded entries and the number of entries overwritten by more recent entries are logged.

26.25.14 HOST_CACHE_COMMAND_LOAD_FAILED

```
cache-load command failed (parameters: %1, reason: %2)
```

The cache-load command has failed. Both the reason as well as the parameters passed are logged.

26.25.15 HOST_CACHE_COMMAND_REMOVE

```
cache-remove command successful (parameters: %1)
```

The cache-remove command has been successful. Parameters of the host deleted are logged.

26.25.16 HOST_CACHE_COMMAND_REMOVE_FAILED

```
cache-remove command failed (parameters: %1, reason: %2)
```

The cache-remove command has failed. Both the reason as well as the parameters passed are logged.

26.25.17 HOST_CACHE_COMMAND_SIZE

```
cache-clear command successful: %1
```

The cache-size command has been successful and returned the number of entries in the host cache.

26.25.18 HOST_CACHE_COMMAND_SIZE_FAILED

```
cache-size command failed (reason: %1)
```

The cache-size command has failed. The reason is logged.

26.25.19 HOST_CACHE_COMMAND_WRITE

```
cache-write command successful (dumped: %1)
```

The cache-write command has been successful. The number of dumped entries is logged.

26.25.20 HOST_CACHE_COMMAND_WRITE_FAILED

```
cache-write command failed (parameters: %1, reason: %2)
```

The cache-write command has failed. Both the reason as well as the parameters passed are logged.

26.25.21 HOST_CACHE_CONFIGURATION_FAILED

```
failed to configure Host Cache hooks library: %1
```

This error message is issued when there is an error configuring the Host Cache hooks library. The argument provides the detailed error message.

26.25.22 HOST_CACHE_DEINIT_OK

```
unloading Host Cache hooks library successful
```

This informational message indicates that the Host Cache hooks library has been unloaded successfully.

26.25.23 HOST_CACHE_DEL_SUBNET_ID_ADDRESS4

```
using subnet id %1 and address %2, delete host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry deleted using the subnet id and IPv4 address.

26.25.24 HOST_CACHE_DEL_SUBNET_ID_ADDRESS6

```
using subnet id %1 and address %2, delete host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry deleted using the subnet id and IPv6 address.

26.25.25 HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER4

```
using subnet id %1 and identifier %2, delete host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry deleted using a subnet id and specific host identifier.

26.25.26 HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER6

```
using subnet id %1 and identifier %2, delete host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry deleted using a subnet id and specific host identifier.

26.25.27 HOST_CACHE_GET_ONE_PREFIX

```
get one host with reservation for prefix %1/%2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host cache entry having a reservation for a specified prefix. The arguments specify a prefix and prefix length.

26.25.28 HOST_CACHE_GET_ONE_PREFIX_HOST

```
using prefix %1/%2, found host: %3
```

Logged at debug log level 45. This debug message includes the details of the host cache entry found using the specific prefix/prefix length. The arguments specify prefix, prefix length and host details respectively.

26.25.29 HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4

```
get one host with reservation for subnet id %1 and IPv4 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host cache entry connected to the specific subnet and having the specific IPv4 address reserved. The arguments specify subnet id and IPv4 address respectively.

26.25.30 HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4_HOST

```
using subnet id %1 and address %2, found host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry found using the subnet id and IPv4 address.

26.25.31 HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6

```
get one host with reservation for subnet id %1 and including IPv6 address %2
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host cache entry connected to the specific subnet and having the specific IPv6 address reserved. The arguments specify subnet id and IPv6 address respectively.

26.25.32 HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6_HOST

```
using subnet id %1 and address %2, found host: %3
```

Logged at debug log level 45. This debug message logs the details of the host cache entry found using the subnet id and IPv6 address.

26.25.33 HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER

```
get one host with %1 reservation for subnet id %2, identified by %3
```

Logged at debug log level 40. This debug message is issued when starting to retrieve a host cache entry holding IPv4 or IPv6 reservations, which is connected to a specific subnet and is identified by a specific unique identifier. The first argument identifies if the IPv4 or IPv6 reservation is desired.

26.25.34 HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER_HOST

```
using subnet id %1 and identifier %2, found host: %3
```

Logged at debug log level 45. This debug message includes the details of a host cache entry found using a subnet id and specific host identifier.

26.25.35 HOST_CACHE_INIT_OK

```
loading Host Cache hooks library successful
```

This info message indicates that the Host Cache hooks library has been loaded successfully. Enjoy!

26.25.36 HOST_CMDS_DEINIT_OK

```
unloading Host Commands hooks library successful
```

This info message indicates that the Host Commands hooks library has been removed successfully.

26.25.37 HOST_CMDS_INIT_FAILED

```
loading Host Commands hooks library failed: %1
```

This error message indicates an error during loading the Host Commands hooks library. The details of the error are provided as argument of the log message.

26.25.38 HOST_CMDS_INIT_OK

```
loading Host Commands hooks library successful
```

This info message indicates that the Host Commands hooks library has been loaded successfully. Enjoy!

26.25.39 HOST_CMDS_RESERV_ADD

```
reservation-add command called (parameters: %1)
```

The reservation-add command has been called. Parameters passed are logged.

26.25.40 HOST_CMDS_RESERV_ADD_FAILED

```
reservation-add command failed (parameters: %1, reason: %2)
```

The reservation-add command has failed. Both the reason as well as the parameters passed are logged.

26.25.41 HOST_CMDS_RESERV_ADD_SUCCESS

```
reservation-add command success (parameters: %1)
```

The reservation-add command has been successful. Parameters passed are logged.

26.25.42 HOST_CMDS_RESERV_DEL

```
reservation-del command called (parameters: %1)
```

The reservation-del command has been called. Parameters passed are logged.

26.25.43 HOST_CMDS_RESERV_DEL_FAILED

```
reservation-del command failed (parameters: %1, reason: %2)
```

The reservation-del command has failed. Both the reason as well as the parameters passed are logged.

26.25.44 HOST_CMDS_RESERV_DEL_SUCCESS

```
reservation-del command success (parameters: %1)
```

The reservation-del command has been successful. Parameters passed are logged.

26.25.45 HOST_CMDS_RESERV_GET

```
reservation-get command called (parameters: %1)
```

The reservation-get command has been called. Parameters passed are logged.

26.25.46 HOST_CMDS_RESERV_GET_ALL

```
reservation-get-all command called (parameters: %1)
```

The reservation-get-all command has been called. Parameters passed are logged.

26.25.47 HOST_CMDS_RESERV_GET_ALL_FAILED

```
reservation-get-all command failed (parameters: %1, reason: %2)
```

The reservation-get-all command has failed. Both the reason as well as the parameters passed are logged.

26.25.48 HOST_CMDS_RESERV_GET_ALL_SUCCESS

```
reservation-get-all command success (parameters: %1)
```

The reservation-get-all command has been successful. Parameters passed are logged.

26.25.49 HOST_CMDS_RESERV_GET_BY_ADDRESS

```
reservation-get-by-address command called (parameters: %1)
```

The reservation-get-by-address command has been called. Parameters passed are logged.

26.25.50 HOST_CMDS_RESERV_GET_BY_ADDRESS_FAILED

```
reservation-get-by-address command failed (parameters: %1, reason: %2)
```

The reservation-get-by-address command has failed. Both the reason as well as the parameters passed are logged.

26.25.51 HOST_CMDS_RESERV_GET_BY_ADDRESS_SUCCESS

```
reservation-get-by-address command success (parameters: %1)
```

The reservation-get-by-address command has been successful. Parameters passed are logged.

26.25.52 HOST_CMDS_RESERV_GET_BY_HOSTNAME

```
reservation-get-by-hostname command called (parameters: %1)
```

The reservation-get-by-hostname command has been called. Parameters passed are logged.

26.25.53 HOST_CMDS_RESERV_GET_BY_HOSTNAME_FAILED

```
reservation-get-by-hostname command failed (parameters: %1, reason: %2)
```

The reservation-get-by-hostname command has failed. Both the reason as well as the parameters passed are logged.

26.25.54 HOST_CMDS_RESERV_GET_BY_HOSTNAME_SUCCESS

```
reservation-get-by-hostname command success (parameters: %1)
```

The reservation-get-by-hostname command has been successful. Parameters passed are logged.

26.25.55 HOST_CMDS_RESERV_GET_BY_ID

```
reservation-get-by-id command called (parameters: %1)
```

The reservation-get-by-id command has been called. Parameters passed are logged.

26.25.56 HOST_CMDS_RESERV_GET_BY_ID_FAILED

```
reservation-get-by-id command failed (parameters: %1, reason: %2)
```

The reservation-get-by-id command has failed. Both the reason as well as the parameters passed are logged.

26.25.57 HOST_CMDS_RESERV_GET_BY_ID_SUCCESS

```
reservation-get-by-id command success (parameters: %1)
```

The reservation-get-by-id command has been successful. Parameters passed are logged.

26.25.58 HOST_CMDS_RESERV_GET_FAILED

```
reservation-get command failed (parameters: %1, reason: %2)
```

The reservation-add command has failed. Both the reason as well as the parameters passed are logged.

26.25.59 HOST_CMDS_RESERV_GET_PAGE

```
reservation-get-page command called (parameters: %1)
```

The reservation-get-page command has been called. Parameters passed are logged.

26.25.60 HOST_CMDS_RESERV_GET_PAGE_FAILED

```
reservation-get-page command failed (parameters: %1, reason: %2)
```

The reservation-get-page command has failed. Both the reason as well as the parameters passed are logged.

26.25.61 HOST_CMDS_RESERV_GET_PAGE_SUCCESS

```
reservation-get-page command success (parameters: %1)
```

The reservation-get-page command has been successful. Parameters passed are logged.

26.25.62 HOST_CMDS_RESERV_GET_SUCCESS

```
reservation-get command success (parameters: %1)
```

The reservation-get command has been successful. Parameters passed are logged.

26.25.63 HOST_CMDS_RESERV_UPDATE

```
reservation-update command called (parameters: %1)
```

The reservation-update command has been called. Parameters passed are logged.

26.25.64 HOST_CMDS_RESERV_UPDATE_FAILED

```
reservation-update command failed (parameters: %1, reason: %2)
```

The reservation-update command has failed. Both the reason as well as the parameters passed are logged.

26.25.65 HOST_CMDS_RESERV_UPDATE_SUCCESS

```
reservation-update command success (parameters: %1)
```

The reservation-update command has been successful. Parameters passed are logged.

26.26 HTTPS

26.26.1 HTTPS_REQUEST_RECEIVE_START

```
start receiving request from %1
```

Logged at debug log level 50. This debug message is issued when the server starts receiving new request over the established connection. The argument specifies the address of the remote endpoint.

26.27 HTTP

26.27.1 HTTP_BAD_CLIENT_REQUEST_RECEIVED

```
bad request received from %1: %2
```

Logged at debug log level 40. This debug message is issued when an HTTP client sends malformed request to the server. This includes HTTP requests using unexpected content types, including malformed JSON etc. The first argument specifies an address of the remote endpoint which sent the request. The second argument provides a detailed error message.

26.27.2 HTTP_BAD_CLIENT_REQUEST_RECEIVED_DETAILS

```
detailed information about bad request received from %1:\n%2
```

Logged at debug log level 45. This debug message is issued when an HTTP client sends malformed request to the server. It includes detailed information about the received request rejected by the server. The first argument specifies an address of the remote endpoint which sent the request. The second argument provides a request in the textual format. The request is truncated by the logger if it is too large to be printed.

26.27.3 HTTP_BAD_SERVER_RESPONSE_RECEIVED

```
bad response received when communicating with %1: %2
```

Logged at debug log level 40. This debug message is issued when an HTTP client fails to receive a response from the server or when this response is malformed. The first argument specifies the server URL. The second argument provides a detailed error message.

26.27.4 HTTP_BAD_SERVER_RESPONSE_RECEIVED_DETAILS

```
detailed information about bad response received from %1:\n%2
```

Logged at debug log level 45. This debug message is issued when an HTTP client receives malformed response from the server. The first argument specifies an URL of the server. The second argument provides a response in the textual format. The request is truncated by the logger if it is too large to be printed.

26.27.5 HTTP_CLIENT_MT_STARTED

```
HttpClient has been started in multi-threaded mode running %1 threads
```

Logged at debug log level 40. This debug message is issued when a multi-threaded HTTP client instance has been created. The argument specifies the maximum number of threads.

26.27.6 HTTP_CLIENT_QUEUE_SIZE_GROWING

```
queue for URL: %1, now has %2 entries and may be growing too quickly
```

This warning message is issued when the queue of pending requests for the given URL appears to be growing more quickly than the requests can be handled. It will be emitted periodically as long as the queue size continues to grow. This may occur with a surge of client traffic creating a momentary backlog which then subsides as the surge subsides. If it happens continually then it most likely indicates a deployment configuration that cannot sustain the client load.

26.27.7 HTTP_CLIENT_REQUEST_AUTHORIZED

```
received HTTP request authorized for '%1'
```

This information message is issued when the server receives with a matching authentication header. The argument provides the user id.

26.27.8 HTTP_CLIENT_REQUEST_BAD_AUTH_HEADER

```
received HTTP request with malformed authentication header: %1
```

This information message is issued when the server receives a request with a malformed authentication header. The argument explains the problem.

26.27.9 HTTP_CLIENT_REQUEST_NOT_AUTHORIZED

```
received HTTP request with not matching authentication header
```

This information message is issued when the server receives a request with authentication header carrying not recognized credential: the user provided incorrect user id and/or password.

26.27.10 HTTP_CLIENT_REQUEST_NO_AUTH_HEADER

```
received HTTP request without required authentication header
```

This information message is issued when the server receives a request without a required authentication header.

26.27.11 HTTP_CLIENT_REQUEST_RECEIVED

```
received HTTP request from %1
```

Logged at debug log level 40. This debug message is issued when the server finished receiving a HTTP request from the remote endpoint. The address of the remote endpoint is specified as an argument.

26.27.12 HTTP_CLIENT_REQUEST_RECEIVED_DETAILS

```
detailed information about well-formed request received from %1:\n%2
```

Logged at debug log level 45. This debug message is issued when the HTTP server receives a well-formed request. It includes detailed information about the received request. The first argument specifies an address of the remote endpoint which sent the request. The second argument provides the request in the textual format. The request is truncated by the logger if it is too large to be printed.

26.27.13 HTTP_CLIENT_REQUEST_SEND

```
sending HTTP request %1 to %2
```

Logged at debug log level 50. This debug message is issued when the client is starting to send a HTTP request to a server. The first argument holds basic information about the request (HTTP version number and status code). The second argument specifies a URL of the server.

26.27.14 HTTP_CLIENT_REQUEST_SEND_DETAILS

```
detailed information about request sent to %1:\n%2
```

Logged at debug log level 55. This debug message is issued right before the client sends an HTTP request to the server. It includes detailed information about the request. The first argument specifies an URL of the server to which the request is being sent. The second argument provides the request in the textual form. The request is truncated by the logger if it is too large to be printed.

26.27.15 HTTP_CLIENT_REQUEST_TIMEOUT_OCCURRED

```
HTTP request timeout occurred when communicating with %1
```

Logged at debug log level 50. This debug message is issued when the HTTP request timeout has occurred and the server is going to send a response with Http Request timeout status code.

26.27.16 HTTP_COMMAND_MGR_HTTPS_SERVICE_REUSE_FAILED

```
failed to reused HTTPS service bound to address: %1 port: %2
```

This error message indicates that the server has failed reusing existing HTTPS service on the specified address and port. The server can not swith from HTTPS to HTTP sockets using the same address and port.

26.27.17 HTTP_COMMAND_MGR_HTTPS_SERVICE_UPDATED

```
reused HTTPS service bound to address: %1 port: %2 and updated TLS settings
```

This informational message indicates that the server has reused existing HTTPS service on the specified address and port. Note that any change in the TLS setup has been applied.

26.27.18 HTTP_COMMAND_MGR_HTTP_SERVICE_REUSE_FAILED

```
failed to reused HTTP service bound to address: %1 port: %2
```

This error message indicates that the server has failed reusing existing HTTP service on the specified address and port. The server can not swith from HTTP to HTTPS sockets using the same address and port.

26.27.19 HTTP_COMMAND_MGR_HTTP_SERVICE_UPDATED

```
reused HTTP service bound to address: %1 port: %2
```

This informational message indicates that the server has reused existing HTTP service on the specified address and port.

26.27.20 HTTP_COMMAND_MGR_SERVICE_STARTED

```
started %1 service bound to address: %2 port: %3
```

This informational message indicates that the server has started HTTP/HTTPS service on the specified address and port for receiving control commands.

26.27.21 HTTP_COMMAND_MGR_SERVICE_STOPPING

```
Server is stopping %1 service %2
```

This informational message indicates that the server has stopped HTTP/HTTPS service. When known the address and port are displayed.

26.27.22 HTTP_CONNECTION_CLOSE_CALLBACK_FAILED

```
Connection close callback threw an exception
```

This is an error message emitted when the close connection callback registered on the connection failed unexpectedly. This is a programmatic error that should be submitted as a bug.

26.27.23 HTTP_CONNECTION_HANDSHAKE_FAILED

```
TLS handshake with %1 failed with %2
```

This information message is issued when the TLS handshake failed at the server side. The client address and the error message are displayed.

26.27.24 HTTP_CONNECTION_HANDSHAKE_START

```
start TLS handshake with %1 with timeout %2
```

Logged at debug log level 50. This debug message is issued when the server starts the TLS handshake with the remote endpoint. The first argument specifies the address of the remote endpoint. The second argument specifies request timeout in seconds.

26.27.25 HTTP_CONNECTION_SHUTDOWN

```
shutting down HTTP connection from %1
```

Logged at debug log level 40. This debug message is issued when one of the HTTP connections is shut down. The connection can be stopped as a result of an error or after the successful message exchange with a client.

26.27.26 HTTP_CONNECTION_SHUTDOWN_FAILED

```
shutting down HTTP connection failed
```

This error message is issued when an error occurred during shutting down a HTTP connection with a client.

26.27.27 HTTP_CONNECTION_STOP

```
stopping HTTP connection from %1
```

Logged at debug log level 40. This debug message is issued when one of the HTTP connections is stopped. The connection can be stopped as a result of an error or after the successful message exchange with a client.

26.27.28 HTTP_CONNECTION_STOP_FAILED

```
stopping HTTP connection failed
```

This error message is issued when an error occurred during closing a HTTP connection with a client.

26.27.29 HTTP_CONNECTION_WATCH_SOCKET_CLEAR_ERROR

```
clearing connection watch socket failed: %1
```

This error message is issued when an error occurred during clearing the watch socket associated with a HTTP connection with a client. The error is displayed.

26.27.30 HTTP_CONNECTION_WATCH_SOCKET_CLOSE_ERROR

```
closing connection watch socket failed: %1
```

This error message is issued when an error occurred during closing the watch socket associated with a HTTP connection with a client. The error is displayed.

26.27.31 HTTP_CONNECTION_WATCH_SOCKET_MARK_READY_ERROR

```
marking ready connection watch socket failed: %1
```

This error message is issued when an error occurred during marking as ready the watch socket associated with a HTTP connection with a client. The error is displayed.

26.27.32 HTTP_DATA_RECEIVED

```
received %1 bytes from %2
```

Logged at debug log level 55. This debug message is issued when the server receives a chunk of data from the remote endpoint. This may include the whole request or only a part of the request. The first argument specifies the amount of received data. The second argument specifies an address of the remote endpoint which produced the data.

26.27.33 HTTP_IDLE_CONNECTION_TIMEOUT_OCCURRED

```
closing persistent connection with %1 as a result of a timeout
```

Logged at debug log level 50. This debug message is issued when the persistent HTTP connection is being closed as a result of being idle.

26.27.34 HTTP_PREMATURE_CONNECTION_TIMEOUT_OCCURRED

```
premature connection timeout occurred: in transaction ? %1, transid: %2, current_  
→transid: %3
```

This warning message is issued when unexpected timeout occurred during the transaction. This is proven to occur when the system clock is moved manually or as a result of synchronization with a time server. Any ongoing transactions will be interrupted. New transactions should be conducted normally.

26.27.35 HTTP_REQUEST_RECEIVE_START

```
start receiving request from %1 with timeout %2
```

Logged at debug log level 50. This debug message is issued when the server starts receiving new request over the established connection. The first argument specifies the address of the remote endpoint. The second argument specifies request timeout in seconds.

26.27.36 HTTP_SERVER_RESPONSE_RECEIVED

```
received HTTP response from %1
```

Logged at debug log level 40. This debug message is issued when the client finished receiving an HTTP response from the server. The URL of the server is specified as an argument.

26.27.37 HTTP_SERVER_RESPONSE_RECEIVED_DETAILS

```
detailed information about well-formed response received from %1:\n%2
```

Logged at debug log level 45. This debug message is issued when the HTTP client receives a well-formed response from the server. It includes detailed information about the received response. The first argument specifies a URL of the server which sent the response. The second argument provides the response in the textual format. The response is truncated by the logger if it is too large to be printed.

26.27.38 HTTP_SERVER_RESPONSE_SEND

```
sending HTTP response %1 to %2
```

Logged at debug log level 40. This debug message is issued when the server is starting to send a HTTP response to a remote endpoint. The first argument holds basic information about the response (HTTP version number and status code). The second argument specifies an address of the remote endpoint.

26.27.39 HTTP_SERVER_RESPONSE_SEND_DETAILS

```
detailed information about response sent to %1:\n%2
```

Logged at debug log level 45. This debug message is issued right before the server sends a HTTP response to the client. It includes detailed information about the response. The first argument specifies an address of the remote endpoint to which the response is being sent. The second argument provides a response in the textual form. The response is truncated by the logger if it is too large to be printed.

26.28 KEY

26.28.1 KEY_LOOKUP_DISABLED

```
hooks library lookup for a key: GSS-TSIG is not enabled for the current DNS server.
```

Logged at debug log level 40. This debug message is issued when the lookup for a GSS-TSIG key was performed for a DNS server where GSS-TSIG is not enabled.

26.28.2 KEY_LOOKUP_FOUND

```
hooks library lookup for a key: return GSS-TSIG key '%1'.
```

Logged at debug log level 40. This debug message is issued when the lookup for a GSS-TSIG key returned an usable key for protecting the DNS update. The key name is displayed.

26.28.3 KEY_LOOKUP_NONE

```
hooks library lookup for a key: found no usable key.
```

Logged at debug log level 40. This debug message is issued when the lookup for a GSS-TSIG key failed to find an usable key.

26.28.4 KEY_PROCESSING_FAILED

```
The GSS-TKEY processing for server %1 failed because of an error: %2
```

This error message is issued when the key processing for a specific server has failed. The first argument specifies the server identifier and the second argument gives more information about the error.

26.28.5 KEY_PROCESSING_FAILED_UNSPECIFIED_ERROR

```
The GSS-TKEY processing for server %1 failed because of an unspecified error
```

This error message is issued when the key processing for a specific server has failed. The first argument specifies the server identifier.

26.29 LEASE

26.29.1 LEASE_CMDS_ADD4

```
lease4-add command successful (address: %1)
```

Logged at debug log level 20. The lease4-add command has been successful. Lease IPv4 address is logged.

26.29.2 LEASE_CMDS_ADD4_CONFLICT

```
lease4-add command failed due to conflict (parameters: %1, reason: %2)
```

The received lease4-add is well-formed and contains valid parameters but the lease could not be created because it is in conflict with the server state or configuration. The reason for a conflict is logged in the message.

26.29.3 LEASE_CMDS_ADD4_FAILED

```
lease4-add command failed (parameters: %1, reason: %2)
```

The lease4-add command has failed. Both the reason as well as the parameters passed are logged.

26.29.4 LEASE_CMDS_ADD6

```
lease6-add command successful (address: %1)
```

Logged at debug log level 20. The lease6-add command has been successful. Lease IPv6 address is logged.

26.29.5 LEASE_CMDS_ADD6_CONFLICT

```
lease6-add command failed due to conflict (parameters: %1, reason: %2)
```

The received lease6-add is well-formed and contains valid parameters but the lease could not be created because it is in conflict with the server state or configuration. The reason for a conflict is logged in the message.

26.29.6 LEASE_CMDS_ADD6_FAILED

```
lease6-add command failed (parameters: %1, reason: %2)
```

The lease6-add command has failed. Both the reason as well as the parameters passed are logged.

26.29.7 LEASE_CMDS_BULK_APPLY6

```
lease6-bulk-apply command successful (applied addresses count: %1)
```

Logged at debug log level 20. The lease6-bulk-apply command has been successful. The number of applied addresses is logged.

26.29.8 LEASE_CMDS_BULK_APPLY6_FAILED

```
lease6-bulk-apply command failed (parameters: %1, reason: %2)
```

The lease6-bulk-apply command has failed. Both the reason as well as the parameters passed are logged.

26.29.9 LEASE_CMDS_DEINIT_OK

```
unloading Lease Commands hooks library successful
```

This info message indicates that the Lease Commands hooks library has been removed successfully.

26.29.10 LEASE_CMDS_DEL4

```
lease4-del command successful (address: %1)
```

Logged at debug log level 20. The attempt to delete an IPv4 lease (lease4-del command) has been successful. Lease IPv4 address is logged.

26.29.11 LEASE_CMDS_DEL4_FAILED

```
lease4-del command failed (parameters: %1, reason: %2)
```

The attempt to delete an IPv4 lease (lease4-del command) has failed. Both the reason as well as the parameters passed are logged.

26.29.12 LEASE_CMDS_DEL6

```
lease4-del command successful (address: %1)
```

Logged at debug log level 20. The attempt to delete an IPv4 lease (lease4-del command) has been successful. Lease IPv6 address is logged.

26.29.13 LEASE_CMDS_DEL6_FAILED

```
lease6-del command failed (parameters: %1, reason: %2)
```

The attempt to delete an IPv6 lease (lease4-del command) has failed. Both the reason as well as the parameters passed are logged.

26.29.14 LEASE_CMDS_GET4_FAILED

```
lease4-get command failed (parameters: %1, reason: %2)
```

The lease4-get command has failed. Both the reason as well as the parameters passed are logged.

26.29.15 LEASE_CMDS_GET6_FAILED

```
lease6-get command failed (parameters: %1, reason: %2)
```

The lease4-get command has failed. Both the reason as well as the parameters passed are logged.

26.29.16 LEASE_CMDS_INIT_OK

```
loading Lease Commands hooks library successful
```

This info message indicates that the Lease Commands hooks library has been loaded successfully. Enjoy!

26.29.17 LEASE_CMDS_LEASE4_OFFER_FAILED

```
processing error occurred evaluating binding variables: %1
```

This error log is emitted when an error occurs in the lease4_offer handler is invoked. The argument provides an explanation.

26.29.18 LEASE_CMDS_LEASES4_COMMITTED_FAILED

```
processing error occurred evaluating binding variables: %1
```

This error log is emitted when an error occurs in the leases4_committed handler is invoked. The argument provides an explanation.

26.29.19 LEASE_CMDS_LEASES6_COMMITTED_CONFLICT

```
could not updating lease: %1 for: %2
```

This error log is emitted by the leases6_committed callback when attempting to update a lease with new binding-variable values but a conflicting change has occurred rendering the update invalid. The arguments provide the lease address and the query details.

26.29.20 LEASE_CMDS_LEASES6_COMMITTED_FAILED

```
reason: %1
```

This error log is emitted when one or more leases associated with a client query failed to be updated with binding-variable values. The argument provides details. Individual errors for each lease should precede this log.

26.29.21 LEASE_CMDS_LEASES6_COMMITTED_LEASE_ERROR

```
evaluating binding-variables for lease: %1 for: %2, reason: %3
```

This error log is emitted by the leases6_committed callback when an unexpected error occurs evaluating the binding-variables for a given lease. The arguments provide the lease address, the query details, and an error explanation.

26.29.22 LEASE_CMDS_LOAD_ERROR

```
loading Lease Commands hooks library failed: %1
```

This error message indicates an error loading the Lease Commands hooks library. The details of the error are provided as argument of the log message.

26.29.23 LEASE_CMDS_RESEND_DDNS4

```
lease4-resend-ddns command successful: %1
```

A request to update DNS for the requested IPv4 lease has been successfully queued for transmission to kea-dhcp-ddns.

26.29.24 LEASE_CMDS_RESEND_DDNS4_FAILED

```
lease4-resend-ddns command failed: %1
```

A request to update DNS for the requested IPv4 lease has failed. The reason for the failure is logged.

26.29.25 LEASE_CMDS_RESEND_DDNS6

```
lease6-resend-ddns command successful: %1
```

A request to update DNS for the requested IPv6 lease has been successfully queued for transmission to kea-dhcp-ddns.

26.29.26 LEASE_CMDS_RESEND_DDNS6_FAILED

```
lease6-resend-ddns command failed: %1
```

A request to update DNS for the requested IPv6 lease has failed. The reason for the failure is logged.

26.29.27 LEASE_CMDS_UPDATE4

```
lease4-update command successful (address: %1)
```

Logged at debug log level 20. The lease4-update command has been successful. Lease IPv4 address is logged.

26.29.28 LEASE_CMDS_UPDATE4_CONFLICT

```
lease4-update command failed due to conflict (parameters: %1, reason: %2)
```

The received lease4-update is well-formed and contains valid parameters but the lease could not be created because it is in conflict with the server state or configuration. The reason for a conflict is logged in the message.

26.29.29 LEASE_CMDS_UPDATE4_FAILED

```
lease4-update command failed (parameters: %1, reason: %2)
```

The lease4-update command has failed. Both the reason as well as the parameters passed are logged.

26.29.30 LEASE_CMDS_UPDATE6

```
lease6-update command successful (address: %1)
```

Logged at debug log level 20. The lease6-update command has been successful. Lease IPv6 address is logged.

26.29.31 LEASE_CMDS_UPDATE6_CONFLICT

```
lease6-update command failed due to conflict (parameters: %1, reason: %2)
```

The received lease6-update is well-formed and contains valid parameters but the lease could not be created because it is in conflict with the server state or configuration. The reason for a conflict is logged in the message.

26.29.32 LEASE_CMDS_UPDATE6_FAILED

```
lease6-add command failed (parameters: %1, reason: %2)
```

The lease6-update command has failed. Both the reason as well as the parameters passed are logged.

26.29.33 LEASE_CMDS_WIPE4

```
lease4-wipe command successful (parameters: %1)
```

The lease4-wipe command has been successful. Parameters of the command are logged.

26.29.34 LEASE_CMDS_WIPE4_FAILED

```
lease4-wipe command failed (parameters: %1, reason: %2)
```

The lease4-wipe command has failed. Both the reason as well as the parameters passed are logged.

26.29.35 LEASE_CMDS_WIPE6

```
lease6-wipe command successful (parameters: %1)
```

The lease6-wipe command has been successful. Parameters of the command are logged.

26.29.36 LEASE_CMDS_WIPE6_FAILED

```
lease6-wipe command failed (parameters: %1, reason: %2)
```

The lease6-wipe command has failed. Both the reason as well as the parameters passed are logged.

26.29.37 LEASE_QUERY_LOAD_FAILED

```
Lease Query hooks library failed to load: %1
```

This error message indicates that an error occurred attempting to load the Lease Query hooks library. The argument details the error.

26.29.38 LEASE_QUERY_LOAD_OK

```
Lease Query hooks library loaded successfully.
```

This info message indicates that the Lease Query hooks library has been loaded successfully.

26.29.39 LEASE_QUERY_UNLOAD_OK

```
Lease Query hooks library unloaded successfully.
```

This info message indicates that the Lease Query hooks library has been unloaded successfully.

26.30 LEGAL

26.30.1 LEGAL_LOG_COMMAND_NO_LEGAL_STORE

```
LegalStore instance is null
```

This is an error message issued when the Legal Log library attempted to write a control command entry to the legal store and the store instance has not been created. This is a programmatic error and should not occur.

26.30.2 LEGAL_LOG_COMMAND_WRITE_ERROR

```
Could not write command entry to the legal store: %1
```

This is an error message issued when the Legal Log library attempted to write a control command entry to the legal store and the write failed. The message content should provide an detailed explanation. error.

26.30.3 LEGAL_LOG_DB_OPEN_CONNECTION_WITH_RETRY_FAILED

```
Failed to connect to database: %1 with error: %2
```

This is an informational message issued when the the server failed to connect to the store database. The operation started a retry to connect procedure. The database access string with password redacted is logged, along with the error and details for the reconnect procedure.

26.30.4 LEGAL_LOG_LEASE4_NO_LEGAL_STORE

```
LegalStore instance is null
```

This is an error message issued when the Legal Log library attempted to write a IPv4 lease entry to the legal store and the store instance has not been created. This is a programmatic error and should not occur.

26.30.5 LEGAL_LOG_LEASE4_WRITE_ERROR

```
Could not write to the legal store: %1
```

This is an error message issued when the Legal Log library attempted to write a IPv4 lease entry to the legal store and the write failed. The message content should include the physical store name and the nature of the error.

26.30.6 LEGAL_LOG_LEASE6_NO_LEGAL_STORE

```
LegalStore instance is null
```

This is an error message issued when the Legal Log library attempted to write a IPv6 lease entry to the legal store and the store instance has not been created. This is a programmatic error and should not occur.

26.30.7 LEGAL_LOG_LEASE6_WRITE_ERROR

```
Could not write to the legal store: %1
```

This is an error message issued when the Legal Log library attempted to write a IPv6 lease entry to the legal store and the write failed. The message content should include the physical store name and the nature of the error.

26.30.8 LEGAL_LOG_LOAD_ERROR

```
LEGAL LOGGING DISABLED! An error occurred loading the library: %1
```

This is an error message issued when the DHCP Legal Log library could not be loaded. The exact cause should be explained in the log message. No existing stores will be altered, nor any legal logging entries emitted.

26.30.9 LEGAL_LOG_MYSQL_COMMIT

```
committing to MySQL database
```

The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the MySQL settings, the committal may not include a write to disk.

26.30.10 LEGAL_LOG_MYSQL_DB_RECONNECT_ATTEMPT_FAILED

```
database reconnect failed: %1
```

An error message issued when an attempt to reconnect has failed.

26.30.11 LEGAL_LOG_MYSQL_DB_RECONNECT_ATTEMPT_SCHEDULE

```
scheduling attempt %1 of %2 in %3 milliseconds
```

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.30.12 LEGAL_LOG_MYSQL_DB_RECONNECT_FAILED

maximum number of database reconnect attempts: %1, has been exhausted without success

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.30.13 LEGAL_LOG_MYSQL_FATAL_ERROR

Unrecoverable MySQL error occurred: %1 for <%2>, reason: %3 (error code: %4).

An error message indicating that communication with the MySQL database server has been lost. When this occurs the server exits immediately with a non-zero exit code. This is most likely due to a network issue.

26.30.14 LEGAL_LOG_MYSQL_GET_VERSION

obtaining schema version information

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the MySQL database.

26.30.15 LEGAL_LOG_MYSQL_INSERT_LOG

Adding a log entry to the database: %1

Logged at debug log level 50. An informational message logged when a log entry is inserted.

26.30.16 LEGAL_LOG_MYSQL_INVALID_ACCESS

invalid database access string: %1

This is logged when an attempt has been made to parse a database access string and the attempt ended in error. The access string in question - which should be of the form 'keyword=value keyword=value...' is included in the message.

26.30.17 LEGAL_LOG_MYSQL_NO_TLS

TLS was required but is not used

This error message is issued when TLS for the connection was required but TLS is not used.

26.30.18 LEGAL_LOG_MYSQL_ROLLBACK

rolling back MySQL database

The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.30.19 LEGAL_LOG_MYSQL_START_TRANSACTION

starting new MySQL transaction

A debug message issued when a new MySQL transaction is being started. This message is typically not issued when inserting data into a single table because the server doesn't explicitly start transactions in this case. This message is issued when data is inserted into multiple tables with multiple INSERT statements and there may be a need to rollback the whole transaction if any of these INSERT statements fail.

26.30.20 LEGAL_LOG_MYSQL_TLS_CIPHER

```
TLS cipher: %1
```

Logged at debug log level 50. A debug message issued when a new MySQL connected is created with TLS. The TLS cipher name is logged.

26.30.21 LEGAL_LOG_PGSQL_COMMIT

```
committing to PostgreSQL database
```

The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the PostgreSQL settings, the committal may not include a write to disk.

26.30.22 LEGAL_LOG_PGSQL_DB_RECONNECT_ATTEMPT_FAILED

```
database reconnect failed: %1
```

An error message issued when an attempt to reconnect has failed.

26.30.23 LEGAL_LOG_PGSQL_DB_RECONNECT_ATTEMPT_SCHEDULE

```
scheduling attempt %1 of %2 in %3 milliseconds
```

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.30.24 LEGAL_LOG_PGSQL_DB_RECONNECT_FAILED

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.30.25 LEGAL_LOG_PGSQL_DEALLOC_ERROR

```
An error occurred deallocating SQL statements while closing the PostgreSQL log database:  
↪ %1
```

This is an error message issued when a legal log hook library experienced an error freeing database SQL resources as part of closing its connection to the PostgreSQL database. The connection is closed as part of normal server shutdown. This error is most likely a programmatic issue that is highly unlikely to occur or negatively impact server operation.

26.30.26 LEGAL_LOG_PGSQL_FATAL_ERROR

```
Unrecoverable PostgreSQL error occurred: Statement: <%1>, reason: %2 (error code: %3).
```

An error message indicating that communication with the PostgreSQL database server has been lost. When this occurs the server exits immediately with a non-zero exit code. This is most likely due to a network issue.

26.30.27 LEGAL_LOG_PGSQL_GET_VERSION

obtaining schema version information

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the PostgreSQL database.

26.30.28 LEGAL_LOG_PGSQL_INSERT_LOG

Adding a log entry to the database: %1

Logged at debug log level 50. An informational message logged when a log entry is inserted.

26.30.29 LEGAL_LOG_PGSQL_INVALID_ACCESS

invalid database access string: %1

This is logged when an attempt has been made to parse a database access string and the attempt ended in error. The access string in question - which should be of the form 'keyword=value keyword=value...' is included in the message.

26.30.30 LEGAL_LOG_PGSQL_ROLLBACK

rolling back PostgreSQL database

The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.30.31 LEGAL_LOG_PGSQL_START_TRANSACTION

starting a new PostgreSQL transaction

A debug message issued when a new PostgreSQL transaction is being started. This message is typically not issued when inserting data into a single table because the server doesn't explicitly start transactions in this case. This message is issued when data is inserted into multiple tables with multiple INSERT statements and there may be a need to rollback the whole transaction if any of these INSERT statements fail.

26.30.32 LEGAL_LOG_STORE_CLOSED

Legal store closed: %1

This is an informational message issued when the Legal Log library has successfully closed the legal store.

26.30.33 LEGAL_LOG_STORE_CLOSE_ERROR

An error occurred closing the store: %1, error: %2

This is an error message issued when the legal log library experienced an error attempting to close a legal store. This is highly unlikely to occur and should not affect the store content or subsequent legal store operations.

26.30.34 LEGAL_LOG_STORE_OPEN

opening Legal Log file: %1

This informational message is logged when a DHCP server (either V4 or V6) is about to open a legal log file. The parameters of the backend are logged.

26.30.35 LEGAL_LOG_STORE_OPENED

Legal store opened: %1

This is an informational message issued when the Legal Log library has successfully opened the legal store.

26.30.36 LEGAL_LOG_SYSLOG

%1

This informational message contains the message being logged to syslog.

26.30.37 LEGAL_LOG_SYSLOG_STORE_OPEN

opening Legal Syslog: %1

This informational message is logged when a DHCP server (either V4 or V6) is about to open a legal syslog store. The parameters of the backend are logged.

26.30.38 LEGAL_LOG_UNLOAD_ERROR

An error occurred unloading the library: %1

This is an error message issued when an error occurs while unloading the Legal Log library. This is unlikely to occur and normal operations of the library will likely resume when it is next loaded.

26.31 LFC

26.31.1 LFC_FAIL_PID_CREATE

: %1

This message is issued if LFC detected a failure when trying to create the PID file. It includes a more specific error string.

26.31.2 LFC_FAIL_PID_DEL

: %1

This message is issued if LFC detected a failure when trying to delete the PID file. It includes a more specific error string.

26.31.3 LFC_FAIL_PROCESS

: %1

This message is issued if LFC detected a failure when trying to process the files. It includes a more specific error string.

26.31.4 LFC_FAIL_ROTATE

: %1

This message is issued if LFC detected a failure when trying to rotate the files. It includes a more specific error string.

26.31.5 LFC_PROCESSING

```
Previous file: %1, copy file: %2
```

This message is issued just before LFC starts processing the lease files.

26.31.6 LFC_READ_STATS

```
Leases: %1, attempts: %2, errors: %3.
```

This message prints out the number of leases that were read, the number of attempts to read leases and the number of errors encountered while reading.

26.31.7 LFC_ROTATING

```
LFC rotating files
```

This message is issued just before LFC starts rotating the lease files - removing the old and replacing them with the new.

26.31.8 LFC_RUNNING

```
LFC instance already running
```

This message is issued if LFC detects that a previous copy of LFC may still be running via the PID check.

26.31.9 LFC_START

```
Starting lease file cleanup
```

This message is issued as the LFC process starts.

26.31.10 LFC_TERMINATE

```
LFC finished processing
```

This message is issued when the LFC process completes. It does not indicate that the process was successful only that it has finished.

26.31.11 LFC_WRITE_STATS

```
Leases: %1, attempts: %2, errors: %3.
```

This message prints out the number of leases that were written, the number of attempts to write leases and the number of errors encountered while writing.

26.32 LIMITS

26.32.1 LIMITS_CONFIGURATION_LEASE_BACKEND_NOT_AVAILABLE

```
Lease backend not available. Could not check JSON support in the database. Continuing↵
↵without checking...
```

Warning message logged to notify that limits might not work if the recovering database does not have JSON support.

26.32.2 LIMITS_CONFIGURATION_LEASE_BACKEND_SHOULD_HAVE_BEEN_AVAILABLE

Lease backend not available when configuration shows it should have been. This is likely ↵
↵ a programmatic error. Continuing...

Error message logged to notify about an unexpected situation where the lease backend was expected to be available, but it was not. Runtime errors might occur.

26.32.3 LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_CLIENT_CLASS

New lease limit of %1 addresses for client class %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of an address limit per client class

26.32.4 LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_SUBNET

New lease limit of %1 addresses for subnet with ID %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of an address limit per subnet

26.32.5 LIMITS_CONFIGURED_PREFIX_LIMIT_BY_CLIENT_CLASS

New lease limit of %1 prefixes for client class %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of a prefix limit per client class

26.32.6 LIMITS_CONFIGURED_PREFIX_LIMIT_BY_SUBNET

New lease limit of %1 prefixes for subnet with ID %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of a prefix limit per subnet

26.32.7 LIMITS_CONFIGURED_RATE_LIMIT_BY_CLIENT_CLASS

New rate limit of %1 for client class %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of a rate limit per client class

26.32.8 LIMITS_CONFIGURED_RATE_LIMIT_BY_SUBNET

New rate limit of %1 for subnet with ID %2 has been configured.

Logged at debug log level 40. Debug message logged to notify about the successful configuration of a rate limit per subnet

26.32.9 LIMITS_LEASE_LIMIT_EXCEEDED

Lease was not allocated due to exceeding %1.

Logged at debug log level 40. Debug message logged to indicate that the current number of leased addresses or prefixes for a client class or a subnet is exceeding the limit.

26.32.10 LIMITS_LEASE_WITHIN_LIMITS

Lease with address %1 is within limits.

Logged at debug log level 40. Debug message logged to indicate that the current number of leased addresses or prefixes for a client class or a subnet is exceeding the limit.

26.32.11 LIMITS_PACKET_WIIH_SUBNET_ID_RATE_NO_SUBNET

Packet is not being rate limited due to no subnet specified.

Logged at debug log level 55. Debug message logged to indicate that the current packet's subnet rate limit, if any, is not being checked due to the subnet not being set in the callout handle. This can happen e.g. if the subnet had been deleted after it was selected for the currently processed packet.

26.32.12 LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_DROPPED

Packet assigned to client classes %1 is being dropped for exceeding the rate limit of %2.  for client class %3.

Logged at debug log level 40. Debug message logged to indicate that the current packet has exceeded one of the rate limits configured under at least one client class

26.32.13 LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_HONORED

Packet assigned to client classes %1 is being honored.

Logged at debug log level 55. Debug message logged to indicate that the current packet has not exceeded any of the rate limits configured under any client class

26.32.14 LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_DROPPED

Packet assigned to subnet with ID %1 is being dropped for exceeding the rate limit of %2.

Logged at debug log level 40. Debug message logged to indicate that the current packet has exceeded the limit configured under the assigned subnet ID, if any is configured

26.32.15 LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_HONORED

Packet under subnet with ID %1 and limit %2 is being honored.

Logged at debug log level 55. Debug message logged to indicate that the current packet has not exceeded the limit configured under the assigned subnet ID. This message is also logged when no rate limit is configured for this subnet ID.

26.33 LOGIMPL

26.33.1 LOGIMPL_ABOVE_MAX_DEBUG

```
debug level of %1 is too high and will be set to the maximum of %2
```

A message from the interface to the underlying logger implementation reporting that the debug level (as set by an internally-created string `DEBUGn`, where `n` is an integer, e.g. `DEBUG22`) is above the maximum allowed value and has been reduced to that value. The appearance of this message may indicate a programming error - please submit a bug report.

26.33.2 LOGIMPL_BAD_DEBUG_STRING

```
debug string '%1' has invalid format
```

A message from the interface to the underlying logger implementation reporting that an internally-created string used to set the debug level is not of the correct format (it should be of the form `DEBUGn`, where `n` is an integer, e.g. `DEBUG22`). The appearance of this message indicates a programming error - please submit a bug report.

26.33.3 LOGIMPL_BELOW_MIN_DEBUG

```
debug level of %1 is too low and will be set to the minimum of %2
```

A message from the interface to the underlying logger implementation reporting that the debug level (as set by an internally-created string `DEBUGn`, where `n` is an integer, e.g. `DEBUG22`) is below the minimum allowed value and has been increased to that value. The appearance of this message may indicate a programming error - please submit a bug report.

26.34 LOG

26.34.1 LOG_BAD_DESTINATION

```
unrecognized log destination: %1
```

This error message is printed when a logger destination value was given that was not recognized. The destination should be one of "console", "file", or "syslog".

26.34.2 LOG_BAD_SEVERITY

```
unrecognized log severity: %1
```

This error message is printed when a logger severity value was given that was not recognized. The severity should be one of "DEBUG", "INFO", "WARN", "ERROR", "FATAL" or "NONE".

26.34.3 LOG_BAD_STREAM

```
bad log console output stream: %1
```

Logging has been configured so that output is written to the terminal (console) but the stream on which it is to be written is not recognized. Allowed values are "stdout" and "stderr".

26.34.4 LOG_DUPLICATE_MESSAGE_ID

```
duplicate message ID (%1) in compiled code
```

During start-up, Kea detected that the given message identification had been defined multiple times in the Kea code. This indicates a programming error; please submit a bug report.

26.34.5 LOG_DUPLICATE_NAMESPACE

```
line %1: duplicate $NAMESPACE directive found
```

When reading a message file, more than one \$NAMESPACE directive was found. (This directive is used to set a C++ namespace when generating header files during software development.) Such a condition is regarded as an error and the read will be abandoned.

26.34.6 LOG_INPUT_OPEN_FAIL

```
unable to open message file %1 for input: %2
```

The program was not able to open the specified input message file for the reason given.

26.34.7 LOG_INVALID_MESSAGE_ID

```
line %1: invalid message identification '%2'
```

An invalid message identification (ID) has been found during the read of a message file. Message IDs should comprise only alphanumeric characters and the underscore, and should not start with a digit.

26.34.8 LOG_LOCK_TEST_MESSAGE

```
this is a test message.
```

This is a log message used in testing.

26.34.9 LOG_NAMESPACE_EXTRA_ARGS

```
line %1: $NAMESPACE directive has too many arguments
```

The \$NAMESPACE directive in a message file takes a single argument, a namespace in which all the generated symbol names are placed. This error is generated when the compiler finds a \$NAMESPACE directive with more than one argument.

26.34.10 LOG_NAMESPACE_INVALID_ARG

```
line %1: $NAMESPACE directive has an invalid argument ('%2')
```

The \$NAMESPACE argument in a message file should be a valid C++ namespace. This message is output if the simple check on the syntax of the string carried out by the reader fails.

26.34.11 LOG_NAMESPACE_NO_ARGS

line %1: no arguments were given to the \$NAMESPACE directive

The \$NAMESPACE directive in a message file takes a single argument, a C++ namespace in which all the generated symbol names are placed. This error is generated when the compiler finds a \$NAMESPACE directive with no arguments.

26.34.12 LOG_NO_MESSAGE_ID

line %1: message definition line found without a message ID

Within a message file, message are defined by lines starting with a "%". The rest of the line should comprise the message ID and text describing the message. This error indicates the message compiler found a line in the message file comprising just the "%" and nothing else.

26.34.13 LOG_NO_MESSAGE_TEXT

line %1: line found containing a message ID ('%2') and no text

Within a message file, message are defined by lines starting with a "%". The rest of the line should comprise the message ID and text describing the message. This error indicates the message compiler found a line in the message file comprising just the "%" and message identification, but no text.

26.34.14 LOG_NO_SUCH_MESSAGE

could not replace message text for '%1': no such message

During start-up a local message file was read. A line with the listed message identification was found in the file, but the identification is not one contained in the compiled-in message dictionary. This message may appear a number of times in the file, once for every such unknown message identification. There are several reasons why this message may appear: - The message ID has been misspelled in the local message file. - The program outputting the message may not use that particular message (e.g. it originates in a module not used by the program). - The local file was written for an earlier version of the Kea software and the later version no longer generates that message. Whatever the reason, there is no impact on the operation of Kea.

26.34.15 LOG_OPEN_OUTPUT_FAIL

unable to open %1 for output: %2

Originating within the logging code, the program was not able to open the specified output file for the reason given.

26.34.16 LOG_PREFIX_EXTRA_ARGS

line %1: \$PREFIX directive has too many arguments

Within a message file, the \$PREFIX directive takes a single argument, a prefix to be added to the symbol names when a C++ file is created. This error is generated when the compiler finds a \$PREFIX directive with more than one argument. Note: the \$PREFIX directive is deprecated and will be removed in a future version of Kea.

26.34.17 LOG_PREFIX_INVALID_ARG

```
line %1: $PREFIX directive has an invalid argument ('%2')
```

Within a message file, the \$PREFIX directive takes a single argument, a prefix to be added to the symbol names when a C++ file is created. As such, it must adhere to restrictions on C++ symbol names (e.g. may only contain alphanumeric characters or underscores, and may not start with a digit). A \$PREFIX directive was found with an argument (given in the message) that violates those restrictions. Note: the \$PREFIX directive is deprecated and will be removed in a future version of Kea.

26.34.18 LOG_READING_LOCAL_FILE

```
reading local message file %1
```

This is an informational message output by Kea when it starts to read a local message file. (A local message file may replace the text of one or more messages; the ID of the message will not be changed though.)

26.34.19 LOG_READ_ERROR

```
error reading from message file %1: %2
```

The specified error was encountered reading from the named message file.

26.34.20 LOG_UNRECOGNIZED_DIRECTIVE

```
line %1: unrecognized directive '%2'
```

Within a message file, a line starting with a dollar symbol was found (indicating the presence of a directive) but the first word on the line (shown in the message) was not recognized.

26.34.21 LOG_WRITE_ERROR

```
error writing to %1: %2
```

The specified error was encountered by the message compiler when writing to the named output file.

26.35 MT

26.35.1 MT_TCP_LISTENER_MGR_STARTED

```
MtTcpListenerMgr started with %1 threads, listening on %2:%3, use TLS: %4
```

Logged at debug log level 40. This debug message is issued when an MtTcpListenerMgr has been started to accept connections. Arguments detail the number of threads that the listener is using, the address and port at which it is listening, and if TLS is used or not.

26.35.2 MT_TCP_LISTENER_MGR_STOPPED

```
MtTcpListenerMgr for %1:%2 stopped.
```

Logged at debug log level 40. This debug message is issued when the MtTcpListenerMgr, listening at the given address and port, has completed shutdown.

26.35.3 MT_TCP_LISTENER_MGR_STOPPING

```
Stopping MtTcpListenerMgr for %1:%2
```

Logged at debug log level 40. This debug messages is issued when the MtTcpListenerMgr, listening at the given address and port, has begun to shutdown.

26.36 MYSQL

26.36.1 MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4

```
create or update option pool start: %1 pool end: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by pool

26.36.2 MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6

```
create or update option pool start: %1 pool end: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by pool

26.36.3 MYSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6

```
create or update option prefix: %1 prefix len: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by prefix

26.36.4 MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4

```
create or update option by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by subnet id

26.36.5 MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6

```
create or update option by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by subnet id

26.36.6 MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS4

```
create or update client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update client class

26.36.7 MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS6

```
create or update client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update client class

26.36.8 MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4

```
create or update global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update global parameter

26.36.9 MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6

```
create or update global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update global parameter

26.36.10 MYSQL_CB_CREATE_UPDATE_OPTION4

```
create or update option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option

26.36.11 MYSQL_CB_CREATE_UPDATE_OPTION6

```
create or update option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option

26.36.12 MYSQL_CB_CREATE_UPDATE_OPTION_DEF4

```
create or update option definition: %1 code: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option definition

26.36.13 MYSQL_CB_CREATE_UPDATE_OPTION_DEF6

```
create or update option definition: %1 code: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option definition

26.36.14 MYSQL_CB_CREATE_UPDATE_SERVER4

```
create or update server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update a DHCPv4 server information.

26.36.15 MYSQL_CB_CREATE_UPDATE_SERVER6

```
create or update server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update a DHCPv6 server information.

26.36.16 MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK4

```
create or update shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network

26.36.17 MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK6

```
create or update shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network

26.36.18 MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4

```
create or update shared network: %1 option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network option

26.36.19 MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6

```
create or update shared network: %1 option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network option

26.36.20 MYSQL_CB_CREATE_UPDATE_SUBNET4

```
create or update subnet: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update subnet

26.36.21 MYSQL_CB_CREATE_UPDATE_SUBNET6

```
create or update subnet: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update subnet

26.36.22 MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4

```
delete all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all client classes

26.36.23 MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all client classes

26.36.24 MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6

```
delete all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all client classes

26.36.25 MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all client classes

26.36.26 MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4

```
delete all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all global parameters

26.36.27 MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all global parameters

26.36.28 MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6

```
delete all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all global parameters

26.36.29 MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all global parameters

26.36.30 MYSQL_CB_DELETE_ALL_OPTION_DEFS4

```
delete all option definitions
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all option definitions

26.36.31 MYSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all option definitions

26.36.32 MYSQL_CB_DELETE_ALL_OPTION_DEFS6

```
delete all option definitions
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all option definitions

26.36.33 MYSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all option definitions

26.36.34 MYSQL_CB_DELETE_ALL_SERVERS4

```
delete all DHCPv4 servers
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all servers.

26.36.35 MYSQL_CB_DELETE_ALL_SERVERS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all servers.

26.36.36 MYSQL_CB_DELETE_ALL_SERVERS6

```
delete all DHCPv6 servers
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all servers.

26.36.37 MYSQL_CB_DELETE_ALL_SERVERS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all servers.

26.36.38 MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4

```
delete all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all shared networks

26.36.39 MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all shared networks

26.36.40 MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6

```
delete all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all shared networks

26.36.41 MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all shared networks

26.36.42 MYSQL_CB_DELETE_ALL_SUBNETS4

```
delete all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all subnets

26.36.43 MYSQL_CB_DELETE_ALL_SUBNETS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all subnets

26.36.44 MYSQL_CB_DELETE_ALL_SUBNETS6

```
delete all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all subnets

26.36.45 MYSQL_CB_DELETE_ALL_SUBNETS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all subnets

26.36.46 MYSQL_CB_DELETE_BY_POOL_OPTION4

```
delete pool start: %1 pool end: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by pool

26.36.47 MYSQL_CB_DELETE_BY_POOL_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by pool

26.36.48 MYSQL_CB_DELETE_BY_POOL_OPTION6

```
delete pool start: %1 pool end: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by pool

26.36.49 MYSQL_CB_DELETE_BY_POOL_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by pool

26.36.50 MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6

```
delete prefix: %1 prefix len: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by prefix

26.36.51 MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by prefix

26.36.52 MYSQL_CB_DELETE_BY_PREFIX_SUBNET4

```
delete subnet by prefix: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by prefix

26.36.53 MYSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by prefix

26.36.54 MYSQL_CB_DELETE_BY_PREFIX_SUBNET6

```
delete subnet by prefix: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by prefix

26.36.55 MYSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by prefix

26.36.56 MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4

```
delete by subnet id: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by subnet id

26.36.57 MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by subnet id

26.36.58 MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6

```
delete by subnet id: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by subnet id

26.36.59 MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by subnet id

26.36.60 MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4

```
delete subnet by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by subnet id

26.36.61 MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by subnet id

26.36.62 MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6

```
delete subnet by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by subnet id

26.36.63 MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by subnet id

26.36.64 MYSQL_CB_DELETE_CLIENT_CLASS4

```
delete client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete client class

26.36.65 MYSQL_CB_DELETE_CLIENT_CLASS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete client class

26.36.66 MYSQL_CB_DELETE_CLIENT_CLASS6

```
delete client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete client class

26.36.67 MYSQL_CB_DELETE_CLIENT_CLASS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete client class

26.36.68 MYSQL_CB_DELETE_GLOBAL_PARAMETER4

```
delete global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete global parameter

26.36.69 MYSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete global parameter

26.36.70 MYSQL_CB_DELETE_GLOBAL_PARAMETER6

```
delete global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete global parameter

26.36.71 MYSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete global parameter

26.36.72 MYSQL_CB_DELETE_OPTION4

```
delete option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option

26.36.73 MYSQL_CB_DELETE_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option

26.36.74 MYSQL_CB_DELETE_OPTION6

```
delete option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option

26.36.75 MYSQL_CB_DELETE_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option

26.36.76 MYSQL_CB_DELETE_OPTION_DEF4

```
delete option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option definition

26.36.77 MYSQL_CB_DELETE_OPTION_DEF4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option definition

26.36.78 MYSQL_CB_DELETE_OPTION_DEF6

```
delete option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option definition

26.36.79 MYSQL_CB_DELETE_OPTION_DEF6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option definition

26.36.80 MYSQL_CB_DELETE_SERVER4

```
delete DHCPv4 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete a server.

26.36.81 MYSQL_CB_DELETE_SERVER4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete a server.

26.36.82 MYSQL_CB_DELETE_SERVER6

```
delete DHCPv6 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete a server.

26.36.83 MYSQL_CB_DELETE_SERVER6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete a server.

26.36.84 MYSQL_CB_DELETE_SHARED_NETWORK4

```
delete shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network

26.36.85 MYSQL_CB_DELETE_SHARED_NETWORK4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network

26.36.86 MYSQL_CB_DELETE_SHARED_NETWORK6

```
delete shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network

26.36.87 MYSQL_CB_DELETE_SHARED_NETWORK6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network

26.36.88 MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4

```
delete shared network: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network option

26.36.89 MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network option

26.36.90 MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6

```
delete shared network: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network option

26.36.91 MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network option

26.36.92 MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4

```
delete shared network: %1 subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network subnets

26.36.93 MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network subnets

26.36.94 MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6

```
delete shared network: %1 subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network subnets

26.36.95 MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network subnets

26.36.96 MYSQL_CB_GET_ALL_CLIENT_CLASSES4

retrieving all client classes

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all client classes

26.36.97 MYSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all client classes

26.36.98 MYSQL_CB_GET_ALL_CLIENT_CLASSES6

retrieving all client classes

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all client classes

26.36.99 MYSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all client classes

26.36.100 MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4

retrieving all global parameters

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all global parameters

26.36.101 MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all global parameters

26.36.102 MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6

retrieving all global parameters

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all global parameters

26.36.103 MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all global parameters

26.36.104 MYSQL_CB_GET_ALL_OPTIONS4

retrieving all options

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all options

26.36.105 MYSQL_CB_GET_ALL_OPTIONS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all options

26.36.106 MYSQL_CB_GET_ALL_OPTIONS6

retrieving all options

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all options

26.36.107 MYSQL_CB_GET_ALL_OPTIONS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all options

26.36.108 MYSQL_CB_GET_ALL_OPTION_DEFS4

retrieving all option definitions

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all option definitions

26.36.109 MYSQL_CB_GET_ALL_OPTION_DEFS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all option definitions

26.36.110 MYSQL_CB_GET_ALL_OPTION_DEFS6

retrieving all option definitions

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all option definitions

26.36.111 MYSQL_CB_GET_ALL_OPTION_DEFS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all option definitions

26.36.112 MYSQL_CB_GET_ALL_SERVERS4

retrieving all servers

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all DHCPv4 servers

26.36.113 MYSQL_CB_GET_ALL_SERVERS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all DHCPv4 servers

26.36.114 MYSQL_CB_GET_ALL_SERVERS6

retrieving all DHCPv6 servers

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all DHCPv6 servers

26.36.115 MYSQL_CB_GET_ALL_SERVERS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all DHCPv6 servers

26.36.116 MYSQL_CB_GET_ALL_SHARED_NETWORKS4

retrieving all shared networks

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all shared networks

26.36.117 MYSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all shared networks

26.36.118 MYSQL_CB_GET_ALL_SHARED_NETWORKS6

retrieving all shared networks

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all shared networks

26.36.119 MYSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all shared networks

26.36.120 MYSQL_CB_GET_ALL_SUBNETS4

retrieving all subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all subnets

26.36.121 MYSQL_CB_GET_ALL_SUBNETS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all subnets

26.36.122 MYSQL_CB_GET_ALL_SUBNETS6

retrieving all subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all subnets

26.36.123 MYSQL_CB_GET_ALL_SUBNETS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all subnets

26.36.124 MYSQL_CB_GET_CLIENT_CLASS4

```
retrieving client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a client class

26.36.125 MYSQL_CB_GET_CLIENT_CLASS6

```
retrieving client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a client class

26.36.126 MYSQL_CB_GET_GLOBAL_PARAMETER4

```
retrieving global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve global parameter

26.36.127 MYSQL_CB_GET_GLOBAL_PARAMETER6

```
retrieving global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve global parameter

26.36.128 MYSQL_CB_GET_HOST4

```
get host
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve host

26.36.129 MYSQL_CB_GET_HOST6

```
get host
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve host

26.36.130 MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4

```
retrieving modified client classes from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified client classes from specified time

26.36.131 MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified client classes from specified time

26.36.132 MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6

```
retrieving modified client classes from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified client classes from specified time

26.36.133 MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified client classes from specified time

26.36.134 MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4

```
retrieving modified global parameters from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified global parameters from specified time

26.36.135 MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified global parameters from specified time

26.36.136 MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6

```
retrieving modified global parameters from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified global parameters from specified time

26.36.137 MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified global parameters from specified time

26.36.138 MYSQL_CB_GET_MODIFIED_OPTIONS4

retrieving modified options from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified options from specified time

26.36.139 MYSQL_CB_GET_MODIFIED_OPTIONS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified options from specified time

26.36.140 MYSQL_CB_GET_MODIFIED_OPTIONS6

retrieving modified options from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified options from specified time

26.36.141 MYSQL_CB_GET_MODIFIED_OPTIONS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified options from specified time

26.36.142 MYSQL_CB_GET_MODIFIED_OPTION_DEFS4

retrieving modified option definitions from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified option definitions from specified time

26.36.143 MYSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified option definitions from specified time

26.36.144 MYSQL_CB_GET_MODIFIED_OPTION_DEFS6

retrieving modified option definitions from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified option definitions from specified time

26.36.145 MYSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified option definitions from specified time

26.36.146 MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4

```
retrieving modified shared networks from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified shared networks from specified time

26.36.147 MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified shared networks from specified time

26.36.148 MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6

```
retrieving modified shared networks from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified shared networks from specified time

26.36.149 MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified shared networks from specified time

26.36.150 MYSQL_CB_GET_MODIFIED_SUBNETS4

```
retrieving modified subnets from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified subnets from specified time

26.36.151 MYSQL_CB_GET_MODIFIED_SUBNETS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified subnets from specified time

26.36.152 MYSQL_CB_GET_MODIFIED_SUBNETS6

```
retrieving modified subnets from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified subnets from specified time

26.36.153 MYSQL_CB_GET_MODIFIED_SUBNETS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified subnets from specified time

26.36.154 MYSQL_CB_GET_OPTION4

```
retrieving option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option

26.36.155 MYSQL_CB_GET_OPTION6

```
retrieving option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option

26.36.156 MYSQL_CB_GET_OPTION_DEF4

```
retrieving option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option definition

26.36.157 MYSQL_CB_GET_OPTION_DEF6

```
retrieving option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option definition

26.36.158 MYSQL_CB_GET_PORT4

```
get port
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve port

26.36.159 MYSQL_CB_GET_PORT6

```
get port
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve port

26.36.160 MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4

```
retrieving audit entries from: %1 %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve audit entries from specified time and id.

26.36.161 MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve audit entries from specified time

26.36.162 MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6

```
retrieving audit entries from: %1 %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve audit entries from specified time and id

26.36.163 MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve audit entries from specified time

26.36.164 MYSQL_CB_GET_SERVER4

```
retrieving DHCPv4 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a DHCPv4 server information.

26.36.165 MYSQL_CB_GET_SERVER6

```
retrieving DHCPv6 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a DHCPv6 server information.

26.36.166 MYSQL_CB_GET_SHARED_NETWORK4

```
retrieving shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network

26.36.167 MYSQL_CB_GET_SHARED_NETWORK6

```
retrieving shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network

26.36.168 MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4

retrieving shared network: %1 subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network subnets

26.36.169 MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve shared network subnets

26.36.170 MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6

retrieving shared network: %1 subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network subnets

26.36.171 MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve shared network subnets

26.36.172 MYSQL_CB_GET_SUBNET4_BY_PREFIX

retrieving subnet by prefix: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by prefix

26.36.173 MYSQL_CB_GET_SUBNET4_BY_SUBNET_ID

retrieving subnet by subnet id: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by subnet id

26.36.174 MYSQL_CB_GET_SUBNET6_BY_PREFIX

retrieving subnet by prefix: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by prefix

26.36.175 MYSQL_CB_GET_SUBNET6_BY_SUBNET_ID

retrieving subnet by subnet id: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by subnet id

26.36.176 MYSQL_CB_GET_TYPE4

get type

Logged at debug log level 40. Debug message issued when triggered an action to retrieve type

26.36.177 MYSQL_CB_GET_TYPE6

```
get type
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve type

26.36.178 MYSQL_CB_NO_TLS

```
TLS was required but is not used
```

This error message is issued when TLS for the connection was required but TLS is not used.

26.36.179 MYSQL_CB_RECONNECT_ATTEMPT_FAILED4

```
database reconnect failed: %1
```

Error message issued when an attempt to reconnect has failed.

26.36.180 MYSQL_CB_RECONNECT_ATTEMPT_FAILED6

```
database reconnect failed: %1
```

Error message issued when an attempt to reconnect has failed.

26.36.181 MYSQL_CB_RECONNECT_ATTEMPT_SCHEDULE4

```
scheduling attempt %1 of %2 in %3 milliseconds
```

Info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.36.182 MYSQL_CB_RECONNECT_ATTEMPT_SCHEDULE6

```
scheduling attempt %1 of %2 in %3 milliseconds
```

Info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.36.183 MYSQL_CB_RECONNECT_FAILED4

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

Error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.36.184 MYSQL_CB_RECONNECT_FAILED6

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

Error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.36.185 MYSQL_CB_REGISTER_BACKEND_TYPE4

```
register backend
```

Logged at debug log level 40. Debug message issued when triggered an action to register backend

26.36.186 MYSQL_CB_REGISTER_BACKEND_TYPE6

```
register backend
```

Logged at debug log level 40. Debug message issued when triggered an action to register backend

26.36.187 MYSQL_CB_TLS_CIPHER

```
TLS cipher: %1
```

Logged at debug log level 40. A debug message issued when a new MySQL connected is created with TLS. The TLS cipher name is logged.

26.36.188 MYSQL_CB_UNREGISTER_BACKEND_TYPE4

```
unregister backend
```

Logged at debug log level 40. Debug message issued when triggered an action to unregister backend

26.36.189 MYSQL_CB_UNREGISTER_BACKEND_TYPE6

```
unregister backend
```

Logged at debug log level 40. Debug message issued when triggered an action to unregister backend

26.36.190 MYSQL_DEINIT_OK

```
unloading MySQL hooks library successful
```

This informational message indicates that the MySQL Backend hooks library has been unloaded successfully.

26.36.191 MYSQL_FB_DB

```
opening MySQL log database: %1
```

This informational message is logged when a legal log hook library is about to open a MySQL log database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.36.192 MYSQL_HB_DB

```
opening MySQL hosts database: %1
```

This informational message is logged when a DHCP server (either V4 or V6) is about to open a MySQL hosts database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.36.193 MYSQL_HB_DB_GET_VERSION

obtaining schema version information for the MySQL hosts database

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the MySQL hosts database.

26.36.194 MYSQL_HB_DB_READONLY

MySQL host database opened for read access only

This informational message is issued when the user has configured the MySQL database in read-only mode. Kea will not be able to insert or modify host reservations but will be able to retrieve existing ones and assign them to the clients communicating with the server.

26.36.195 MYSQL_HB_DB_RECONNECT_ATTEMPT_FAILED

database reconnect failed: %1

An error message issued when an attempt to reconnect has failed.

26.36.196 MYSQL_HB_DB_RECONNECT_ATTEMPT_SCHEDULE

scheduling attempt %1 of %2 in %3 milliseconds

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.36.197 MYSQL_HB_DB_RECONNECT_FAILED

maximum number of database reconnect attempts: %1, has been exhausted without success

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.36.198 MYSQL_HB_NO_TLS

TLS was required but is not used

This error message is issued when TLS for the connection was required but TLS is not used.

26.36.199 MYSQL_HB_TLS_CIPHER

TLS cipher: %1

Logged at debug log level 40. A debug message issued when a new MySQL connection is created with TLS. The TLS cipher name is logged.

26.36.200 MYSQL_INIT_OK

loading MySQL hooks library successful

This informational message indicates that the MySQL Backend hooks library has been loaded successfully. Enjoy!

26.36.201 MYSQL_LB_ADD_ADDR4

adding IPv4 lease with address %1

Logged at debug log level 50. A debug message issued when the server is about to add an IPv4 lease with the specified address to the MySQL backend database.

26.36.202 MYSQL_LB_ADD_ADDR6

adding IPv6 lease with address %1, lease type %2

Logged at debug log level 50. A debug message issued when the server is about to add an IPv6 lease with the specified address to the MySQL backend database.

26.36.203 MYSQL_LB_COMMIT

committing to MySQL database

Logged at debug log level 50. The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the MySQL settings, the commit may not include a write to disk.

26.36.204 MYSQL_LB_DB

opening MySQL lease database: %1

This informational message is logged when a DHCP server (either V4 or V6) is about to open a MySQL lease database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.36.205 MYSQL_LB_DB_RECONNECT_ATTEMPT_FAILED

database reconnect failed: %1

An error message issued when an attempt to reconnect has failed.

26.36.206 MYSQL_LB_DB_RECONNECT_ATTEMPT_SCHEDULE

scheduling attempt %1 of %2 in %3 milliseconds

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.36.207 MYSQL_LB_DB_RECONNECT_FAILED

maximum number of database reconnect attempts: %1, has been exhausted without success

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.36.208 MYSQL_LB_DELETED_EXPIRED_RECLAIMED

```
deleted %1 reclaimed leases from the database
```

Logged at debug log level 50. A debug message issued when the server has removed a number of reclaimed leases from the database. The number of removed leases is included in the message.

26.36.209 MYSQL_LB_DELETED_SUBNET4_ID

```
deleted %1 leases that match subnet ID %2.
```

Logged at debug log level 50. A debug message issued when the server is removing leases which match respective subnet ID.

26.36.210 MYSQL_LB_DELETED_SUBNET6_ID

```
deleted %1 leases that match subnet ID %2.
```

Logged at debug log level 50. A debug message issued when the server is removing leases which match respective subnet ID.

26.36.211 MYSQL_LB_DELETE_ADDR4

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the MySQL database for the specified address.

26.36.212 MYSQL_LB_DELETE_ADDR6

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the MySQL database for the specified address.

26.36.213 MYSQL_LB_DELETE_EXPIRED_RECLAIMED4

```
deleting reclaimed IPv4 leases that expired more than %1 seconds ago
```

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv4 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.36.214 MYSQL_LB_DELETE_EXPIRED_RECLAIMED6

```
deleting reclaimed IPv6 leases that expired more than %1 seconds ago
```

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv6 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.36.215 MYSQL_LB_GET4

obtaining all IPv4 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases from the MySQL database.

26.36.216 MYSQL_LB_GET6

obtaining all IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases from the MySQL database.

26.36.217 MYSQL_LB_GET_ADDR4

obtaining IPv4 lease for address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the MySQL database for the specified address.

26.36.218 MYSQL_LB_GET_ADDR6

obtaining IPv6 lease for address %1, lease type %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the MySQL database for the specified address.

26.36.219 MYSQL_LB_GET_CLIENTID

obtaining IPv4 leases for client ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the MySQL database for a client with the specified client identification.

26.36.220 MYSQL_LB_GET_DUID

obtaining IPv6 lease for duid %1,

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the MySQL database for the specified duid.

26.36.221 MYSQL_LB_GET_EXPIRED4

obtaining maximum %1 of expired IPv4 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv4 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.36.222 MYSQL_LB_GET_EXPIRED6

obtaining maximum %1 of expired IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv6 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.36.223 MYSQL_LB_GET_HOSTNAME4

obtaining IPv4 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the MySQL database for a client with the specified hostname.

26.36.224 MYSQL_LB_GET_HOSTNAME6

obtaining IPv6 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the MySQL database for a client with the specified hostname.

26.36.225 MYSQL_LB_GET_HWADDR

obtaining IPv4 leases for hardware address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the MySQL database for a client with the specified hardware address.

26.36.226 MYSQL_LB_GET_IAID_DUID

obtaining IPv6 leases for IAID %1, DUID %2, lease type %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the MySQL database for a client with the specified IAID (Identity Association ID) and DUID (DHCP Unique Identifier).

26.36.227 MYSQL_LB_GET_IAID_SUBID_DUID

obtaining IPv6 leases for IAID %1, subnet ID %2, DUID %3, lease type %4

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the MySQL database for a client with the specified IAID (Identity Association ID), subnet ID and DUID (DHCP Unique Identifier).

26.36.228 MYSQL_LB_GET_PAGE4

obtaining at most %1 IPv4 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.36.229 MYSQL_LB_GET_PAGE6

obtaining at most %1 IPv6 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.36.230 MYSQL_LB_GET_RELAYID4

obtaining at most %1 IPv4 leases starting from address %2 with relay id %3 and cltt↵
↵between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a relay id and client transaction time between start and end dates.

26.36.231 MYSQL_LB_GET_RELAYID6

obtaining at most %1 IPv6 leases starting from address %2 with relay id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a relay id.

26.36.232 MYSQL_LB_GET_REMOTEID4

obtaining at most %1 IPv4 leases starting from address %2 with remote id %3 and cltt↵
↵between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a remote id and client transaction time between start and end dates.

26.36.233 MYSQL_LB_GET_REMOTEID6

obtaining at most %1 IPv6 leases starting from address %2 with remote id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a remote id.

26.36.234 MYSQL_LB_GET_SUBID4

obtaining IPv4 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases for a given subnet identifier from the MySQL database.

26.36.235 MYSQL_LB_GET_SUBID6

obtaining IPv6 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases for a given subnet identifier from the MySQL database.

26.36.236 MYSQL_LB_GET_SUBID_CLIENTID

obtaining IPv4 lease for subnet ID %1 and client ID %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the MySQL database for a client with the specified subnet ID and client ID.

26.36.237 MYSQL_LB_GET_SUBID_HWADDR

obtaining IPv4 lease for subnet ID %1 and hardware address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the MySQL database for a client with the specified subnet ID and hardware address.

26.36.238 MYSQL_LB_GET_SUBID_PAGE6

obtaining at most %1 IPv6 leases starting from address %2 for subnet ID %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases from the MySQL database beginning with the specified address for the specified subnet identifier.

26.36.239 MYSQL_LB_GET_VERSION

obtaining schema version information

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the MySQL database.

26.36.240 MYSQL_LB_NEGATIVE_LEASES_STAT

recount of leases returned a negative value

This warning message is issued when the recount of leases using counters in the MySQL database returned a negative value. This shows a problem which can be fixed only by an offline direct recount on the database. This message is issued only once.

26.36.241 MYSQL_LB_NO_TLS

TLS was required but is not used

This error message is issued when TLS for the connection was required but TLS is not used.

26.36.242 MYSQL_LB_ROLLBACK

rolling back MySQL database

Logged at debug log level 50. The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.36.243 MYSQL_LB_TLS_CIPHER

```
TLS cipher: %1
```

Logged at debug log level 40. A debug message issued when a new MySQL connected is created with TLS. The TLS cipher name is logged.

26.36.244 MYSQL_LB_UPDATE_ADDR4

```
updating IPv4 lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv4 lease from the MySQL database for the specified address.

26.36.245 MYSQL_LB_UPDATE_ADDR6

```
updating IPv6 lease for address %1, lease type %2
```

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv6 lease from the MySQL database for the specified address.

26.36.246 MYSQL_LB_UPGRADE_EXTENDED_INFO4

```
upgrading IPv4 leases done in %1 pages with %2 updated leases
```

Logged at debug log level 40. The server upgraded extended info. The number of pages and the final count of updated leases are displayed.

26.36.247 MYSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR

```
upgrading extending info for IPv4 lease at %1 failed with %2
```

Logged at debug log level 40. A debug message issued when the server failed to upgrade an extended info. The address of the lease and the error message are displayed.

26.36.248 MYSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE

```
upgrading IPv4 lease extended info at page %1 starting at %2 (updated %3)
```

Logged at debug log level 50. A debug message issued when the server upgrades IPv4 lease extended info. The page number and started address, and the count of already updated leases are displayed.

26.36.249 MYSQL_LB_UPGRADE_EXTENDED_INFO6

```
upgrading IPv6 leases done in %1 pages with %2 updated leases
```

Logged at debug log level 40. The server upgraded extended info. The number of pages and the final count of updated leases are displayed.

26.36.250 MYSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR

```
upgrading extending info for IPv6 lease at %1 failed with %2
```

Logged at debug log level 40. A debug message issued when the server failed to upgrade the extended info for a lease. The address of the lease and the error message are displayed.

26.36.251 MYSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE

```
upgrading IPv6 lease extended info at page %1 starting at %2 (updated %3)
```

Logged at debug log level 50. A debug message issued when the server upgrades IPv6 lease extended info. The page number and started address, and the count of already updated leases are displayed.

26.37 NETCONF

26.37.1 NETCONF_BOOT_UPDATE_COMPLETED

```
Boot-update configuration completed for server %1
```

This informational message is issued when the initial configuration was retrieved using NETCONF and successfully applied to Kea server.

26.37.2 NETCONF_CONFIG_CHANGED_DETAIL

```
YANG configuration changed: %1
```

Logged at debug log level 55. This debug message indicates a YANG configuration change. The format is the change operation (created, modified, deleted or moved) followed by xpaths and values of old and new nodes.

26.37.3 NETCONF_CONFIG_CHANGE_EVENT

```
Received YANG configuration change %1 event
```

This informational message is issued when kea-netconf receives a YANG configuration change event. The type of event is printed.

26.37.4 NETCONF_CONFIG_CHECK_FAIL

```
NETCONF configuration check failed: %1
```

This error message indicates that kea-netconf had failed configuration check. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.37.5 NETCONF_CONFIG_FAIL

```
NETCONF configuration failed: %1
```

This error message indicates that kea-netconf had failed configuration attempt. Details are provided. Additional details may be available in earlier log entries, possibly on lower levels.

26.37.6 NETCONF_CONFIG_SYNTAX_WARNING

```
NETCONF configuration syntax warning: %1
```

This warning message indicates that the NETCONF configuration had a minor syntax error. The error was displayed and the configuration parsing resumed.

26.37.7 NETCONF_FAILED

```
application experienced a fatal error: %1
```

This is a fatal error message issued when kea-netconf got an unrecoverable error from within the event loop.

26.37.8 NETCONF_GET_CONFIG

```
got configuration from %1 server: %2
```

Logged at debug log level 55. This debug message indicates that kea-netconf got the configuration from a Kea server. The server name and the retrieved configuration are printed.

26.37.9 NETCONF_GET_CONFIG_FAILED

```
getting configuration from %1 server failed: %2
```

The error message indicates that kea-netconf got an error getting the configuration from a Kea server. Make sure that the server is up and running, has appropriate control socket defined and that the controls socket configuration on the server matches that of kea-netconf. The name of the server and the error are printed.

26.37.10 NETCONF_GET_CONFIG_STARTED

```
getting configuration from %1 server
```

This informational message indicates that kea-netconf is trying to get the configuration from a Kea server.

26.37.11 NETCONF_MODULE_CHANGE_INTERNAL_ERROR

```
an internal error occurred while processing changes for module %1: %2
```

The error message indicates that kea-netconf got an error while sysrepo was processing modules changes. This usually follows a config validation failure, and can be recovered from. The name of the module and the internal error message are printed.

26.37.12 NETCONF_MODULE_MISSING_ERR

```
Missing essential module %1 in sysrepo
```

This fatal error message indicates that a module required by Netconf configuration is not available in the sysrepo repository. The name of the module is printed.

26.37.13 NETCONF_MODULE_MISSING_WARN

```
Missing module %1 in sysrepo
```

This warning message indicates that a module used by Kea is not available in the sysrepo repository. The name of the module is printed.

26.37.14 NETCONF_MODULE_REVISION_ERR

```
Essential module %1 does NOT have the right revision: expected %2, got %3
```

This fatal error message indicates that a module required by Netconf configuration is not at the right revision in the sysrepo repository. The name, expected and available revisions of the module are printed.

26.37.15 NETCONF_MODULE_REVISION_WARN

```
Module %1 does NOT have the right revision: expected %2, got %3
```

This warning message indicates that a module used by Kea is not at the right revision in the sysrepo repository. The name, expected and available revisions of the module are printed.

26.37.16 NETCONF_NOTIFICATION_INTERNAL_ERROR

```
an internal error occurred while sending an event notification for module %1: %2
```

The error message indicates that kea-netconf got an error while sysrepo was sending an event notification. This error is not fatal and can be recovered from. The name of the module and the internal error message are printed.

26.37.17 NETCONF_NOTIFICATION_RECEIVED

```
Received notification of type %1 for module %2: '%3'
```

This informational message logs any YANG notification that has been signaled by the server, sent to kea-netconf which then was forwarded to subscribed clients. To achieve this, kea-netconf subscribes itself as a client to all notifications for the configured module.

26.37.18 NETCONF_NOT_SUBSCRIBED_TO_NOTIFICATIONS

```
subscribing to notifications for %1 server with %2 module failed: %3
```

The warning message indicates that kea-netconf got an error subscribing to notifications for a Kea server. The most probable cause is probably that the model that kea-netconf subscribed to does not have any notification nodes, but there may be other more unexpected causes as well. The server name, module name and the error are printed.

26.37.19 NETCONF_RUN_EXIT

```
application is exiting the event loop
```

Logged at debug log level 0. This is a debug message issued when kea-netconf exits its event loop. This is a normal step during kea-netconf shutdown.

26.37.20 NETCONF_SET_CONFIG

```
set configuration to %1 server: %2
```

Logged at debug log level 55. This debug message indicates that kea-netconf set the configuration to a Kea server. The server name and the applied configuration are printed.

26.37.21 NETCONF_SET_CONFIG_FAILED

```
setting configuration to %1 server failed: %2
```

The error message indicates that kea-netconf got an error setting the configuration to a Kea server. Make sure that the server is up and running, has appropriate control socket defined and that the controls socket configuration on the server matches that of kea-netconf. The name of the server and the error are printed.

26.37.22 NETCONF_SET_CONFIG_STARTED

```
setting configuration to %1 server
```

This informational message indicates that kea-netconf is trying to set the configuration to a Kea server.

26.37.23 NETCONF_STARTED

```
kea-netconf (version %1) started
```

This informational message indicates that kea-netconf has processed all configuration information and is ready to begin processing. The version is also printed.

26.37.24 NETCONF_SUBSCRIBE_CONFIG

```
subscribing configuration changes for %1 server with %2 module
```

This information message indicates that kea-netconf is trying to subscribe configuration changes for a Kea server. The names of the server and the module are printed.

26.37.25 NETCONF_SUBSCRIBE_CONFIG_FAILED

```
subscribe configuration changes for %1 server with %2 module failed: %3
```

The error message indicates that kea-netconf got an error subscribing configuration changes for a Kea server. The names of the server and the module, and the error are printed.

26.37.26 NETCONF_SUBSCRIBE_NOTIFICATIONS

```
subscribing to notifications for %1 server with %2 module
```

This information message indicates that kea-netconf is trying to subscribe to notifications for a Kea server. The server name and module name are printed.

26.37.27 NETCONF_UPDATE_CONFIG

```
updating configuration with %1 server: %2
```

Logged at debug log level 55. This debug message indicates that kea-netconf update the configuration of a Kea server. The server name and the updated configuration are printed.

26.37.28 NETCONF_UPDATE_CONFIG_COMPLETED

```
completed updating configuration for %1 server
```

This informational message indicates that kea-netconf updated with success the configuration of a Kea server.

26.37.29 NETCONF_UPDATE_CONFIG_FAILED

```
updating configuration with %1 server: %2
```

The error message indicates that kea-netconf got an error updating the configuration of a Kea server. This includes a configuration rejected by a Kea server when it tried to apply it. The name of the server and the error are printed.

26.37.30 NETCONF_UPDATE_CONFIG_STARTED

```
started updating configuration for %1 server
```

This informational message indicates that kea-netconf is trying to update the configuration of a Kea server.

26.37.31 NETCONF_VALIDATE_CONFIG

```
validating configuration with %1 server: %2
```

Logged at debug log level 55. This debug message indicates that kea-netconf is validating the configuration with a Kea server. The server name and the validated configuration are printed.

26.37.32 NETCONF_VALIDATE_CONFIG_COMPLETED

```
completed validating configuration for %1 server
```

This informational message indicates that kea-netconf validated with success the configuration with a Kea server.

26.37.33 NETCONF_VALIDATE_CONFIG_FAILED

```
validating configuration with %1 server got an error: %2
```

The error message indicates that kea-netconf got an error validating the configuration with a Kea server. This message is produced when exception is thrown during an attempt to validate received configuration. Additional explanation may be provided as a parameter. You may also take a look at earlier log messages. The name of the server and the error are printed.

26.37.34 NETCONF_VALIDATE_CONFIG_REJECTED

```
validating configuration with %1 server was rejected: %2
```

The warning message indicates that kea-netconf got an error validating the configuration with a Kea server. This message is printed when the configuration was rejected during normal processing. Additional explanation may be provided as a parameter. You may also take a look at earlier log messages. The name of the server and the error are printed.

26.37.35 NETCONF_VALIDATE_CONFIG_STARTED

```
started validating configuration for %1 server
```

This informational message indicates that kea-netconf is trying to validate the configuration with a Kea server.

26.38 PERFMON

26.38.1 PERFMON_ALARM_CLEARED

```
Alarm for %1 has been cleared, reported mean duration %2 is now below low-water-ms: %3
```

This info message is emitted when the reported mean duration for an alarm that has been triggered has fallen below the value of its low-water-ms parameter. The arguments detail the alarm's key and the most recently reported mean.

26.38.2 PERFMON_ALARM_TRIGGERED

```
Alarm for %1 has been triggered since %2, reported mean duration %3 exceeds high-water-  
ms: %4
```

This warning message is emitted when the reported mean duration for an alarm exceeds its high-water-ms value. As long as the reported averages remain above the low-water-ms value, the alarm will remain triggered and this message will be repeated every alarm-report-secs. Arguments detail the alarm's key, the time the alarm was first triggered, the most recent reported mean, and the high-water-ms value.

26.38.3 PERFMON_CMDS_CONTROL_ERROR

```
perfmon-control command processing failed: %1
```

This error message is issued when the PerfMon hook library encounters an error processing a perfmon-control command. The argument explains the command error.

26.38.4 PERFMON_CMDS_CONTROL_OK

```
perfmon-control command success: active monitoring: %1, stats-mgr-reporting: %2
```

This info log is issued when perfmon-control command has successfully enabled/disabled active monitoring and/or statistics mgr reporting. Arguments reflect the current state of both.

26.38.5 PERFMON_CMDS_GET_ALL_DURATIONS_ERROR

```
perfmon-get-all-durations command processing failed: %1
```

This error message is issued when the PerfMon hook library encounters an error processing a perfmon-get-all-durations command. The argument explains the command error.

26.38.6 PERFMON_CMDS_GET_ALL_DURATIONS_OK

```
perfmon-get-all-durations returning %1 durations
```

This info log is issued when perfmon-get-all-durations command has completed successfully. The argument contains the number of durations returned.

26.38.7 PERFMON_DEINIT_OK

```
unloading PerfMon hooks library successful
```

This info message indicates that the PerfMon hooks library has been removed successfully.

26.38.8 PERFMON_DHCP4_PKT_EVENTS

```
query: %1 events=[%2]
```

Logged at debug log level 50. This debug message is emitted after an inbound DHCPv4 query has been processed, the arguments are the query label and the dump of the query's packet event stack.

26.38.9 PERFMON_DHCP4_PKT_PROCESS_ERROR

```
Packet event stack was not processed for query %1, reason %2
```

Logged at debug log level 50. This debug message is emitted when the query's event stack could not be processed. This is most likely a programmatic error and should be reported. The arguments identify the query and the reason it could not be processed. These errors should not affect server's normal operations.

26.38.10 PERFMON_DHCP4_SOCKET_RECEIVED_TIME_SUPPORT

```
Kernel supports socket received time? %1
```

Logged at debug log level 40. This debug message is emitted after a (re)configuration and indicates whether or not the packet filter being used by kea-dhcp4 can supply the timestamp a packet was received by the kernel for recording SOCKET_RECEIVED events. If it does not, perfmon will still function but will not have data available to determine kernel buffer wait times.

26.38.11 PERFMON_DHCP6_PKT_EVENTS

```
query: %1 events=[%2]
```

Logged at debug log level 50. The debug message is emitted after an inbound DHCPv6 query has been processed, the arguments are the query label and the dump of the query's packet event stack.

26.38.12 PERFMON_DHCP6_PKT_PROCESS_ERROR

```
Packet event stack was not processed for query %1, reason %2
```

Logged at debug log level 50. This debug message is emitted when the query's event stack could not be processed. This is most likely a programmatic error and should be reported. The arguments identify the query and the reason it could not be processed. These errors should not affect server's normal operations.

26.38.13 PERFMON_DHCP6_SOCKET_RECEIVED_TIME_SUPPORT

```
Kernel supports socket received time? %1
```

Logged at debug log level 40. This debug message is emitted after a (re)configuration and indicates whether or not the packet filter being used by kea-dhcp6 can supply the timestamp a packet was received by the kernel for recording SOCKET_RECEIVED events. If it does not, perfmon will still function but will not have data available to determine kernel buffer wait times.

26.38.14 PERFMON_INIT_FAILED

```
loading PerfMon hooks library failed: %1
```

This error message indicates an error during loading the PerfMon hooks library. The details of the error are provided as argument of the log message.

26.38.15 PERFMON_INIT_OK

```
loading PerfMon hooks library successful
```

This info message indicates that the PerfMon hooks library has been loaded successfully. Enjoy!

26.39 PGSQL

26.39.1 PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4

```
create or update option pool start: %1 pool end: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by pool

26.39.2 PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6

```
create or update option pool start: %1 pool end: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by pool

26.39.3 PGSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6

```
create or update option prefix: %1 prefix len: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by prefix

26.39.4 PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4

```
create or update option by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by subnet id

26.39.5 PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6

```
create or update option by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option by subnet id

26.39.6 PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS4

```
create or update client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update client class

26.39.7 PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS6

```
create or update client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update client class

26.39.8 PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4

```
create or update global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update global parameter

26.39.9 PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6

```
create or update global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update global parameter

26.39.10 PGSQL_CB_CREATE_UPDATE_OPTION4

```
create or update option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option

26.39.11 PGSQL_CB_CREATE_UPDATE_OPTION6

```
create or update option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option

26.39.12 PGSQL_CB_CREATE_UPDATE_OPTION_DEF4

```
create or update option definition: %1 code: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option definition

26.39.13 PGSQL_CB_CREATE_UPDATE_OPTION_DEF6

```
create or update option definition: %1 code: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update option definition

26.39.14 PGSQL_CB_CREATE_UPDATE_SERVER4

```
create or update server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update a DHCPv4 server information.

26.39.15 PGSQL_CB_CREATE_UPDATE_SERVER6

```
create or update server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update a DHCPv6 server information.

26.39.16 PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK4

```
create or update shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network

26.39.17 PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK6

```
create or update shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network

26.39.18 PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4

```
create or update shared network: %1 option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network option

26.39.19 PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6

```
create or update shared network: %1 option
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update shared network option

26.39.20 PGSQL_CB_CREATE_UPDATE_SUBNET4

```
create or update subnet: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update subnet

26.39.21 PGSQL_CB_CREATE_UPDATE_SUBNET6

```
create or update subnet: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to create or update subnet

26.39.22 PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4

```
delete all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all client classes

26.39.23 PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all client classes

26.39.24 PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6

```
delete all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all client classes

26.39.25 PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all client classes

26.39.26 PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4

```
delete all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all global parameters

26.39.27 PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all global parameters

26.39.28 PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6

```
delete all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all global parameters

26.39.29 PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all global parameters

26.39.30 PGSQL_CB_DELETE_ALL_OPTION_DEFS4

```
delete all option definitions
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all option definitions

26.39.31 PGSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all option definitions

26.39.32 PGSQL_CB_DELETE_ALL_OPTION_DEFS6

```
delete all option definitions
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all option definitions

26.39.33 PGSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all option definitions

26.39.34 PGSQL_CB_DELETE_ALL_SERVERS4

```
delete all DHCPv4 servers
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all servers.

26.39.35 PGSQL_CB_DELETE_ALL_SERVERS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all servers.

26.39.36 PGSQL_CB_DELETE_ALL_SERVERS6

```
delete all DHCPv6 servers
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all servers.

26.39.37 PGSQL_CB_DELETE_ALL_SERVERS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all servers.

26.39.38 PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4

```
delete all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all shared networks

26.39.39 PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all shared networks

26.39.40 PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6

```
delete all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all shared networks

26.39.41 PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all shared networks

26.39.42 PGSQL_CB_DELETE_ALL_SUBNETS4

```
delete all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all subnets

26.39.43 PGSQL_CB_DELETE_ALL_SUBNETS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all subnets

26.39.44 PGSQL_CB_DELETE_ALL_SUBNETS6

```
delete all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete all subnets

26.39.45 PGSQL_CB_DELETE_ALL_SUBNETS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete all subnets

26.39.46 PGSQL_CB_DELETE_BY_POOL_OPTION4

```
delete pool start: %1 pool end: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by pool

26.39.47 PGSQL_CB_DELETE_BY_POOL_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by pool

26.39.48 PGSQL_CB_DELETE_BY_POOL_OPTION6

```
delete pool start: %1 pool end: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by pool

26.39.49 PGSQL_CB_DELETE_BY_POOL_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by pool

26.39.50 PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6

```
delete prefix: %1 prefix len: %2 option code: %3 space: %4
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by prefix

26.39.51 PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by prefix

26.39.52 PGSQL_CB_DELETE_BY_PREFIX_SUBNET4

```
delete subnet by prefix: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by prefix

26.39.53 PGSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by prefix

26.39.54 PGSQL_CB_DELETE_BY_PREFIX_SUBNET6

```
delete subnet by prefix: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by prefix

26.39.55 PGSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by prefix

26.39.56 PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4

```
delete by subnet id: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by subnet id

26.39.57 PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by subnet id

26.39.58 PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6

```
delete by subnet id: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option by subnet id

26.39.59 PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option by subnet id

26.39.60 PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4

```
delete subnet by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by subnet id

26.39.61 PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by subnet id

26.39.62 PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6

```
delete subnet by subnet id: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete subnet by subnet id

26.39.63 PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete subnet by subnet id

26.39.64 PGSQL_CB_DELETE_CLIENT_CLASS4

```
delete client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete client class

26.39.65 PGSQL_CB_DELETE_CLIENT_CLASS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete client class

26.39.66 PGSQL_CB_DELETE_CLIENT_CLASS6

```
delete client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete client class

26.39.67 PGSQL_CB_DELETE_CLIENT_CLASS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete client class

26.39.68 PGSQL_CB_DELETE_GLOBAL_PARAMETER4

```
delete global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete global parameter

26.39.69 PGSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete global parameter

26.39.70 PGSQL_CB_DELETE_GLOBAL_PARAMETER6

```
delete global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete global parameter

26.39.71 PGSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete global parameter

26.39.72 PGSQL_CB_DELETE_OPTION4

```
delete option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option

26.39.73 PGSQL_CB_DELETE_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option

26.39.74 PGSQL_CB_DELETE_OPTION6

```
delete option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option

26.39.75 PGSQL_CB_DELETE_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option

26.39.76 PGSQL_CB_DELETE_OPTION_DEF4

```
delete option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option definition

26.39.77 PGSQL_CB_DELETE_OPTION_DEF4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option definition

26.39.78 PGSQL_CB_DELETE_OPTION_DEF6

```
delete option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to delete option definition

26.39.79 PGSQL_CB_DELETE_OPTION_DEF6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete option definition

26.39.80 PGSQL_CB_DELETE_SERVER4

```
delete DHCPv4 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete a server.

26.39.81 PGSQL_CB_DELETE_SERVER4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete a server.

26.39.82 PGSQL_CB_DELETE_SERVER6

```
delete DHCPv6 server: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete a server.

26.39.83 PGSQL_CB_DELETE_SERVER6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete a server.

26.39.84 PGSQL_CB_DELETE_SHARED_NETWORK4

```
delete shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network

26.39.85 PGSQL_CB_DELETE_SHARED_NETWORK4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network

26.39.86 PGSQL_CB_DELETE_SHARED_NETWORK6

```
delete shared network: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network

26.39.87 PGSQL_CB_DELETE_SHARED_NETWORK6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network

26.39.88 PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4

```
delete shared network: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network option

26.39.89 PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network option

26.39.90 PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6

```
delete shared network: %1 option code: %2 space: %3
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network option

26.39.91 PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network option

26.39.92 PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4

```
delete shared network: %1 subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network subnets

26.39.93 PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network subnets

26.39.94 PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6

```
delete shared network: %1 subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to delete shared network subnets

26.39.95 PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT

```
deleted: %1 entries
```

Logged at debug log level 40. Debug message indicating the result of an action to delete shared network subnets

26.39.96 PGSQL_CB_GET_ALL_CLIENT_CLASSES4

```
retrieving all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all client classes

26.39.97 PGSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all client classes

26.39.98 PGSQL_CB_GET_ALL_CLIENT_CLASSES6

```
retrieving all client classes
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all client classes

26.39.99 PGSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all client classes

26.39.100 PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4

```
retrieving all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all global parameters

26.39.101 PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all global parameters

26.39.102 PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6

```
retrieving all global parameters
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all global parameters

26.39.103 PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all global parameters

26.39.104 PGSQL_CB_GET_ALL_OPTIONS4

retrieving all options

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all options

26.39.105 PGSQL_CB_GET_ALL_OPTIONS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all options

26.39.106 PGSQL_CB_GET_ALL_OPTIONS6

retrieving all options

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all options

26.39.107 PGSQL_CB_GET_ALL_OPTIONS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all options

26.39.108 PGSQL_CB_GET_ALL_OPTION_DEFS4

retrieving all option definitions

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all option definitions

26.39.109 PGSQL_CB_GET_ALL_OPTION_DEFS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all option definitions

26.39.110 PGSQL_CB_GET_ALL_OPTION_DEFS6

retrieving all option definitions

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all option definitions

26.39.111 PGSQL_CB_GET_ALL_OPTION_DEFS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all option definitions

26.39.112 PGSQL_CB_GET_ALL_SERVERS4

retrieving all servers

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all DHCPv4 servers

26.39.113 PGSQL_CB_GET_ALL_SERVERS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all DHCPv4 servers

26.39.114 PGSQL_CB_GET_ALL_SERVERS6

```
retrieving all DHCPv6 servers
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all DHCPv6 servers

26.39.115 PGSQL_CB_GET_ALL_SERVERS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all DHCPv6 servers

26.39.116 PGSQL_CB_GET_ALL_SHARED_NETWORKS4

```
retrieving all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all shared networks

26.39.117 PGSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all shared networks

26.39.118 PGSQL_CB_GET_ALL_SHARED_NETWORKS6

```
retrieving all shared networks
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all shared networks

26.39.119 PGSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all shared networks

26.39.120 PGSQL_CB_GET_ALL_SUBNETS4

```
retrieving all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all subnets

26.39.121 PGSQL_CB_GET_ALL_SUBNETS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all subnets

26.39.122 PGSQL_CB_GET_ALL_SUBNETS6

```
retrieving all subnets
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve all subnets

26.39.123 PGSQL_CB_GET_ALL_SUBNETS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve all subnets

26.39.124 PGSQL_CB_GET_CLIENT_CLASS4

```
retrieving client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a client class

26.39.125 PGSQL_CB_GET_CLIENT_CLASS6

```
retrieving client class: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a client class

26.39.126 PGSQL_CB_GET_GLOBAL_PARAMETER4

```
retrieving global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve global parameter

26.39.127 PGSQL_CB_GET_GLOBAL_PARAMETER6

```
retrieving global parameter: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve global parameter

26.39.128 PGSQL_CB_GET_HOST4

```
get host
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve host

26.39.129 PGSQL_CB_GET_HOST6

```
get host
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve host

26.39.130 PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4

```
retrieving modified client classes from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified client classes from specified time

26.39.131 PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified client classes from specified time

26.39.132 PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6

retrieving modified client classes from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified client classes from specified time

26.39.133 PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified client classes from specified time

26.39.134 PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4

retrieving modified global parameters from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified global parameters from specified time

26.39.135 PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified global parameters from specified time

26.39.136 PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6

retrieving modified global parameters from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified global parameters from specified time

26.39.137 PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified global parameters from specified time

26.39.138 PGSQL_CB_GET_MODIFIED_OPTIONS4

retrieving modified options from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified options from specified time

26.39.139 PGSQL_CB_GET_MODIFIED_OPTIONS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified options from specified time

26.39.140 PGSQL_CB_GET_MODIFIED_OPTIONS6

retrieving modified options from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified options from specified time

26.39.141 PGSQL_CB_GET_MODIFIED_OPTIONS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified options from specified time

26.39.142 PGSQL_CB_GET_MODIFIED_OPTION_DEFS4

retrieving modified option definitions from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified option definitions from specified time

26.39.143 PGSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified option definitions from specified time

26.39.144 PGSQL_CB_GET_MODIFIED_OPTION_DEFS6

retrieving modified option definitions from: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified option definitions from specified time

26.39.145 PGSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified option definitions from specified time

26.39.146 PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4

```
retrieving modified shared networks from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified shared networks from specified time

26.39.147 PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified shared networks from specified time

26.39.148 PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6

```
retrieving modified shared networks from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified shared networks from specified time

26.39.149 PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified shared networks from specified time

26.39.150 PGSQL_CB_GET_MODIFIED_SUBNETS4

```
retrieving modified subnets from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified subnets from specified time

26.39.151 PGSQL_CB_GET_MODIFIED_SUBNETS4_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified subnets from specified time

26.39.152 PGSQL_CB_GET_MODIFIED_SUBNETS6

```
retrieving modified subnets from: %1
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve modified subnets from specified time

26.39.153 PGSQL_CB_GET_MODIFIED_SUBNETS6_RESULT

```
retrieving: %1 elements
```

Logged at debug log level 40. Debug message indicating the result of an action to retrieve modified subnets from specified time

26.39.154 PGSQL_CB_GET_OPTION4

```
retrieving option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option

26.39.155 PGSQL_CB_GET_OPTION6

```
retrieving option code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option

26.39.156 PGSQL_CB_GET_OPTION_DEF4

```
retrieving option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option definition

26.39.157 PGSQL_CB_GET_OPTION_DEF6

```
retrieving option definition code: %1 space: %2
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve option definition

26.39.158 PGSQL_CB_GET_PORT4

```
get port
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve port

26.39.159 PGSQL_CB_GET_PORT6

```
get port
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve port

26.39.160 PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4

retrieving audit entries from: %1 %2

Logged at debug log level 40. Debug message issued when triggered an action to retrieve audit entries from specified time and id.

26.39.161 PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve audit entries from specified time

26.39.162 PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6

retrieving audit entries from: %1 %2

Logged at debug log level 40. Debug message issued when triggered an action to retrieve audit entries from specified time and id

26.39.163 PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve audit entries from specified time

26.39.164 PGSQL_CB_GET_SERVER4

retrieving DHCPv4 server: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a DHCPv4 server information.

26.39.165 PGSQL_CB_GET_SERVER6

retrieving DHCPv6 server: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve a DHCPv6 server information.

26.39.166 PGSQL_CB_GET_SHARED_NETWORK4

retrieving shared network: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network

26.39.167 PGSQL_CB_GET_SHARED_NETWORK6

retrieving shared network: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network

26.39.168 PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4

retrieving shared network: %1 subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network subnets

26.39.169 PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve shared network subnets

26.39.170 PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6

retrieving shared network: %1 subnets

Logged at debug log level 40. Debug message issued when triggered an action to retrieve shared network subnets

26.39.171 PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT

retrieving: %1 elements

Logged at debug log level 40. Debug message indicating the result of an action to retrieve shared network subnets

26.39.172 PGSQL_CB_GET_SUBNET4_BY_PREFIX

retrieving subnet by prefix: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by prefix

26.39.173 PGSQL_CB_GET_SUBNET4_BY_SUBNET_ID

retrieving subnet by subnet id: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by subnet id

26.39.174 PGSQL_CB_GET_SUBNET6_BY_PREFIX

retrieving subnet by prefix: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by prefix

26.39.175 PGSQL_CB_GET_SUBNET6_BY_SUBNET_ID

retrieving subnet by subnet id: %1

Logged at debug log level 40. Debug message issued when triggered an action to retrieve subnet by subnet id

26.39.176 PGSQL_CB_GET_TYPE4

get type

Logged at debug log level 40. Debug message issued when triggered an action to retrieve type

26.39.177 PGSQL_CB_GET_TYPE6

```
get type
```

Logged at debug log level 40. Debug message issued when triggered an action to retrieve type

26.39.178 PGSQL_CB_RECONNECT_ATTEMPT_FAILED4

```
database reconnect failed: %1
```

Error message issued when an attempt to reconnect has failed.

26.39.179 PGSQL_CB_RECONNECT_ATTEMPT_FAILED6

```
database reconnect failed: %1
```

Error message issued when an attempt to reconnect has failed.

26.39.180 PGSQL_CB_RECONNECT_ATTEMPT_SCHEDULE4

```
scheduling attempt %1 of %2 in %3 milliseconds
```

Info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.39.181 PGSQL_CB_RECONNECT_ATTEMPT_SCHEDULE6

```
scheduling attempt %1 of %2 in %3 milliseconds
```

Info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.39.182 PGSQL_CB_RECONNECT_FAILED4

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

Error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.39.183 PGSQL_CB_RECONNECT_FAILED6

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

Error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.39.184 PGSQL_CB_REGISTER_BACKEND_TYPE4

```
register backend
```

Logged at debug log level 40. Debug message issued when triggered an action to register backend

26.39.185 PGSQL_CB_REGISTER_BACKEND_TYPE6

register backend

Logged at debug log level 40. Debug message issued when triggered an action to register backend

26.39.186 PGSQL_CB_UNREGISTER_BACKEND_TYPE4

unregister backend

Logged at debug log level 40. Debug message issued when triggered an action to unregister backend

26.39.187 PGSQL_CB_UNREGISTER_BACKEND_TYPE6

unregister backend

Logged at debug log level 40. Debug message issued when triggered an action to unregister backend

26.39.188 PGSQL_DEINIT_OK

unloading PostgreSQL hooks library successful

This informational message indicates that the PostgreSQL Backend hooks library has been unloaded successfully.

26.39.189 PGSQL_FB_DB

opening PostgreSQL log database: %1

This informational message is logged when a legal log hook library is about to open a PostgreSQL log database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.39.190 PGSQL_HB_DB

opening PostgreSQL hosts database: %1

This informational message is logged when a DHCP server (either V4 or V6) is about to open a PostgreSQL hosts database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.39.191 PGSQL_HB_DB_GET_VERSION

obtaining schema version information for the PostgreSQL hosts database

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the PostgreSQL hosts database.

26.39.192 PGSQL_HB_DB_READONLY

PostgreSQL host database opened for read access only

This informational message is issued when the user has configured the PostgreSQL database in read-only mode. Kea will not be able to insert or modify host reservations but will be able to retrieve existing ones and assign them to the clients communicating with the server.

26.39.193 PGSQL_HB_DB_RECONNECT_ATTEMPT_FAILED

```
database reconnect failed: %1
```

An error message issued when an attempt to reconnect has failed.

26.39.194 PGSQL_HB_DB_RECONNECT_ATTEMPT_SCHEDULE

```
scheduling attempt %1 of %2 in %3 milliseconds
```

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.39.195 PGSQL_HB_DB_RECONNECT_FAILED

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.39.196 PGSQL_INIT_OK

```
loading PostgreSQL hooks library successful
```

This informational message indicates that the PostgreSQL Backend hooks library has been loaded successfully. Enjoy!

26.39.197 PGSQL_LB_ADD_ADDR4

```
adding IPv4 lease with address %1
```

Logged at debug log level 50. A debug message issued when the server is about to add an IPv4 lease with the specified address to the PostgreSQL backend database.

26.39.198 PGSQL_LB_ADD_ADDR6

```
adding IPv6 lease with address %1, lease type %2
```

Logged at debug log level 50. A debug message issued when the server is about to add an IPv6 lease with the specified address to the PostgreSQL backend database.

26.39.199 PGSQL_LB_COMMIT

```
committing to PostgreSQL database
```

Logged at debug log level 50. The code has issued a commit call. All outstanding transactions will be committed to the database. Note that depending on the PostgreSQL settings, the commit may not include a write to disk.

26.39.200 PGSQL_LB_DB

```
opening PostgreSQL lease database: %1
```

This informational message is logged when a DHCP server (either V4 or V6) is about to open a PostgreSQL lease database. The parameters of the connection including database name and username needed to access it (but not the password if any) are logged.

26.39.201 PGSQL_LB_DB_RECONNECT_ATTEMPT_FAILED

```
database reconnect failed: %1
```

An error message issued when an attempt to reconnect has failed.

26.39.202 PGSQL_LB_DB_RECONNECT_ATTEMPT_SCHEDULE

```
scheduling attempt %1 of %2 in %3 milliseconds
```

An info message issued when the server is scheduling the next attempt to reconnect to the database. This occurs when the server has lost database connectivity and is attempting to reconnect automatically.

26.39.203 PGSQL_LB_DB_RECONNECT_FAILED

```
maximum number of database reconnect attempts: %1, has been exhausted without success
```

An error message issued when the server failed to reconnect. Loss of connectivity is typically a network or database server issue.

26.39.204 PGSQL_LB_DELETED_SUBNET4_ID

```
deleted %1 leases that match subnet ID %2.
```

Logged at debug log level 50. A debug message issued when the server is removing leases which match respective subnet ID.

26.39.205 PGSQL_LB_DELETED_SUBNET6_ID

```
deleted %1 leases that match subnet ID %2.
```

Logged at debug log level 50. A debug message issued when the server is removing leases which match respective subnet ID.

26.39.206 PGSQL_LB_DELETE_ADDR4

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the PostgreSQL database for the specified address.

26.39.207 PGSQL_LB_DELETE_ADDR6

```
deleting lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to delete a lease for the specified address from the PostgreSQL database for the specified address.

26.39.208 PGSQL_LB_DELETE_EXPIRED_RECLAIMED4

```
deleting reclaimed IPv4 leases that expired more than %1 seconds ago
```

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv4 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.39.209 PGSQL_LB_DELETE_EXPIRED_RECLAIMED6

deleting reclaimed IPv6 leases that expired more than %1 seconds ago

Logged at debug log level 50. A debug message issued when the server is removing reclaimed DHCPv6 leases which have expired longer than a specified period of time. The argument is the amount of time Kea waits after a reclaimed lease expires before considering its removal.

26.39.210 PGSQL_LB_GET4

obtaining all IPv4 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases from the PostgreSQL database.

26.39.211 PGSQL_LB_GET6

obtaining all IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases from the PostgreSQL database.

26.39.212 PGSQL_LB_GET_ADDR4

obtaining IPv4 lease for address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the PostgreSQL database for the specified address.

26.39.213 PGSQL_LB_GET_ADDR6

obtaining IPv6 lease for address %1 (lease type %2)

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the PostgreSQL database for the specified address.

26.39.214 PGSQL_LB_GET_CLIENTID

obtaining IPv4 leases for client ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the PostgreSQL database for a client with the specified client identification.

26.39.215 PGSQL_LB_GET_DUID

obtaining IPv6 leases for DUID %1,

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the PostgreSQL database for a client with the specified DUID (DHCP Unique Identifier).

26.39.216 PGSQL_LB_GET_EXPIRED4

obtaining maximum %1 of expired IPv4 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv4 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.39.217 PGSQL_LB_GET_EXPIRED6

obtaining maximum %1 of expired IPv6 leases

Logged at debug log level 50. A debug message issued when the server is attempting to obtain expired IPv6 leases to reclaim them. The maximum number of leases to be retrieved is logged in the message.

26.39.218 PGSQL_LB_GET_HOSTNAME4

obtaining IPv4 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the PostgreSQL database for a client with the specified hostname.

26.39.219 PGSQL_LB_GET_HOSTNAME6

obtaining IPv6 leases for hostname %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the PostgreSQL database for a client with the specified hostname.

26.39.220 PGSQL_LB_GET_HWADDR

obtaining IPv4 leases for hardware address %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv4 leases from the PostgreSQL database for a client with the specified hardware address.

26.39.221 PGSQL_LB_GET_IAID_DUID

obtaining IPv4 leases for IAID %1 and DUID %2, lease type %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a set of IPv6 leases from the PostgreSQL database for a client with the specified IAID (Identity Association ID) and DUID (DHCP Unique Identifier).

26.39.222 PGSQL_LB_GET_IAID_SUBID_DUID

obtaining IPv4 leases for IAID %1, subnet ID %2, DUID %3, and lease type %4

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv6 lease from the PostgreSQL database for a client with the specified IAID (Identity Association ID), subnet ID and DUID (DHCP Unique Identifier).

26.39.223 PGSQL_LB_GET_PAGE4

obtaining at most %1 IPv4 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.39.224 PGSQL_LB_GET_PAGE6

obtaining at most %1 IPv6 leases starting from address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of leases beginning with the specified address.

26.39.225 PGSQL_LB_GET_RELAYID4

obtaining at most %1 IPv4 leases starting from address %2 with relay id %3 and cltt↵
↵ between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a relay id and client transaction time between start and end dates.

26.39.226 PGSQL_LB_GET_RELAYID6

obtaining at most %1 IPv6 leases starting from address %2 with relay id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a relay id.

26.39.227 PGSQL_LB_GET_REMOTEID4

obtaining at most %1 IPv4 leases starting from address %2 with remote id %3 and cltt↵
↵ between %4 and %5

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv4 leases beginning with the specified address with a remote id and client transaction time between start and end dates.

26.39.228 PGSQL_LB_GET_REMOTEID6

obtaining at most %1 IPv6 leases starting from address %2 with remote id %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases beginning with the specified address with a remote id.

26.39.229 PGSQL_LB_GET_SUBID4

obtaining IPv4 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv4 leases for a given subnet identifier from the PostgreSQL database.

26.39.230 PGSQL_LB_GET_SUBID6

obtaining IPv6 leases for subnet ID %1

Logged at debug log level 50. A debug message issued when the server is attempting to obtain all IPv6 leases for a given subnet identifier from the PostgreSQL database.

26.39.231 PGSQL_LB_GET_SUBID_CLIENTID

obtaining IPv4 lease for subnet ID %1 and client ID %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the PostgreSQL database for a client with the specified subnet ID and client ID.

26.39.232 PGSQL_LB_GET_SUBID_HWADDR

obtaining IPv4 lease for subnet ID %1 and hardware address %2

Logged at debug log level 50. A debug message issued when the server is attempting to obtain an IPv4 lease from the PostgreSQL database for a client with the specified subnet ID and hardware address.

26.39.233 PGSQL_LB_GET_SUBID_PAGE6

obtaining at most %1 IPv6 leases starting from address %2 for subnet ID %3

Logged at debug log level 50. A debug message issued when the server is attempting to obtain a page of IPv6 leases from the PostgreSQL database beginning with the specified address for the specified subnet identifier.

26.39.234 PGSQL_LB_GET_VERSION

obtaining schema version information

Logged at debug log level 50. A debug message issued when the server is about to obtain schema version information from the PostgreSQL database.

26.39.235 PGSQL_LB_NEGATIVE_LEASES_STAT

recount of leases returned a negative value

This warning message is issued when the recount of leases using counters in the PostgreSQL database returned a negative value. This shows a problem which can be fixed only by an offline direct recount on the database. This message is issued only once.

26.39.236 PGSQL_LB_ROLLBACK

rolling back PostgreSQL database

Logged at debug log level 50. The code has issued a rollback call. All outstanding transaction will be rolled back and not committed to the database.

26.39.237 PGSQL_LB_UPDATE_ADDR4

```
updating IPv4 lease for address %1
```

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv4 lease from the PostgreSQL database for the specified address.

26.39.238 PGSQL_LB_UPDATE_ADDR6

```
updating IPv6 lease for address %1, lease type %2
```

Logged at debug log level 50. A debug message issued when the server is attempting to update IPv6 lease from the PostgreSQL database for the specified address.

26.39.239 PGSQL_LB_UPGRADE_EXTENDED_INFO4

```
upgrading IPv4 leases done in %1 pages with %2 updated leases
```

Logged at debug log level 40. The server upgraded extended info. The number of pages and the final count of updated leases are displayed.

26.39.240 PGSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR

```
upgrading extending info for IPv4 lease at %1 failed with %2
```

Logged at debug log level 40. A debug message issued when the server failed to upgrade an extended info. The address of the lease and the error message are displayed.

26.39.241 PGSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE

```
upgrading IPv4 lease extended info at page %1 starting at %2 (updated %3)
```

Logged at debug log level 50. A debug message issued when the server upgrades IPv4 lease extended info. The page number and started address, and the count of already updated leases are displayed.

26.39.242 PGSQL_LB_UPGRADE_EXTENDED_INFO6

```
upgrading IPv6 leases done in %1 pages with %2 updated leases
```

Logged at debug log level 40. The server upgraded extended info. The number of pages and the final count of updated leases are displayed.

26.39.243 PGSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR

```
upgrading extending info for IPv6 lease at %1 failed with %2
```

Logged at debug log level 40. A debug message issued when the server failed to upgrade the extended info for a lease. The address of the lease and the error message are displayed.

26.39.244 PGSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE

upgrading IPv6 lease extended info at page %1 starting at %2 (updated %3)

Logged at debug log level 50. A debug message issued when the server upgrades IPv6 lease extended info. The page number and started address, and the count of already updated leases are displayed.

26.40 PING

26.40.1 PING_CHECK_CB4_UPDATE_FAILED

A subnet ping-check parameters failed to parse after being updated %1

This error message is emitted when an error occurs trying to parse a subnet ping-check parameters after the subnet was updated via configuration backend. This implies one or more of the parameters is invalid and must be corrected.

26.40.2 PING_CHECK_CHANNEL_ECHO_REPLY_RECEIVED

from address %1, id %2, sequence %3

Logged at debug log level 50. This debug message is issued when an ECHO REPLY has been received on the ping channel's ICMP socket.

26.40.3 PING_CHECK_CHANNEL_ECHO_REQUEST_SENT

to address %1, id %2, sequence %3

Logged at debug log level 50. This debug message is issued when an ECHO REQUEST has been written to the ping channel's ICMP socket.

26.40.4 PING_CHECK_CHANNEL_MALFORMED_PACKET_RECEIVED

error occurred unpacking message %1, discarding it

Logged at debug log level 40. This debug message is emitted when an ICMP packet has been received that could not be unpacked.

26.40.5 PING_CHECK_CHANNEL_NETWORK_WRITE_ERROR

occurred trying to ping %1, error %2

This error message occurs when an asynchronous write on the ICMP socket failed trying to send on the ping target's network. This may mean an interface is down or there is a configuration error. The lease address to ping and the type of the error are provided in the arguments.

26.40.6 PING_CHECK_CHANNEL_SOCKET_CLOSED

ICMP socket has been closed.

Logged at debug log level 40. This debug message is emitted when the ICMP socket for carrying out ping checks has been closed.

26.40.7 PING_CHECK_CHANNEL_SOCKET_CLOSE_ERROR

```
an attempt to close the ICMP socket failed %1
```

This error message is emitted when an unexpected error occurred while closing the ping check ICMP socket. The error detail is provided as an argument of the log message.

26.40.8 PING_CHECK_CHANNEL_SOCKET_OPENED

```
ICMP socket been opened successfully.
```

Logged at debug log level 40. This debug message is emitted when the ICMP socket for carrying out ping checks has been successfully opened.

26.40.9 PING_CHECK_CHANNEL_SOCKET_READ_FAILED

```
socket read completed with an error %1
```

This error message occurs when an asynchronous read on the ICMP socket failed. The details of the error are provided as an argument of the log message.

26.40.10 PING_CHECK_CHANNEL_SOCKET_WRITE_FAILED

```
socket write completed with an error %1
```

This error message occurs when an asynchronous write on the ICMP socket failed. The details of the error are provided as an argument of the log message.

26.40.11 PING_CHECK_CHANNEL_STOP

```
channel is stopping operations.
```

Logged at debug log level 40. This debug message indicates that the channel is stopping operations and closing the ICMP socket. The reason for stopping should be apparent in preceding log messages.

26.40.12 PING_CHECK_CHANNEL_WATCH_SOCKET_CLEAR_ERROR

```
an attempt to clear the WatchSocket associated with
```

the single-threaded ping-channel failed %1 This error message is emitted when an unexpected error occurred while clearing the ready marker of the WatchSocket associated with the ping check channel. This can only occur when running in single-threaded mode. The error detail is provided as an argument of the log message.

26.40.13 PING_CHECK_CHANNEL_WATCH_SOCKET_CLOSE_ERROR

```
an attempt to close the WatchSocket associated with
```

the single-threaded ping-channel failed %1 This error message is emitted when an unexpected error occurred while closing the WatchSocket associated with the ping check channel. This can only occur when running in single-threaded mode. The error detail is provided as an argument of the log message.

26.40.14 PING_CHECK_DHCP4_SRV_CONFIGURED_FAILED

```
dhcp4_srv_configured callout failed %1
```

This error message indicates an error during the Ping Check hook library dhcp4_srv_configured callout. The details of the error are provided as argument of the log message.

26.40.15 PING_CHECK_DUPLICATE_CHECK

```
Ping check already in progress for %1, initiated by %2
```

Logged at debug log level 40. This debug message is emitted when a duplicate request to test an address is received. When this occurs the duplicate test will be skipped and the associated DHCPOFFER will be dropped.

26.40.16 PING_CHECK_LEASE4_OFFER_FAILED

```
lease4_offer callout failed for query %1, lease address %2, reason %3
```

This error message indicates an error during the Ping Check hook library lease4_offer callout. The details of the error are provided as argument of the log message.

26.40.17 PING_CHECK_LOAD_ERROR

```
loading Ping Check hooks library failed %1
```

This error message indicates an error during loading the Ping Check hooks library. The details of the error are provided as argument of the log message.

26.40.18 PING_CHECK_LOAD_OK

```
Ping Check hooks library loaded successfully.
```

This info message indicates that the Ping Check hooks library has been loaded successfully.

26.40.19 PING_CHECK_MGR_CHANNEL_DOWN

```
Ping Channel has shutdown, ping checking will be skipped
```

This error message is emitted when the underlying ICMP channel has stopped due to an unrecoverable error. DHCP service may continue to function but without performing ping checks. Prior log messages should provide details.

26.40.20 PING_CHECK_MGR_LEASE_FREE_TO_USE

```
address %1 is free to use for %2
```

Logged at debug log level 40. This debug message is emitted when ping check has deemed an address is free to use. The log arguments detail the lease address checked and the query which initiated the check.

26.40.21 PING_CHECK_MGR_NEXT_ECHO_SCHEDULED

```
for %1, scheduling ECHO_REQUEST %2 of %3
```

Logged at debug log level 50. This debug message is emitted when the minimum number of ECHO REQUESTs is greater than 1 and the next ECHO REQUEST for a given lease address has been scheduled.

26.40.22 PING_CHECK_MGR_RECEIVED_ECHO_REPLY

```
from %1, id %2, sequence %3
```

Logged at debug log level 40. This debug message is emitted when an ECHO REPLY message has been received. The log argument details the source IP address, id, and sequence number of the ECHO REPLY.

26.40.23 PING_CHECK_MGR_RECEIVED_UNEXPECTED_ECHO_REPLY

```
from %1, id %2, sequence %3 received after reply-timeout expired
```

Logged at debug log level 50. This debug message is emitted when an ECHO REPLY has been received after the reply-timeout has expired and is no longer of interest. This may be an errant ECHO REPLY or it may indicate that the reply-timeout value is too short. The log argument details the source IP address, id, and sequence number of the reply.

26.40.24 PING_CHECK_MGR_RECEIVED_UNEXPECTED_UNREACHABLE_MSG

```
for %1, id %2, sequence %3 received after reply-timeout expired
```

Logged at debug log level 50. This debug message is emitted when an UNREACHABLE message has been received after the reply-timeout has expired and is no longer of interest. This may be an errant message or it may indicate that the reply-timeout value is too short.

26.40.25 PING_CHECK_MGR_RECEIVED_UNREACHABLE_MSG

```
for %1, id %2, sequence %3
```

Logged at debug log level 50. This debug message is emitted when an UNREACHABLE message has been received. The log argument details the target IP address, id, and sequence number from the embedded ECHO REQUEST.

26.40.26 PING_CHECK_MGR_REPLY_RECEIVED_ERROR

```
an error occurred processing an ICMP reply message %1
```

This debug message is emitted when an error occurred while processing an inbound ICMP message. The log argument describes the specific error.

26.40.27 PING_CHECK_MGR_REPLY_TIMEOUT_EXPIRED

```
for %1, ECHO REQUEST %2 of %3, reply-timeout %4
```

Logged at debug log level 50. This debug message is emitted when no reply is received to an ECHO REQUEST before the configured timeout value, *reply-timeout* was reached. The log arguments provides details.

26.40.28 PING_CHECK_MGR_SEND_COMPLETED_ERROR

```
an error occurred in the send completion callback %1
```

This error message is emitted when an unexpected error occurred after the completion of a successful write to the PingChannel socket. The log argument describes the specific error.

26.40.29 PING_CHECK_MGR_STARTED

```
ping channel operations are running, number of threads %1
```

This message is emitted when the ping check channel has been opened and is ready to process requests. The log argument includes the number of threads in the channel's thread pool.

26.40.30 PING_CHECK_MGR_STARTED_SINGLE_THREADED

```
single-threaded ping channel operations are running
```

This message is emitted when the ping check channel has been opened and is ready to process requests in single-threaded mode.

26.40.31 PING_CHECK_MGR_START_PING_CHECK

```
for %1, initiated by %2
```

Logged at debug log level 40. This debug message is emitted when a ping check for an address has been initiated. The log arguments detail the lease address to ping and the query which initiated the check.

26.40.32 PING_CHECK_MGR_STOPPED

```
channel operations have stopped
```

This message is emitted when the ping check channel operations have been stopped.

26.40.33 PING_CHECK_MGR_STOPPING

```
ping channel operations are stopping
```

Logged at debug log level 40. This debug message is emitted when the ping check channel is stopping operations, typically due to configuration event or server shutdown.

26.40.34 PING_CHECK_MGR_SUBNET_CONFIG_FAILED

```
user-context for subnet id %1, contains invalid ping-check %2
```

This error message indicates that a subnet was updated via subnet commands and its 'user-context' contains invalid 'ping-check' configuration. The server will log the error once and then use global ping-check parameters for the subnet until the configuration is corrected.

26.40.35 PING_CHECK_NO_LEASE_OR_LEASE_REUSED

```
Ping check skipped: no lease
```

Logged at debug log level 50. This debug message is emitted when the ping check request made by the server does not contain a lease. This typically happens when a lease is being reused. The ping check will be skipped and the offer processing will continue as normal.

26.40.36 PING_CHECK_PAUSE_FAILED

Pausing ping channel operations failed %1

This error message is emitted when an unexpected error occurred while attempting to pause the ping channel's thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as defect.

26.40.37 PING_CHECK_PAUSE_ILLEGAL

Pausing ping channel operations not allowed %1

This error message is emitted when attempting to pause the ping channel's thread pool. This indicates that a channel thread attempted to use a critical section which would result in a dead-lock. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.40.38 PING_CHECK_PAUSE_PERMISSIONS_FAILED

Permissions check for ping-channel pause failed %1

This error message is emitted when an unexpected error occurred while validating an attempt to pause the ping channel's thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.40.39 PING_CHECK_RESUME_FAILED

Resuming ping channel operations failed %1

This error message is emitted when an unexpected error occurred while attempting to resume operation of the ping channel's thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as defect.

26.40.40 PING_CHECK_UNEXPECTED_READ_ERROR

could not start next socket read %1

This error message occurs when initiating an asynchronous read on the ICMP socket failed in an unexpected fashion. The details of the error are provided as an argument of the log message.

26.40.41 PING_CHECK_UNEXPECTED_WRITE_ERROR

could not start next socket write %1

This error message occurs when initiating an asynchronous write on the ICMP socket failed in an unexpected fashion. The details of the error are provided as an argument of the log message.

26.40.42 PING_CHECK_UNLOAD

Ping Check hooks library has been unloaded

This info message indicates that the Ping Check hooks library has been unloaded.

26.41 RADIUS

26.41.1 RADIUS_ACCESS_BUILD_FAILED

```
building Access-Request failed: %1 for incoming message %2
```

This error message is issued when an error was raised during building of Access-Request.

26.41.2 RADIUS_ACCESS_CACHE_GET

```
host %1 with attributes %2 was retrieved from the cache
```

Logged at debug log level 40. This debug message is issued when a host is retrieved from the host cache.

26.41.3 RADIUS_ACCESS_CACHE_INSERT

```
host %1 with attributes %2 was inserted into the cache
```

Logged at debug log level 40. This debug message is issued when a new host is inserted into the host cache.

26.41.4 RADIUS_ACCESS_CONFLICT

```
query %1 triggers a conflict for %2
```

Logged at debug log level 40. This debug message is issued when a query triggers a conflict with a pending access request for the same identifier.

26.41.5 RADIUS_ACCESS_DROP_PARKED_QUERY

```
access request terminate callback decided to drop the parked query %1
```

Logged at debug log level 40. This debug message is issued when an access request terminates with the decision to drop the parked query.

26.41.6 RADIUS_ACCESS_ERROR

```
Access-Request failed with %1 (%2)
```

This error message is issued when no valid Access-Accept or Access-Reject message was received. The return code and name are logged.

26.41.7 RADIUS_ACCESS_GET_IDENTIFIER

```
identifier %1 of type %2 and User-Name %3 were set from incoming message %4
```

Logged at debug log level 40. This debug message is issued when the host identifier and User-Name attribute were set from an incoming message.

26.41.8 RADIUS_ACCESS_GET_IDENTIFIER_FAILED

```
no identifier of type %1 can be set from incoming message %2, reason: %3
```

This error message is issued when it was not possible to build the host identifier from the incoming message.

26.41.9 RADIUS_ACCESS_HOST_BACKEND_ERROR

```
Configuring access failed during host backend '%1' setup, reason: %2
```

This error message is issued when access/authentication is enabled in the configuration but something in host backend setup went wrong. The name of the host backend and the reason are logged.

26.41.10 RADIUS_ACCESS_MAX_PENDING_REQUESTS

```
query '%1' with identifier '%2' was dropped for too many pending access requests
```

Logged at debug log level 40. This debug message is issued when the number of pending access requests is over the configured limit. The query and its identifier are displayed.

26.41.11 RADIUS_ACCESS_NO_HOST_CACHE

```
Configuring access failed: host cache library not loaded.
```

This error message is issued when access/authentication is enabled in the configuration but no host cache was found. The Radius hook requires Host Cache hook to be loaded to store (cache) parameters received from exchanges with RADIUS server.

26.41.12 RADIUS_ACCESS_ORPHAN

```
orphan pending access request for %1
```

This error message is issued when an access request terminates without the corresponding pending request for the identifier.

26.41.13 RADIUS_ACCESS_RESUME_PARKED_QUERY

```
access request terminate callback resumes processing of parked query %1 in %2
```

Logged at debug log level 40. This debug message is issued when access request terminate callback resumes the processing of a parked query after the subnet select callout point.

26.41.14 RADIUS_ACCESS_SUBNET_RESELECT

```
subnet was reselected from 'ID %1' to 'ID %2'
```

Logged at debug log level 40. This debug message is issued when access/authentication triggered a subnet reselect. The original and new subnet IDs are logged.

26.41.15 RADIUS_ACCESS_TERMINATE_ERROR

```
access request terminate callback got an error: %1
```

This error message is issued when an access request terminates with an unexpected internal error.

26.41.16 RADIUS_ACCOUNTING_ASYNC

Asynchronous send Accounting-Request for NAS port %1 with %2

Logged at debug log level 40. This debug message is issued when starting to send an Accounting-Request message to accounting servers. The NAS port and message attributes are logged.

26.41.17 RADIUS_ACCOUNTING_ASYNC_FAILED

Asynchronous Accounting-Request failed: return code %1 (%2)

Logged at debug log level 40. This debug message is issued when no valid Accounting-Response message was received.

26.41.18 RADIUS_ACCOUNTING_ASYNC_SUCCEED

received valid Accounting-Response (asynchronously)

Logged at debug log level 40. This debug message indicates that a valid Accounting-Response was received.

26.41.19 RADIUS_ACCOUNTING_ERROR

Accounting-Request failed for %1 on event %2 (%3) failed with %4 (%5)

This error message is issued when accounting communication failed. The session Id, the event code and name, and return code and name are logged.

26.41.20 RADIUS_ACCOUNTING_HISTORY_UPDATE_FAILED

failed to insert a record for %1 in the history container

This error message is issued when it was not possible to insert a record in the create timestamp aka history container. This should break the session, i.e. it will not be possible for instance to match start and stop status messages.

26.41.21 RADIUS_ACCOUNTING_NO_HISTORY

failed to find the date the lease for %1 was created

Logged at debug log level 40. This debug message is issued when an address was not found in create timestamp aka history container. This should lead to a accounting session without a start status message.

26.41.22 RADIUS_ACCOUNTING_SYNC

Synchronous send Accounting-Request for NAS port %1 with %2

Logged at debug log level 40. This debug message is issued when starting to send an Accounting-Request message to accounting servers. The NAS port and message attributes are logged.

26.41.23 RADIUS_ACCOUNTING_SYNC_FAILED

Synchronous Accounting-Request failed: return code %1 (%2)

Logged at debug log level 40. This debug message is issued when no valid Accounting-Response message was received.

26.41.24 RADIUS_ACCOUNTING_SYNC_SUCCEED

```
received valid Accounting-Response (synchronously)
```

Logged at debug log level 40. This debug message indicates that a valid Accounting-Response was received.

26.41.25 RADIUS_AUTHENTICATION_ASYNC

```
send Access-Request for NAS port %1 with %2
```

Logged at debug log level 40. This debug message is issued when starting to send an Access-Request message to access servers. The NAS port and message attributes are logged.

26.41.26 RADIUS_AUTHENTICATION_ASYNC_ACCEPTED

```
received valid Access-Accept with %1
```

Logged at debug log level 40. This debug message indicates that a valid Access-Accept message was received. Attributes from the message are logged.

26.41.27 RADIUS_AUTHENTICATION_ASYNC_FAILED

```
Access-Request failed: return code %1 (%2)
```

Logged at debug log level 40. This debug message is issued when no correct Access-Accept or Access-Reject message was received.

26.41.28 RADIUS_AUTHENTICATION_ASYNC_REJECTED

```
received valid Access-Reject with %1
```

Logged at debug log level 40. This debug message indicates that a valid Access-Reject message was received. Attributes from the message are logged.

26.41.29 RADIUS_AUTHENTICATION_SYNC

```
send Access-Request for NAS port %1 with %2
```

Logged at debug log level 40. This debug message is issued when starting to send an Access-Request message to access servers. The NAS port and message attributes are logged.

26.41.30 RADIUS_AUTHENTICATION_SYNC_ACCEPTED

```
received valid Access-Accept with %1
```

Logged at debug log level 40. This debug message indicates that a valid Access-Accept message was received. Attributes from the message are logged.

26.41.31 RADIUS_AUTHENTICATION_SYNC_FAILED

```
Access-Request failed: return code %1 (%2)
```

Logged at debug log level 40. This debug message is issued when no correct Access-Accept or Access-Reject message was received.

26.41.32 RADIUS_AUTHENTICATION_SYNC_REJECTED

```
received valid Access-Reject with %1
```

Logged at debug log level 40. This debug message indicates that a valid Access-Reject message was received. Attributes from the message are logged.

26.41.33 RADIUS_BACKEND_GET4

```
spurious lookup for IPv4 subnet %1 and id %2 of type %3
```

Logged at debug log level 40. This debug message is issued when the radius host backend is unexpectedly called for looking for an IPv4 entry. Details of the lookup are logged.

26.41.34 RADIUS_BACKEND_GET6

```
spurious lookup for IPv6 subnet %1 and id %2 of type %3
```

Logged at debug log level 40. This debug message is issued when the radius host backend is unexpectedly called for looking for an IPv6 entry. Details of the lookup are logged.

26.41.35 RADIUS_CLEANUP_EXCEPTION

```
Exception on RADIUS cleanup: %1
```

This warning message is issued when there is an exception thrown when destroying an object in the RADIUS hook library. The exception is not allowed to continue propagating to not obfuscate another exception, so it is logged. It generally means a programmatic error and should be reported to ISC, but could also be harmless. The argument provides the detailed error message.

26.41.36 RADIUS_CONFIGURATION_FAILED

```
failed to configure Radius hooks library: %1
```

This error message is issued when there is an error configuring the Radius hooks library. The argument provides the detailed error message.

26.41.37 RADIUS_DECODE_MESSAGE

```
Decoded message '%1' (%2) id %3 length %4 with %5 attributes.
```

Logged at debug log level 40. This debug message is issued when a message is decoded. The message type name and value, the identifier, the length and the number of attributes are displayed.

26.41.38 RADIUS_DEINIT_OK

```
unloading Radius hooks library successful
```

This informational message indicates that the Radius hooks library has been unloaded successfully.

26.41.39 RADIUS_ENCODE_MESSAGE

Encoded message '%1' (%2) id %3 length %4 with %5 attributes.

Logged at debug log level 40. This debug message is issued when a message is encoded. The message type name and value, the identifier, the length and the number of attributes are displayed.

26.41.40 RADIUS_EXCHANGE_FAILED

Exchange %1 failed: %2

This error message is issued when an exchange terminates with an error. The exchange identifier and the error message are displayed.

26.41.41 RADIUS_EXCHANGE_OPEN_FAILED

Open socket for exchange %1 failed: %2

This error message is issued when an exchange failed to open a new socket. The exchange identifier and the error message are displayed.

26.41.42 RADIUS_EXCHANGE_RECEIVED

Exchange %1 received %2 bytes.

Logged at debug log level 40. This debug message is issued when an exchange received a response. The exchange identifier and the response size are displayed.

26.41.43 RADIUS_EXCHANGE_RECEIVED_ACCESS_ACCEPT

Exchange %1 received an Access-Accept.

Logged at debug log level 40. This debug message is issued when an exchange received an Access-Accept response. The exchange identifier is displayed.

26.41.44 RADIUS_EXCHANGE_RECEIVED_ACCESS_REJECT

Exchange %1 received an Access-Reject.

Logged at debug log level 40. This debug message is issued when an exchange received an Access-Reject response. The exchange identifier is displayed.

26.41.45 RADIUS_EXCHANGE_RECEIVED_ACCOUNTING_RESPONSE

Exchange %1 received an Accounting-Response.

Logged at debug log level 40. This debug message is issued when an exchange received an Accounting-Response response. The exchange identifier is displayed.

26.41.46 RADIUS_EXCHANGE_RECEIVED_BAD_RESPONSE

Exchange %1 received a bad response: %2

This error message is issued when an exchange received a bad response. The exchange identifier and the error message are displayed.

26.41.47 RADIUS_EXCHANGE_RECEIVED_MISMATCH

Exchange %1: received response with identifier %2 when %3 was expected.

This error message is issued when the sent request and the received response have different identifiers. The exchange identifier and the two RADIUS message identifiers are displayed.

26.41.48 RADIUS_EXCHANGE_RECEIVED_RESPONSE

Exchange %1 received response: %2

Logged at debug log level 40. This debug message is issued at the end of the reception routine. The exchange identifier and the error code are displayed.

26.41.49 RADIUS_EXCHANGE_RECEIVED_UNEXPECTED

Exchange %1: sent %2, received unexpected %3

This error message is issued when the sent request and the received response do not match. The exchange identifier and the two RADIUS message codes are displayed.

26.41.50 RADIUS_EXCHANGE_RECEIVE_FAILED

Receive for exchange %1 failed: %2

This error message is issued when an exchange failed to receive a message. The exchange identifier and the error message are displayed.

26.41.51 RADIUS_EXCHANGE_SEND_FAILED

Send for exchange %1 failed: %2

This error message is issued when an exchange failed to send a message. The exchange identifier and the error message are displayed.

26.41.52 RADIUS_EXCHANGE_SEND_NEW

Exchange %1 sends %2 bytes to new server[%3] %4 on port %5

Logged at debug log level 40. This debug message is issued when an exchange sends a message to a new server. The exchange identifier, message size, server index, address and port are displayed.

26.41.53 RADIUS_EXCHANGE_SEND_RETRY

Exchange %1 sends %2 bytes for the %3 try.

Logged at debug log level 40. This debug message is issued when an exchange sends a message to a new server. The exchange identifier, message size and retry counter are displayed.

26.41.54 RADIUS_EXCHANGE_SENT

Exchange %1 sent %2 bytes.

Logged at debug log level 40. This debug message is issued when an exchange sent a request and is ready to receive the response. The exchange identifier and request size are displayed.

26.41.55 RADIUS_EXCHANGE_START

Start exchange %1

Logged at debug log level 40. This debug message is issued when an exchange starts. The exchange identifier is displayed.

26.41.56 RADIUS_EXCHANGE_SYNC_RETURN

Synchronous exchange %1 returns with %2

Logged at debug log level 40. This debug message is issued when a synchronous exchange returns. The exchange identifier and the error/return code are displayed.

26.41.57 RADIUS_EXCHANGE_TERMINATE

Exchange %1 terminates with %2

Logged at debug log level 40. This debug message is issued when an exchange terminates with success. The exchange identifier and the return code are displayed.

26.41.58 RADIUS_EXCHANGE_TIMEOUT

Exchange %1 timeout

This error message is issued when an exchange failed on timeout. The exchange identifier is displayed.

26.41.59 RADIUS_HOOK_FAILED

processing for hook %1 failed: %2

This error message is issued when processing at a standard hook point failed. The reason of the failure is displayed.

26.41.60 RADIUS_INIT_OK

loading Radius hooks library successful

This informational message indicates that the Radius hooks library has been loaded successfully. Enjoy!

26.41.61 RADIUS_INTEGER_ATTRIBUTE_FROM_BYTES_FAILED

Creating an integer attribute %1 '%2' failed: %3

This error message is issued when an integer attribute can't be created. Attribute type, name and error message are displayed.

26.41.62 RADIUS_INTEGER_ATTRIBUTE_FROM_TEXT_FAILED

Creating an integer attribute %1 '%2' from %3 failed.

This error message is issued when an integer attribute can't be created. Attribute type, name and bad submitted value are displayed.

26.41.63 RADIUS_IPADDR_ATTRIBUTE_FROM_BYTES_FAILED

Creating an IP address attribute %1 '%2' failed: %3

This error message is issued when an IP address attribute can't be created. Attribute type, name and error message are displayed.

26.41.64 RADIUS_IPADDR_ATTRIBUTE_FROM_TEXT_FAILED

Creating an IP address attribute %1 '%2' from %3 failed.

This error message is issued when an IP address attribute can't be created. Attribute type, name and bad submitted value are displayed.

26.41.65 RADIUS_IPV6ADDR_ATTRIBUTE_FROM_BYTES_FAILED

Creating an IPv6 address attribute %1 '%2' failed: %3

This error message is issued when an IPv6 address attribute can't be created. Attribute type, name and error message are displayed.

26.41.66 RADIUS_IPV6ADDR_ATTRIBUTE_FROM_TEXT_FAILED

Creating an IPv6 address attribute %1 '%2' from %3 failed.

This error message is issued when an IPv6 address attribute can't be created. Attribute type, name and bad submitted value are displayed.

26.41.67 RADIUS_IPV6PREFIX_ATTRIBUTE_FROM_BYTES_FAILED

Creating an IPv6 prefix attribute %1 '%2' failed: %3

This error message is issued when an IPv6 prefix attribute can't be created. Attribute type, name and error message are displayed.

26.41.68 RADIUS_IPV6PREFIX_ATTRIBUTE_FROM_TEXT_FAILED

Creating an IPv6 prefix attribute %1 '%2' from %3 failed.

This error message is issued when an IPv6 prefix attribute can't be created. Attribute type, name and bad submitted value are displayed.

26.41.69 RADIUS_PAUSE_FAILED

Pausing the RADIUS thread pool failed: %1

This error message is emitted when an unexpected error occurred while validating an attempt to pause the thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.41.70 RADIUS_PAUSE_ILLEGAL

Pausing the RADIUS thread pool not allowed: %1

This error message is emitted when attempting to pause the thread pool. This indicates that a thread attempted to use a critical section which would result in a dead-lock. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.41.71 RADIUS_PAUSE_PERMISSIONS_FAILED

Checking for permissions to pause the RADIUS thread pool failed: %1

This error message is emitted when an unexpected error occurred while validating an attempt to pause the thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as a defect.

26.41.72 RADIUS_REPLY_MESSAGE_ATTRIBUTE

Message %1 on exchange %2 has a Reply-Message attribute with value '%3'.

This informational message is issued when a Reply-Message attribute is found in a message received from the RADIUS server. One log message is printed per attribute. A log message can contain multiple attributes, so there can be multiple log messages per RADIUS message. It displays the message identifier, the exchange identifier and the value of the Reply-Message attribute.

26.41.73 RADIUS_RESUME_FAILED

Resuming the RADIUS thread pool failed: %1

This error message is emitted when an unexpected error occurred while attempting to resume the thread pool. This error is highly unlikely and indicates a programmatic issue that should be reported as defect.

26.41.74 RADIUS_SERVER_CONFIGURED

configured an %1 server: %2

This informational message is issued when a RADIUS server is configured. The kind of the server (access or accounting) and configuration details are logged with the secret replaced by stars.

26.41.75 RADIUS_SESSION_HISTORY_APPEND_FAILED

```
appending of a new record for %1 to the session history file failed: %2
```

This warning message is issued when appending a new record to the session history file failed. The address of the new record and the reason of the failure are displayed.

26.41.76 RADIUS_SESSION_HISTORY_LOADED

```
loading of the session history file succeeded: read %1 records including %2 active_↵
↵records
```

This informational message is issued when loading of the session history file completed with success. Numbers of loaded and active records are displayed.

26.41.77 RADIUS_SESSION_HISTORY_LOAD_FAILED

```
loading of the session history file
```

failed: loaded %1, skipped %2 and active %3 records. This error message is issued when loading of the session history file did not completed with success. Numbers of loaded, skipped and active records displayed.

26.41.78 RADIUS_SESSION_HISTORY_OPENED

```
opening of the session history file %1 succeeded
```

This informational message is issued when opening of the CSV file providing session history persistence succeeded. The name of the file is displayed.

26.41.79 RADIUS_SESSION_HISTORY_OPEN_FAILED

```
opening of the session history file %1 failed: %2
```

This error message is issued when opening of the CSV file providing session history persistence failed. The name of the file and the reason of the failure are displayed.

26.41.80 RADIUS_SESSION_HISTORY_STORED

```
Storing to the session history file succeeded: stored %1 records
```

This informational message is issued when writing to a new session history file completed with success. The number of stored records is displayed.

26.41.81 RADIUS_SESSION_HISTORY_STORE_FAILED

```
Writing to the session history file %1 failed: %2 (stored %3 over %4 records)
```

This error message is issued when writing to a new session history file failed. The name of the file, the reason of the failure, the number of stored records before the failure and the expected number of records are displayed.

26.41.82 RADIUS_THREAD_POOL_STARTED

RADIUS thread pool started with %1 threads.

This informational message is issued when the thread pool is started. The number of threads is displayed.

26.42 RBAC

26.42.1 RBAC_CONFIGURED_ACLS

RBAC hooks library has configured %1 named ACLs.

Logged at debug log level 40. This debug message indicates that some named access control lists have been configured. The count is displayed.

26.42.2 RBAC_CONFIGURED_COMMANDS

RBAC hooks library has configured %1 extra commands.

Logged at debug log level 40. This debug message indicates that some extra commands have been configured. The count is displayed.

26.42.3 RBAC_CONFIGURED_ROLES

hooks library has configured %1 roles.

Logged at debug log level 40. This debug message indicates that some roles have been configured. The count is displayed.

26.42.4 RBAC_HTTP_AUTH_ERROR

Error in http_auth callout: %1.

This error message indicates that an error has been raised in http_auth callout by the RBAC hooks library. The argument details the error.

26.42.5 RBAC_HTTP_AUTH_RESPONSE

RBAC hooks library has returned response: %1.

This info message indicates that the RBAC hooks library has returned a response in http_auth callout. The response is summarized.

26.42.6 RBAC_LOAD_FAILED

RBAC hooks library failed to load: %1.

This error message indicates that an error occurred attempting to load the RBAC hooks library. The argument details the error.

26.42.7 RBAC_LOAD_OK

```
RBAC hooks library loaded successfully.
```

This info message indicates that the RBAC hooks library has been loaded successfully.

26.42.8 RBAC_READ_API_FILES

```
RBAC hooks library read API files from '%1' getting %2 commands, %3 access types and %4 hooks.
```

This info message indicates that the RBAC hooks library has read API files from the configured directory. The name of the directory and some statistics are displayed.

26.42.9 RBAC_TRACE_HTTP_AUTH_ACCEPT

```
Role configuration '%1' for role '%2' has accepted command '%3'.
```

Logged at debug log level 40. The command has been accepted in http_auth callout. The role configuration name, the role name and the command are displayed.

26.42.10 RBAC_TRACE_HTTP_AUTH_BAD_BODY_TYPE

```
Bad body type in JSON request in http_auth callout.
```

Logged at debug log level 40. The http_auth callout has been called with a bad body type in the JSON request. The RBAC hooks library immediately returns. This is an error condition.

26.42.11 RBAC_TRACE_HTTP_AUTH_BAD_COMMAND_TYPE

```
Bad command type in JSON request in http_auth callout.
```

Logged at debug log level 40. The http_auth callout has been called with a bad command type in the JSON request. The RBAC hooks library immediately returns. This is an error condition.

26.42.12 RBAC_TRACE_HTTP_AUTH_COMMAND

```
Command '%1' in http_auth callout.
```

Logged at debug log level 40. The http_auth callout has been called with the displayed command.

26.42.13 RBAC_TRACE_HTTP_AUTH_DISABLED

```
RBAC hooks library is disabled in http_auth callout.
```

Logged at debug log level 40. The http_auth callout has been called with the RBAC hooks library disabled i.e. with no role assigned.

26.42.14 RBAC_TRACE_HTTP_AUTH_EMPTY_BODY

```
Empty body in JSON request in http_auth callout.
```

Logged at debug log level 40. The http_auth callout has been called with an empty body in the JSON request. The RBAC hooks library immediately returns. This is an error condition.

26.42.15 RBAC_TRACE_HTTP_AUTH_NO_COMMAND

No command entry in JSON request in http_auth callout.

Logged at debug log level 40. The http_auth callout has been called without a command entry in the JSON request. The RBAC hooks library immediately returns. This is an error condition.

26.42.16 RBAC_TRACE_HTTP_AUTH_NO_JSON

No JSON request in http_auth callout.

Logged at debug log level 40. The http_auth callout has been called with a non JSON request. The RBAC hooks library immediately returns. This is an error condition.

26.42.17 RBAC_TRACE_HTTP_AUTH_NO_REQUEST

No request in http_auth callout.

Logged at debug log level 40. The http_auth callout has been called without a request. The RBAC hooks library immediately returns. This is an error condition.

26.42.18 RBAC_TRACE_HTTP_AUTH_NO_TLS_REJECT

Non TLS request has been rejected.

Logged at debug log level 40. The non TLS request has been rejected in http_auth callout.

26.42.19 RBAC_TRACE_HTTP_AUTH_REJECT

Role configuration '%1' for role '%2' has rejected command '%3'.

Logged at debug log level 40. The command has been rejected in http_auth callout. The role configuration name, the role name and the command are displayed.

26.42.20 RBAC_TRACE_HTTP_AUTH_RESPONSE

Response in http_auth callout: %1.

Logged at debug log level 40. The http_auth callout has been called with a response. The RBAC hooks library immediately returns. The response is summarized.

26.42.21 RBAC_TRACE_HTTP_AUTH_ROLE

Assigned role '%1' in http_auth callout.

Logged at debug log level 40. The displayed role has been assigned in the http_auth callout.

26.42.22 RBAC_TRACE_HTTP_RESPONSE_BAD_BODY_TYPE

Bad body type in JSON response in http_response callout.

Logged at debug log level 40. The http_response callout has been called with a bad body type in the JSON response which is likely an error response. The RBAC hooks library immediately returns.

26.42.23 RBAC_TRACE_HTTP_RESPONSE_CONTEXT

Retrieved command '%1' and role config '%2' in http_response callout.

Logged at debug log level 40. The command and the role config have been retrieved from the request context. They are displayed.

26.42.24 RBAC_TRACE_HTTP_RESPONSE_DISABLED

RBAC hooks library is disabled in http_response callout.

Logged at debug log level 40. The http_response callout has been called with the RBAC hooks library disabled i.e. with no role assigned.

26.42.25 RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY

Empty body in JSON response in http_response callout.

Logged at debug log level 40. The http_response callout has been called with an empty body in the JSON response. The RBAC hooks library immediately returns. This is an error condition.

26.42.26 RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY_LIST

Empty list in JSON response in http_response callout.

Logged at debug log level 40. The http_response callout has been called with an empty body list in the JSON response. The RBAC hooks library immediately returns. This is an error condition.

26.42.27 RBAC_TRACE_HTTP_RESPONSE_MODIFIED

The response has been modified by a response filter in http_response callout.

Logged at debug log level 40. A response filter has modified the HTTP response in http_response callout.

26.42.28 RBAC_TRACE_HTTP_RESPONSE_NO_ARGUMENTS

No request or response in http_response callout.

Logged at debug log level 40. The http_response callout has been called without request or response. The RBAC hooks library immediately returns. This is an error condition.

26.42.29 RBAC_UNLOAD_OK

RBAC hooks library unloaded successfully.

This info message indicates that the RBAC hooks library has been unloaded successfully.

26.43 RUN

26.43.1 RUN_SCRIPT_LOAD

```
Run Script hooks library has been loaded
```

This info message indicates that the Run Script hooks library has been loaded.

26.43.2 RUN_SCRIPT_LOAD_ERROR

```
Run Script hooks library failed: %1
```

This error message indicates an error during loading the Run Script hooks library. The details of the error are provided as argument of the log message.

26.43.3 RUN_SCRIPT_UNLOAD

```
Run Script hooks library has been unloaded
```

This info message indicates that the RunScript hooks library has been unloaded.

26.44 START

26.44.1 START_REKEY_TIMER

```
started timer handling rekey for server %1 in %2 seconds.
```

Logged at debug log level 40. This debug message is issued when starting the rekey timer to handle new keys for this server when at least one key is currently available. The first argument specifies the server identifier and the second argument specifies the time interval when the next key processing will be attempted.

26.44.2 START_RETRY_TIMER

```
started timer handling retry for server %1 in %2 seconds.
```

Logged at debug log level 40. This debug message is issued when starting the retry timer to handle new keys for this server when there is no key currently available. The first argument specifies the server identifier and the second argument specifies the time interval when the next key processing will be attempted.

26.45 STAT

26.45.1 STAT_CMDS_DEINIT_OK

```
unloading Stat Commands hooks library successful
```

This info message indicates that the Stat Commands hooks library has been removed successfully.

26.45.2 STAT_CMDS_INIT_OK

```
loading Stat Commands hooks library successful
```

This info message indicates that the Stat Commands hooks library has been loaded successfully. Enjoy!

26.45.3 STAT_CMDS_LEASE4_FAILED

```
stat-lease4-get command failed: reason: %1
```

The stat-lease4-get command has failed. The reason for failure is logged.

26.45.4 STAT_CMDS_LEASE4_GET

```
stat-lease4-get command successful, parameters: %1 rows found: %2
```

The stat-lease4-get command has been successful. The log will contain the parameters supplied and the number of rows found.

26.45.5 STAT_CMDS_LEASE4_GET_FAILED

```
stat-lease4-get command failed: parameters: %1, reason: %2
```

The stat-lease4-get command has failed. Both the parameters supplied and the reason for failure are logged.

26.45.6 STAT_CMDS_LEASE4_GET_INVALID

```
stat-lease4-get command is malformed or invalid, reason: %1
```

The stat-lease4-get command was either malformed or contained invalid parameters. A detailed explanation should be logged.

26.45.7 STAT_CMDS_LEASE4_GET_NO_SUBNETS

```
stat-lease4-get, parameters: %1, %2"
```

The parameters submitted with stat-lease4-get were valid but excluded all known subnets. The parameters supplied along with an explanation should be logged.

26.45.8 STAT_CMDS_LEASE4_ORPHANED_STATS

```
stat-lease4-get command omitted statistics for one or more non-existent subnets
```

Logged at debug log level 40. During processing the stat-lease4-get found statistics for subnet IDs for non-existent subnets. These values were omitted from the command response returned to the user. This may occur when subnets have been removed from the configuration in a manner that did not also remove the statistics. While the existence of such statistics is not harmful, steps should be considered to remove them. For memfile lease storage, the problem should disappear upon configuration reload or server restart. For database lease storage the issue is more complicated and as of Kea 2.0.0 we do not yet have a clean solution.

26.45.9 STAT_CMDS_LEASE6_FAILED

```
stat-lease6-get command failed: reason: %1
```

The stat-lease6-get command has failed. The reason for failure is logged.

26.45.10 STAT_CMDS_LEASE6_GET

```
stat-lease6-get command successful, parameters: %1 rows found: %2
```

The stat-lease6-get command has been successful. The log will contain the parameters supplied and the number of rows found.

26.45.11 STAT_CMDS_LEASE6_GET_FAILED

```
stat-lease6-get command failed: parameters: %1, reason: %2
```

The stat-lease6-get command has failed. Both the parameters supplied and the reason for failure are logged.

26.45.12 STAT_CMDS_LEASE6_GET_INVALID

```
stat-lease6-get command is malformed or invalid, reason: %1
```

The stat-lease6-get command was either malformed or contained invalid parameters. A detailed explanation should be logged.

26.45.13 STAT_CMDS_LEASE6_GET_NO_SUBNETS

```
stat-lease6-get, parameters: %1, %2"
```

The parameters submitted with stat-lease6-get were valid but excluded all known subnets. The parameters supplied along with an explanation should be logged.

26.45.14 STAT_CMDS_LEASE6_ORPHANED_STATS

```
stat-lease6-get command omitted statistics for one or more non-existent subnets
```

Logged at debug log level 40. During processing the stat-lease4-get found statistics for subnet IDs for non-existent subnets. These values were omitted from the command response returned to the user. This may occur when subnets have been removed from the configuration in a manner that did not also remove the statistics. While the existence of such statistics is not harmful, steps should be considered to remove them. For memfile lease storage, the problem should disappear upon configuration reload or server restart. For database lease storage the issue is more complicated and as of Kea 2.0.0 we do not yet have a clean solution.

26.46 SUBNET

26.46.1 SUBNET_CMDS_DEINIT_OK

```
unloading Subnet Commands hooks library successful
```

This informational message indicates that the Host Commands hooks library has been unloaded successfully.

26.46.2 SUBNET_CMDS_INIT_FAILED

```
loading Subnet Commands hooks library failed: %1
```

This error message indicates an error during loading the Subnet Commands hooks library. The details of the error are provided as argument of the log message.

26.46.3 SUBNET_CMDS_INIT_OK

```
loading Subnet Commands hooks library successful
```

This informational message indicates that the Host Commands hooks library has been loaded successfully. Enjoy!

26.46.4 SUBNET_CMDS_NETWORK4_ADD_FAILED

```
failed to add new IPv4 network: %1
```

This error message is issued when the Subnet Commands hooks library fails to add a new IPv4 network to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.5 SUBNET_CMDS_NETWORK4_DEL_FAILED

```
failed to delete IPv4 network: %1
```

This error message is issued when the Subnet Commands hooks library fails to delete an IPv4 network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.6 SUBNET_CMDS_NETWORK4_GET_FAILED

```
failed to return an IPv4 network: %1
```

This error message is issued when the server fails to return an IPv4 network in response to 'network4-get' command. The argument details the reason for failure. The error message will be returned to the controlling client with the error status code. This error may occur when the received command has invalid structure, has not allowed parameters or lacks required parameters. It will also be returned when the command syntax is correct but no network was found.

26.46.7 SUBNET_CMDS_NETWORK4_LIST_FAILED

```
failed to return a list of IPv4 networks: %1
```

This error message is issued when the Subnet Commands hooks library fails to return a list of IPv4 networks requested with 'network4-list' command. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.8 SUBNET_CMDS_NETWORK4_SUBNET_ADD_FAILED

```
failed to add existing IPv4 subnet to a shared network: %1
```

This error message is issued when the Subnet Commands hooks library fails to add existing IPv4 subnet to existing shared network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.9 SUBNET_CMDS_NETWORK4_SUBNET_DEL_FAILED

```
failed to remove a IPv4 subnet from a shared network: %1
```

This error message is issued when the Subnet Commands hooks library fails to remove existing IPv4 subnet to existing shared network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.10 SUBNET_CMDS_NETWORK6_ADD_FAILED

failed to add new IPv6 network: %1

This error message is issued when the Subnet Commands hooks library fails to add a new IPv6 network to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.11 SUBNET_CMDS_NETWORK6_DEL_FAILED

failed to delete IPv6 network: %1

This error message is issued when the Subnet Commands hooks library fails to delete an IPv6 network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.12 SUBNET_CMDS_NETWORK6_GET_FAILED

failed to return an IPv6 network: %1

This error message is issued when the server fails to return an IPv6 network in response to 'network4-get' command. The argument details the reason for failure. The error message will be returned to the controlling client with the error status code. This error may occur when the received command has invalid structure, has not allowed parameters or lacks required parameters. It will also be returned when the command syntax is correct but no network was found.

26.46.13 SUBNET_CMDS_NETWORK6_LIST_FAILED

failed to return a list of IPv6 networks: %1

This error message is issued when the Subnet Commands hooks library fails to return a list of IPv4 networks requested with 'network6-list' command. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.14 SUBNET_CMDS_NETWORK6_SUBNET_ADD_FAILED

failed to add existing IPv6 subnet to a shared network: %1

This error message is issued when the Subnet Commands hooks library fails to add existing IPv6 subnet to existing shared network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.15 SUBNET_CMDS_NETWORK6_SUBNET_DEL_FAILED

failed to remove a IPv6 subnet from a shared network: %1

This error message is issued when the Subnet Commands hooks library fails to remove existing IPv6 subnet to existing shared network. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.16 SUBNET_CMDS_NETWORK_ADD

```
successfully added shared network %1
```

This informational message is issued when the Subnet Commands hooks library successfully adds a shared network as a result of receiving a 'network4-add' or 'network6-add' command'. The argument represents the name of added shared network.

26.46.17 SUBNET_CMDS_NETWORK_DEL

```
successfully deleted shared network %1
```

This informational message is issued when the Subnet Commands hooks library successfully deletes a shared network as a result of receiving a 'network4-del' or 'network6-del' command'. The argument represents the name of deleted shared network.

26.46.18 SUBNET_CMDS_NETWORK_GET

```
successfully retrieved shared network %1
```

This informational message is issued when the Subnet Commands hooks library successfully retrieves a shared network as a result of receiving a 'network4-get' or 'network6-get' command'. The argument represents the name of retrieved shared network.

26.46.19 SUBNET_CMDS_NETWORK_GET_EMPTY

```
specified shared network is not found: %1
```

This informational message is issued when the Subnet Commands hooks library found no matching shared network as a result of receiving a 'network4-get' or 'network6-get' command'.

26.46.20 SUBNET_CMDS_NETWORK_LIST

```
successfully retrieved list of %1 %2 shared networks
```

This informational message is issued when the Subnet Commands hooks library successfully retrieves a list of shared networks as a result of receiving 'network4-list' or 'network6-list' command. The first argument specifies a number of networks retrieved. The second parameter specifies a protocol type: 'IPv4' or 'IPv6'.

26.46.21 SUBNET_CMDS_NETWORK_LIST_EMPTY

```
no %1 shared networks listed
```

This informational message is issued when the Subnet Commands hooks library successfully processes the 'network4-list' or 'network6-list' command but no shared network has been found. This indicates that the server configuration contains no shared networks of the specific type. The argument specifies a protocol type: 'IPv4' or 'IPv6'.

26.46.22 SUBNET_CMDS_NETWORK_SUBNET_ADD

```
%1 subnet %2 (id %3) added to shared network %4
```

This informational message indicates that specified subnet (address family given in parameter 1, details in parameters 2 and 3) is now part of a shared network. This is a successful result of either network4-subnet-add or network6-subnet-add commands.

26.46.23 SUBNET_CMDS_NETWORK_SUBNET_DEL

```
%1 subnet %2 (id %3) removed from shared network %4
```

This informational message indicates that specified subnets (address family given in parameter 1, details in parameters 2 and 3) is no longer part of a shared network. The subnet remains in configuration, but is a stand alone subnet. This is a successful result of either network4-subnet-del or network6-subnet-del commands.

26.46.24 SUBNET_CMDS_SUBNET4_ADD_FAILED

```
failed to add new IPv4 subnet: %1
```

This error message is issued when the Subnet Commands hooks library fails to add a new IPv4 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.25 SUBNET_CMDS_SUBNET4_DELTA_ADD_FAILED

```
failed to update IPv4 subnet: %1
```

This error message is issued when the Subnet Commands hooks library fails to update by adding a delta in a IPv4 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.26 SUBNET_CMDS_SUBNET4_DELTA_DEL_FAILED

```
failed to update IPv4 subnet: %1
```

This error message is issued when the Subnet Commands hooks library fails to update by removing a delta in a IPv4 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.27 SUBNET_CMDS_SUBNET4_DEL_FAILED

```
failed to delete IPv4 subnet: %1
```

This error message is issued when the Subnet Commands hooks library fails to delete an IPv4 subnet. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.28 SUBNET_CMDS_SUBNET4_GET_FAILED

```
failed to return an IPv4 subnet: %1
```

This error message is issued when the server fails to return an IPv4 subnet in response to 'subnet4-get' command. The argument details the reason for failure. The error message will be returned to the controlling client with the error status code. This error may occur when the received command has invalid structure, has not allowed parameters or lacks required parameters. It will also be returned when the command syntax is correct but no subnet was found.

26.46.29 SUBNET_CMDS_SUBNET4_LIST_FAILED

failed to return a list of IPv4 subnets: %1

This error message is issued when the Subnet Commands hooks library fails to return a list of IPv4 subnets requested with 'subnet4-list' command. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.30 SUBNET_CMDS_SUBNET4_UPDATE_FAILED

failed to update IPv4 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to update a IPv4 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.31 SUBNET_CMDS_SUBNET6_ADD_FAILED

failed to add new IPv6 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to add a new IPv6 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.32 SUBNET_CMDS_SUBNET6_DELTA_ADD_FAILED

failed to update IPv6 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to update by adding a delta in a IPv6 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.33 SUBNET_CMDS_SUBNET6_DELTA_DEL_FAILED

failed to update IPv6 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to update by removing a delta in a IPv6 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.34 SUBNET_CMDS_SUBNET6_DEL_FAILED

failed to delete IPv6 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to delete an IPv6 subnet. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.35 SUBNET_CMDS_SUBNET6_GET_FAILED

failed to return an IPv6 subnet: %1

This error message is issued when the server fails to return an IPv6 subnet in response to 'subnet4-get' command. The argument details the reason for failure. The error message will be returned to the controlling client with the error status code. This error may occur when the received command has invalid structure, has not allowed parameters or lacks required parameters. It will also be returned when the command syntax is correct but no subnet was found.

26.46.36 SUBNET_CMDS_SUBNET6_LIST_FAILED

failed to return a list of IPv6 subnets: %1

This error message is issued when the Subnet Commands hooks library fails to return a list of IPv6 subnets requested with 'subnet6-list' command. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.37 SUBNET_CMDS_SUBNET6_UPDATE_FAILED

failed to update IPv6 subnet: %1

This error message is issued when the Subnet Commands hooks library fails to update a IPv6 subnet to the server configuration. The reason for failure is provided within the error message. The error message will be returned to the controlling client with the error status code.

26.46.38 SUBNET_CMDS_SUBNET_ADD

successfully added subnet %1 having id %2

This informational message is issued when the Subnet Commands hooks library successfully adds a subnet as a result of receiving a 'subnet4-add' or 'subnet6-add' command'. The first parameter specifies an added subnet prefix. The second parameter specifies a subnet identifier.

26.46.39 SUBNET_CMDS_SUBNET_DEL

successfully deleted subnet %1 having id %2

This informational message is issued when the Subnet Commands hooks library successfully deletes a subnet as a result of receiving a 'subnet4-del' or 'subnet6-del' command'. The first parameter specifies a deleted subnet prefix. The second parameter specifies a subnet identifier.

26.46.40 SUBNET_CMDS_SUBNET_GET

successfully retrieved subnet %1 having id %2

This informational message is issued when the Subnet Commands hooks library successfully retrieves a subnet as a result of receiving a 'subnet4-get' or 'subnet6-get' command'. The first parameter specifies a retrieved subnet prefix. The second parameter specifies a subnet identifier.

26.46.41 SUBNET_CMDS_SUBNET_GET_EMPTY

specified subnet is not found: %1

This informational message is issued when the Subnet Commands hooks library found no a matching subnet as a result of receiving a 'subnet4-get' or 'subnet6-get' command'.

26.46.42 SUBNET_CMDS_SUBNET_LIST

```
successfully retrieved list of %1 %2 subnets
```

This informational message is issued when the Subnet Commands hooks library successfully retrieves a list of subnets as a result of receiving 'subnet4-list' or 'subnet6-list' command. The first argument specifies a number of subnets retrieved. The second parameter specifies a protocol type: 'IPv4' or 'IPv6'.

26.46.43 SUBNET_CMDS_SUBNET_LIST_EMPTY

```
no %1 subnets listed
```

This informational message is issued when the Subnet Commands hooks library successfully processes the 'subnet4-list' or 'subnet6-list' command but no subnets have been found. This indicates that the server configuration contains no subnets of the specific type. This is not an error condition but it is unusual case for the DHCP service. The argument specifies a type of the subnets being listed, i.e. 'IPv4' or 'IPv6'.

26.46.44 SUBNET_CMDS_SUBNET_UPDATE

```
successfully updated subnet %1 having id %2
```

This informational message is issued when the Subnet Commands hooks library successfully updates a subnet as a result of receiving a 'subnet4-update' or 'subnet6-update' command'. The first parameter specifies an updated subnet prefix. The second parameter specifies a subnet identifier.

26.47 TCP

26.47.1 TCP_CLIENT_REQUEST_RECEIVED

```
received TCP request from %1
```

Logged at debug log level 40. This debug message is issued when the server finished receiving a TCP request from the remote endpoint. The address of the remote endpoint is specified as an argument.

26.47.2 TCP_CONNECTION_REJECTED_BY_FILTER

```
connection from %1 has been denied by the connection filter.
```

Logged at debug log level 50. This debug message is issued when the server's connection filter rejects a new connection based on the client's ip address.

26.47.3 TCP_CONNECTION_SHUTDOWN

```
shutting down TCP connection from %1
```

Logged at debug log level 40. This debug message is issued when one of the TCP connections is shut down. The connection can be stopped as a result of an error or after the successful message exchange with a client.

26.47.4 TCP_CONNECTION_SHUTDOWN_FAILED

```
shutting down TCP connection failed
```

This error message is issued when an error occurred during shutting down a TCP connection with a client.

26.47.5 TCP_CONNECTION_STOP

```
stopping TCP connection from %1
```

Logged at debug log level 40. This debug message is issued when one of the TCP connections is stopped. The connection can be stopped as a result of an error or after the successful message exchange with a client.

26.47.6 TCP_CONNECTION_STOP_FAILED

```
stopping TCP connection failed
```

This error message is issued when an error occurred during closing a TCP connection with a client.

26.47.7 TCP_DATA_RECEIVED

```
received %1 bytes from %2
```

Logged at debug log level 55. This debug message is issued when the server receives a chunk of data from the remote endpoint. This may include the whole request or only a part of the request. The first argument specifies the amount of received data. The second argument specifies an address of the remote endpoint which produced the data.

26.47.8 TCP_DATA_SENT

```
send %1 bytes to %2
```

Logged at debug log level 55. This debug message is issued when the server sends a chunk of data to the remote endpoint. This may include the whole response or only a part of the response. The first argument specifies the amount of sent data. The second argument specifies an address of the remote endpoint.

26.47.9 TCP_IDLE_CONNECTION_TIMEOUT_OCCURRED

```
closing connection with %1 as a result of a timeout
```

Logged at debug log level 50. This debug message is issued when the TCP connection is being closed as a result of being idle.

26.47.10 TCP_REQUEST_RECEIVED_FAILED

```
An unexpected error occurred processing a request from %1, error: %2
```

This error message is issued when an unexpected error occurred while the server attempted to process a received request. The first argument specifies the address of the remote endpoint. The second argument describes the nature error.

26.47.11 TCP_REQUEST_RECEIVE_START

```
start receiving request from %1 with timeout %2
```

Logged at debug log level 50. This debug message is issued when the server starts receiving new request over the established connection. The first argument specifies the address of the remote endpoint. The second argument specifies request timeout in seconds.

26.47.12 TCP_SERVER_RESPONSE_SEND

```
sending TCP response to %1
```

Logged at debug log level 40. This debug message is issued when the server is starting to send a TCP response to a remote endpoint. The argument specifies an address of the remote endpoint.

26.48 TKEY

26.48.1 TKEY_EXCHANGE_ANSWER_CLASS

```
GSS-TKEY exchange received a response with answer class: %1.
```

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange received a response with specified answer class.

26.48.2 TKEY_EXCHANGE_FAILED_TO_VERIFY

```
GSS-TKEY exchange failed because the response failed to verify.
```

This error message indicated that GSS-TKEY exchange failed because the response failed to verify.

26.48.3 TKEY_EXCHANGE_FAIL_EMPTY_IN_TOKEN

```
GSS-TKEY exchange failed because input token is empty.
```

This error message indicated that GSS-TKEY exchange failed because input token is empty.

26.48.4 TKEY_EXCHANGE_FAIL_EMPTY_OUT_TOKEN

```
GSS-TKEY exchange failed because output token is empty.
```

This error message indicated that GSS-TKEY exchange failed because output token is empty.

26.48.5 TKEY_EXCHANGE_FAIL_EMPTY_RESPONSE

```
GSS-TKEY exchange failed because the response is empty.
```

This error message indicated that GSS-TKEY exchange failed because the response is empty.

26.48.6 TKEY_EXCHANGE_FAIL_IO_ERROR

```
GSS-TKEY exchange failed because of the IO error: %1.
```

This error message indicated that GSS-TKEY exchange failed because of an IO error. The argument details the IO error.

26.48.7 TKEY_EXCHANGE_FAIL_IO_STOPPED

```
GSS-TKEY exchange failed because the IO service was stopped.
```

This error message indicated that GSS-TKEY exchange failed because the IO service was stopped.

26.48.8 TKEY_EXCHANGE_FAIL_IO_TIMEOUT

GSS-TKEY exchange failed because of IO timeout.

This error message indicated that GSS-TKEY exchange failed because of IO timeout.

26.48.9 TKEY_EXCHANGE_FAIL_NOT_SIGNED

GSS-TKEY exchange failed because the response is not signed.

This error message indicated that GSS-TKEY exchange failed because the response is not signed.

26.48.10 TKEY_EXCHANGE_FAIL_NO_RDATA

GSS-TKEY exchange failed because the response contains no rdata.

This error message indicated that GSS-TKEY exchange failed because the response contains no rdata.

26.48.11 TKEY_EXCHANGE_FAIL_NO_RESPONSE_ANSWER

GSS-TKEY exchange failed because the response contains no answer.

This error message indicated that GSS-TKEY exchange failed because the response contains no answer.

26.48.12 TKEY_EXCHANGE_FAIL_NULL_RESPONSE

GSS-TKEY exchange failed because the response is null.

This error message indicated that GSS-TKEY exchange failed because the response is null.

26.48.13 TKEY_EXCHANGE_FAIL_RESPONSE_ERROR

GSS-TKEY exchange failed because the response contains an error: %1.

This error message indicated that GSS-TKEY exchange failed because the response contains an error. The argument details the response error.

26.48.14 TKEY_EXCHANGE_FAIL_TKEY_ERROR

GSS-TKEY exchange failed because the response contains TKEY error: %1.

This error message indicated that GSS-TKEY exchange failed because the response contains TKEY error. The argument details the TKEY error.

26.48.15 TKEY_EXCHANGE_FAIL_TO_INIT

GSS-TKEY exchange failed to initialize because of the error: %1.

This error message indicated that GSS-TKEY exchange failed in the initialization phase, for instance because the server principal does not exist. The argument details the error.

26.48.16 TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_ANSWER_COUNT

GSS-TKEY exchange failed because the response contains invalid number of RRs: %1.

This error message indicated that GSS-TKEY exchange failed because the response contains invalid number of RRs. The argument contains the wrong number of RRs.

26.48.17 TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_ANSWER_TYPE

GSS-TKEY exchange failed because the response contains wrong answer type: %1.

This error message indicated that GSS-TKEY exchange failed because the response contains wrong answer type. The argument contains the wrong answer type.

26.48.18 TKEY_EXCHANGE_FAIL_WRONG_RESPONSE_OPCODE

GSS-TKEY exchange failed because the response contains invalid opcode: %1.

This error message indicated that GSS-TKEY exchange failed because the response contains invalid opcode. The argument contains the wrong opcode.

26.48.19 TKEY_EXCHANGE_NOT_A_RESPONSE

GSS-TKEY exchange received a non response type.

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange received a non response type.

26.48.20 TKEY_EXCHANGE_OUT_TOKEN_NOT_EMPTY

GSS-TKEY exchange output token is not empty.

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange output token is not empty.

26.48.21 TKEY_EXCHANGE_RDATA_COUNT

GSS-TKEY exchange received a response with rdata count: %1.

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange received a response with specified rdata count.

26.48.22 TKEY_EXCHANGE_RECEIVE_MESSAGE

GSS-TKEY exchange receives a message of size: %1.

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange receives a message of specified size.

26.48.23 TKEY_EXCHANGE_RESPONSE_TTL

GSS-TKEY exchange received a response with TTL of: %1 seconds.

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange received a response with specified TTL.

26.48.24 TKEY_EXCHANGE_SEND_MESSAGE

```
GSS-TKEY exchange sends a message of size: %1.
```

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange sends a message of specified size.

26.48.25 TKEY_EXCHANGE_VALID

```
GSS-TKEY exchange retrieved a TKEY valid for: %1 seconds.
```

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange retrieved a TKEY valid for the specified time period expressed in seconds.

26.48.26 TKEY_EXCHANGE_VERIFIED

```
GSS-TKEY exchange verified.
```

Logged at debug log level 40. This debug message indicates that GSS-TKEY exchange is verified.

26.49 TLS

26.49.1 TLS_CONNECTION_HANDSHAKE_FAILED

```
TLS handshake with %1 failed with %2
```

This information message is issued when the TLS handshake failed at the server side. The client address and the error message are displayed.

26.49.2 TLS_CONNECTION_HANDSHAKE_START

```
start TLS handshake with %1 with timeout %2
```

Logged at debug log level 50. This debug message is issued when the server starts the TLS handshake with the remote endpoint. The first argument specifies the address of the remote endpoint. The second argument specifies request timeout in seconds.

26.49.3 TLS_REQUEST_RECEIVE_START

```
start receiving request from %1 with timeout %2
```

Logged at debug log level 50. This debug message is issued when the server starts receiving new request over the established connection. The first argument specifies the address of the remote endpoint. The second argument specifies request timeout in seconds.

26.49.4 TLS_SERVER_RESPONSE_SEND

```
sending TLS response to %1
```

Logged at debug log level 40. This debug message is issued when the server is starting to send a TLS response to a remote endpoint. The argument specifies an address of the remote endpoint.

26.50 USER

26.50.1 USER_CHK_HOOK_LOAD_ERROR

```
DHCP UserCheckHook could not be loaded: %1
```

This is an error message issued when the DHCP UserCheckHook could not be loaded. The exact cause should be explained in the log message. User subnet selection will revert to default processing.

26.50.2 USER_CHK_HOOK_UNLOAD_ERROR

```
DHCP UserCheckHook an error occurred unloading the library: %1
```

This is an error message issued when an error occurs while unloading the UserCheckHook library. This is unlikely to occur and normal operations of the library will likely resume when it is next loaded.

26.50.3 USER_CHK_SUBNET4_SELECT_ERROR

```
DHCP UserCheckHook an unexpected error occurred in subnet4_select callout: %1
```

This is an error message issued when the DHCP UserCheckHook subnet4_select hook encounters an unexpected error. The message should contain a more detailed explanation.

26.50.4 USER_CHK_SUBNET4_SELECT_REGISTRY_NULL

```
DHCP UserCheckHook UserRegistry has not been created.
```

This is an error message issued when the DHCP UserCheckHook subnet4_select hook has been invoked but the UserRegistry has not been created. This is a programmatic error and should not occur.

26.50.5 USER_CHK_SUBNET6_SELECT_ERROR

```
DHCP UserCheckHook an unexpected error occurred in subnet6_select callout: %1
```

This is an error message issued when the DHCP UserCheckHook subnet6_select hook encounters an unexpected error. The message should contain a more detailed explanation.

26.50.6 USER_CHK_SUBNET6_SELECT_REGISTRY_NULL

```
DHCP UserCheckHook UserRegistry has not been created.
```

This is an error message issued when the DHCP UserCheckHook subnet6_select hook has been invoked but the UserRegistry has not been created. This is a programmatic error and should not occur.

26.51 Kea Debug Messages By Log Level

26.51.1 Messages printed on debuglevel 0

- CTRL_AGENT_RUN_EXIT
- DCTL_INIT_PROCESS
- DCTL_RUN_PROCESS

- DCTL_SHUTDOWN
- DCTL_SHUTDOWN_SIGNAL_RECV
- DCTL_STANDALONE
- DHCP4_OPEN_SOCKET
- DHCP4_START_INFO
- DHCP6_OPEN_SOCKET
- DHCP6_START_INFO
- DHCP_DDNS_CLEARED_FOR_SHUTDOWN
- DHCP_DDNS_QUEUE_MGR_STARTED
- DHCP_DDNS_QUEUE_MGR_STOPPING
- DHCP_DDNS_RUN_EXIT
- DHCP_DDNS_SHUTDOWN_COMMAND
- NETCONF_RUN_EXIT

26.51.2 Messages printed on debuglevel 10

- COMMAND_DEREGISTERED
- COMMAND_EXTENDED_REGISTERED
- COMMAND_HTTP_LISTENER_COMMAND_REJECTED
- COMMAND_HTTP_LISTENER_STARTED
- COMMAND_HTTP_LISTENER_STOPPED
- COMMAND_HTTP_LISTENER_STOPPING
- COMMAND_REGISTERED
- COMMAND_SOCKET_CONNECTION_CLOSED
- COMMAND_SOCKET_CONNECTION_OPENED
- COMMAND_SOCKET_READ
- COMMAND_SOCKET_WRITE
- CTRL_AGENT_COMMAND_FORWARD_BEGIN
- CTRL_AGENT_COMMAND_FORWARD_FAILED
- DCTL_CONFIG_START
- DHCP4_CONFIG_RECEIVED
- DHCP4_CONFIG_START
- DHCP6_CONFIG_RECEIVED
- DHCP6_CONFIG_START

26.51.3 Messages printed on debuglevel 15

- DHCP4_HOOK_BUFFER_RCVD_DROP
- DHCP4_HOOK_DDNS_UPDATE
- DHCP4_HOOK_DECLINE_SKIP
- DHCP4_HOOK_LEASE4_RELEASE_SKIP
- DHCP4_HOOK_PACKET_SEND_DROP
- DHCP4_HOOK_SUBNET4_SELECT_406_PARKING_LOT_FULL
- DHCP4_HOOK_SUBNET4_SELECT_PARKING_LOT_FULL
- DHCP4_HOOK_SUBNET6_SELECT_PARKING_LOT_FULL
- DHCP4_PACKET_DROP_0001
- DHCP4_PACKET_DROP_0002
- DHCP4_PACKET_DROP_0003
- DHCP4_PACKET_DROP_0004
- DHCP4_PACKET_DROP_0005
- DHCP4_PACKET_DROP_0006
- DHCP4_PACKET_DROP_0007
- DHCP4_PACKET_DROP_0008
- DHCP4_PACKET_DROP_0009
- DHCP4_PACKET_DROP_0010
- DHCP4_PACKET_DROP_0011
- DHCP4_PACKET_DROP_0012
- DHCP4_PACKET_DROP_0013
- DHCP4_PACKET_DROP_0014
- DHCP6_HOOK_BUFFER_RCVD_DROP
- DHCP6_HOOK_DDNS_UPDATE
- DHCP6_HOOK_DECLINE_DROP
- DHCP6_HOOK_LEASES6_COMMITTED_DROP
- DHCP6_HOOK_LEASES6_PARKING_LOT_FULL
- DHCP6_HOOK_PACKET_SEND_DROP
- DHCP6_PACKET_DROP_DHCP_DISABLED
- DHCP6_PACKET_DROP_DROP_CLASS
- DHCP6_PACKET_DROP_DROP_CLASS2
- DHCP6_PACKET_DROP_DROP_CLASS_EARLY
- DHCP6_PACKET_DROP_DUPLICATE
- DHCP6_PACKET_DROP_PARSE_FAIL
- DHCP6_PACKET_DROP_SERVERID_MISMATCH

- DHCP6_PACKET_DROP_UNICAST

26.51.4 Messages printed on debuglevel 20

- LEASE_CMDS_ADD4
- LEASE_CMDS_ADD6
- LEASE_CMDS_BULK_APPLY6
- LEASE_CMDS_DEL4
- LEASE_CMDS_DEL6
- LEASE_CMDS_UPDATE4
- LEASE_CMDS_UPDATE6

26.51.5 Messages printed on debuglevel 40

- ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS
- ALLOC_ENGINE_IGNOREING_UNSUITABLE_GLOBAL_ADDRESS6
- ALLOC_ENGINE_LEASE_RECLAIMED
- ALLOC_ENGINE_V4_DISCOVER_HR
- ALLOC_ENGINE_V4_LEASES_RECLAMATION_COMPLETE
- ALLOC_ENGINE_V4_LEASES_RECLAMATION_START
- ALLOC_ENGINE_V4_LEASES_RECLAMATION_TIMEOUT
- ALLOC_ENGINE_V4_LEASE_RECLAIM
- ALLOC_ENGINE_V4_NO_MORE_EXPIRED_LEASES
- ALLOC_ENGINE_V4_OFFER_EXISTING_LEASE
- ALLOC_ENGINE_V4_OFFER_NEW_LEASE
- ALLOC_ENGINE_V4_OFFER_REQUESTED_LEASE
- ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE
- ALLOC_ENGINE_V4_RECLAIMED_LEASES_DELETE_COMPLETE
- ALLOC_ENGINE_V4_REQUEST_ADDRESS_RESERVED
- ALLOC_ENGINE_V4_REQUEST_ALLOC_REQUESTED
- ALLOC_ENGINE_V4_REQUEST_EXTEND_LEASE
- ALLOC_ENGINE_V4_REQUEST_INVALID
- ALLOC_ENGINE_V4_REQUEST_IN_USE
- ALLOC_ENGINE_V4_REQUEST_OUT_OF_POOL
- ALLOC_ENGINE_V4_REQUEST_PICK_ADDRESS
- ALLOC_ENGINE_V4_REQUEST_REMOVE_LEASE
- ALLOC_ENGINE_V4_REQUEST_USE_HR
- ALLOC_ENGINE_V6_ALLOC_HR_LEASE_EXISTS
- ALLOC_ENGINE_V6_ALLOC_LEASES_HR

- ALLOC_ENGINE_V6_ALLOC_LEASES_NO_HR
- ALLOC_ENGINE_V6_ALLOC_NO_LEASES_HR
- ALLOC_ENGINE_V6_ALLOC_NO_V6_HR
- ALLOC_ENGINE_V6_ALLOC_UNRESERVED
- ALLOC_ENGINE_V6_CALCULATED_PREFERRED_LIFETIME
- ALLOC_ENGINE_V6_EXPIRED_HINT_RESERVED
- ALLOC_ENGINE_V6_EXTEND_ALLOC_UNRESERVED
- ALLOC_ENGINE_V6_HINT_RESERVED
- ALLOC_ENGINE_V6_LEASES_RECLAMATION_COMPLETE
- ALLOC_ENGINE_V6_LEASES_RECLAMATION_START
- ALLOC_ENGINE_V6_LEASES_RECLAMATION_TIMEOUT
- ALLOC_ENGINE_V6_LEASE_RECLAIM
- ALLOC_ENGINE_V6_NO_MORE_EXPIRED_LEASES
- ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE
- ALLOC_ENGINE_V6_RECLAIMED_LEASES_DELETE_COMPLETE
- ALLOC_ENGINE_V6_RENEW_HR
- ALLOC_ENGINE_V6_RENEW_REMOVE_RESERVED
- ASIODNS_FETCH_STOPPED
- BOOTP_BOOTP_QUERY
- BOOTP_PACKET_OPTIONS_SKIPPED
- BOOTP_PACKET_PACK
- BOOTP_PACKET_UNPACK_FAILED
- BULK_LEASE_QUERY_AT_MAX_CONCURRENT_QUERIES
- BULK_LEASE_QUERY_DEQUEUED
- BULK_LEASE_QUERY_QUERY_RECEIVED
- BULK_LEASE_QUERY_RESPONSE_SENT
- DDNS_TUNING4_CALCULATED_HOSTNAME
- DDNS_TUNING4_SKIPPING_DDNS
- DDNS_TUNING6_CALCULATED_HOSTNAME
- DDNS_TUNING6_SKIPPING_DDNS
- DDNS_TUNING_SUBNET_EXPRESSION_PARSE
- DDNS_TUNING_SUBNET_EXPR_CACHED
- DHCP4_ADDITIONAL_CLASS_NO_TEST
- DHCP4_ADDITIONAL_CLASS_UNDEFINED
- DHCP4_BUFFER_RECEIVED
- DHCP4_CLASSES_ASSIGNED

- DHCP4_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION
- DHCP4_CLASS_ASSIGNED
- DHCP4_CLASS_UNCONFIGURED
- DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_DROP
- DHCP4_DHCP4O6_HOOK_SUBNET4_SELECT_SKIP
- DHCP4_DHCP4O6_PACKET_RECEIVED
- DHCP4_DHCP4O6_PACKET_SEND
- DHCP4_FLEX_ID
- DHCP4_HOOK_BUFFER_SEND_SKIP
- DHCP4_HOOK_PACKET_RCVD_SKIP
- DHCP4_HOOK_PACKET_SEND_SKIP
- DHCP4_HOOK_SUBNET4_SELECT_DROP
- DHCP4_HOOK_SUBNET4_SELECT_PARK
- DHCP4_HOOK_SUBNET4_SELECT_SKIP
- DHCP4_LEASE_QUERY_PACKET_UNPACK_FAILED
- DHCP4_LEASE_QUERY_PROCESS_FAILED
- DHCP4_LEASE_QUERY_RECEIVED
- DHCP4_LEASE_QUERY_RESPONSE_SENT
- DHCP4_PACKET_QUEUE_FULL
- DHCP4_RECLAIM_EXPIRED_LEASES_SKIPPED
- DHCP4_SHUTDOWN
- DHCP4_SHUTDOWN_REQUEST
- DHCP6_ADDITIONAL_CLASS_NO_TEST
- DHCP6_ADDITIONAL_CLASS_UNDEFINED
- DHCP6_BUFFER_RECEIVED
- DHCP6_CLASSES_ASSIGNED
- DHCP6_CLASSES_ASSIGNED_AFTER_SUBNET_SELECTION
- DHCP6_CLASS_ASSIGNED
- DHCP6_CLASS_UNCONFIGURED
- DHCP6_DHCP4O6_PACKET_RECEIVED
- DHCP6_FLEX_ID
- DHCP6_HOOK_ADDR6_REGISTER_DROP
- DHCP6_HOOK_ADDR6_REGISTER_SKIP
- DHCP6_HOOK_BUFFER_SEND_SKIP
- DHCP6_HOOK_LEASE6_RELEASE_NA_SKIP
- DHCP6_HOOK_LEASE6_RELEASE_PD_SKIP

- DHCP6_HOOK_LEASES6_COMMITTED_PARK
- DHCP6_HOOK_PACKET_RCVD_SKIP
- DHCP6_HOOK_PACKET_SEND_SKIP
- DHCP6_HOOK_SUBNET6_SELECT_DROP
- DHCP6_HOOK_SUBNET6_SELECT_PARK
- DHCP6_HOOK_SUBNET6_SELECT_SKIP
- DHCP6_LEASE_QUERY_PACKET_UNPACK_FAILED
- DHCP6_LEASE_QUERY_PREFIX_LENGTH_LIST
- DHCP6_LEASE_QUERY_PROCESS_FAILED
- DHCP6_LEASE_QUERY_RECEIVED
- DHCP6_LEASE_QUERY_REPLY_SENT
- DHCP6_PACKET_PROCESS_FAIL
- DHCP6_PACKET_QUEUE_FULL
- DHCP6_RECLAIM_EXPIRED_LEASES_SKIPPED
- DHCP6_REQUIRED_OPTIONS_CHECK_FAIL
- DHCP6_SHUTDOWN
- DHCP6_SHUTDOWN_REQUEST
- DHCP6_UNKNOWN_MSG_RECEIVED
- DHCPSRV_CFGMGR_ADD_SUBNET4
- DHCPSRV_CFGMGR_ADD_SUBNET6
- DHCPSRV_CFGMGR_ALL_IFACES_ACTIVE
- DHCPSRV_CFGMGR_CFG_DHCP_DDNS
- DHCPSRV_CFGMGR_DEL_SUBNET4
- DHCPSRV_CFGMGR_DEL_SUBNET6
- DHCPSRV_CFGMGR_SUBNET4
- DHCPSRV_CFGMGR_SUBNET4_ADDR
- DHCPSRV_CFGMGR_SUBNET4_IFACE
- DHCPSRV_CFGMGR_SUBNET4_RELAY
- DHCPSRV_CFGMGR_SUBNET6
- DHCPSRV_CFGMGR_SUBNET6_IFACE
- DHCPSRV_CFGMGR_SUBNET6_IFACE_ID
- DHCPSRV_CFGMGR_SUBNET6_RELAY
- DHCPSRV_CFGMGR_UPDATE_SUBNET4
- DHCPSRV_CFGMGR_UPDATE_SUBNET6
- DHCPSRV_CLOSE_DB
- DHCPSRV_FORENSIC_BACKEND_DEREGISTER

- DHCPSRV_FORENSIC_BACKEND_REGISTER
- DHCPSRV_HOOK_LEASE4_RECOVER_SKIP
- DHCPSRV_HOOK_LEASE4_RENEW_SKIP
- DHCPSRV_HOOK_LEASE4_SELECT_SKIP
- DHCPSRV_HOOK_LEASE6_EXTEND_SKIP
- DHCPSRV_HOOK_LEASE6_RECOVER_SKIP
- DHCPSRV_HOOK_LEASE6_SELECT_SKIP
- DHCPSRV_LEASE4_EXTENDED_INFO_UPGRADED
- DHCPSRV_LEASE6_EXTENDED_INFO_UPGRADED
- DHCPSRV_LEASE_MGR_BACKEND_DEREGISTER
- DHCPSRV_LEASE_MGR_BACKEND_REGISTER
- DHCPSRV_MEMFILE_BEGIN_BUILD_EXTENDED_INFO_TABLES6
- DHCPSRV_MEMFILE_BEGIN_EXTRACT_EXTENDED_INFO4
- DHCPSRV_MEMFILE_EXTRACT_EXTENDED_INFO4
- DHCPSRV_MEMFILE_EXTRACT_EXTENDED_INFO4_ERROR
- DHCPSRV_MEMFILE_LFC_UNREGISTER_TIMER_FAILED
- DHCPSRV_SUBNET4O6_SELECT_FAILED
- DHCPSRV_SUBNET4_SELECT_BY_ADDRESS_NO_MATCH
- DHCPSRV_SUBNET4_SELECT_BY_INTERFACE_NO_MATCH
- DHCPSRV_SUBNET4_SELECT_BY_RELAY_ADDRESS_NO_MATCH
- DHCPSRV_SUBNET4_SELECT_NO_RAI_OPTIONS
- DHCPSRV_SUBNET4_SELECT_NO_RELAY_ADDRESS
- DHCPSRV_SUBNET4_SELECT_NO_USABLE_ADDRESS
- DHCPSRV_SUBNET6_SELECT_BY_ADDRESS_NO_MATCH
- DHCPSRV_SUBNET6_SELECT_BY_INTERFACE_ID_NO_MATCH
- DHCPSRV_SUBNET6_SELECT_BY_INTERFACE_NO_MATCH
- DHCPSRV_TIMERMGR_REGISTER_TIMER
- DHCPSRV_TIMERMGR_START_TIMER
- DHCPSRV_TIMERMGR_STOP_TIMER
- DHCPSRV_TIMERMGR_UNREGISTER_ALL_TIMERS
- DHCPSRV_TIMERMGR_UNREGISTER_TIMER
- DHCP_DDNS_CONFIGURE
- DHCP_DDNS_NCR_UDP_RECV_CANCELED
- DHCP_DDNS_QUEUE_MGR_RECONFIGURING
- DHCP_DDNS_QUEUE_MGR_STOPPED
- FLEX_ID_EXPRESSION_EVALUATED

- FLEX_ID_EXPRESSION_HEX
- FLEX_ID_IGNORE_IAID_APPLIED_ON_NA
- FLEX_ID_IGNORE_IAID_APPLIED_ON_PD
- FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_NA
- FLEX_ID_IGNORE_IAID_NOT_APPLIED_ON_PD
- FLEX_ID_RESTORE_CLIENT_ID
- FLEX_ID_RESTORE_DUID
- FLEX_ID_USED_AS_CLIENT_ID
- FLEX_ID_USED_AS_DUID
- FLEX_OPTION_PROCESS_ADD
- FLEX_OPTION_PROCESS_CLIENT_CLASS
- FLEX_OPTION_PROCESS_REMOVE
- FLEX_OPTION_PROCESS_SUB_ADD
- FLEX_OPTION_PROCESS_SUB_CLIENT_CLASS
- FLEX_OPTION_PROCESS_SUB_REMOVE
- FLEX_OPTION_PROCESS_SUB_SUPERSEDE
- FLEX_OPTION_PROCESS_SUPERSEDE
- FLEX_OPTION_PROCESS_VENDOR_ID_MISMATCH
- GSS_TSIG_MANAGER_STARTED
- GSS_TSIG_MANAGER_STOPPED
- GSS_TSIG_NEW_KEY
- GSS_TSIG_NEW_KEY_SETUP_SUCCEED
- GSS_TSIG_OLD_KEY_REMOVED
- GSS_TSIG_VERIFIED
- HA_BUFFER4_RECEIVE_NOT_FOR_US
- HA_BUFFER4_RECEIVE_PACKET_OPTIONS_SKIPPED
- HA_BUFFER4_RECEIVE_UNPACK_FAILED
- HA_BUFFER6_RECEIVE_NOT_FOR_US
- HA_BUFFER6_RECEIVE_PACKET_OPTIONS_SKIPPED
- HA_BUFFER6_RECEIVE_UNPACK_FAILED
- HA_LEASE4_EXPIRE_RECLAMATION_SKIP
- HA_LEASE6_EXPIRE_RECLAMATION_SKIP
- HA_LEASES4_COMMITTED_NOTHING_TO_UPDATE
- HA_LEASES6_COMMITTED_NOTHING_TO_UPDATE
- HA_LEASE_SYNC_STALE_LEASE4_SKIP
- HA_LEASE_SYNC_STALE_LEASE6_SKIP

- HA_LOAD_BALANCING_DUID_MISSING
- HA_LOAD_BALANCING_IDENTIFIER_MISSING
- HA_LOAD_BALANCING_LEASE_DUID_MISSING
- HA_LOAD_BALANCING_LEASE_IDENTIFIER_MISSING
- HA_SUBNET4_SELECT_NOT_FOR_US
- HA_SUBNET4_SELECT_NO_SUBNET_SELECTED
- HA_SUBNET6_SELECT_NOT_FOR_US
- HA_SUBNET6_SELECT_NO_SUBNET_SELECTED
- HOOKS_LIBRARY_LOADING
- HOOKS_LIBRARY_UNLOADING
- HOOKS_LOAD_SUCCESS
- HOOKS_NO_LOAD
- HOOKS_NO_UNLOAD
- HOOKS_UNLOAD_SUCCESS
- HOSTS_BACKEND_DEREGISTER
- HOSTS_BACKEND_REGISTER
- HOSTS_CFG_ADD_HOST
- HOSTS_CFG_CLOSE_HOST_DATA_SOURCE
- HOSTS_CFG_DEL
- HOSTS_CFG_DEL4
- HOSTS_CFG_DEL6
- HOSTS_CFG_DEL_ALL_SUBNET4
- HOSTS_CFG_DEL_ALL_SUBNET6
- HOSTS_CFG_GET_ALL
- HOSTS_CFG_GET_ALL_ADDRESS4
- HOSTS_CFG_GET_ALL_ADDRESS6
- HOSTS_CFG_GET_ALL_HOSTNAME
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6
- HOSTS_CFG_GET_ALL_IDENTIFIER
- HOSTS_CFG_GET_ALL_SUBNET_ID4
- HOSTS_CFG_GET_ALL_SUBNET_ID6
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6
- HOSTS_CFG_GET_ONE_PREFIX
- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4

- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6
- HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER
- HOSTS_CFG_UPDATE_ADD
- HOSTS_CFG_UPDATE_DEL4
- HOSTS_CFG_UPDATE_DEL6
- HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_ADDRESS4
- HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER
- HOSTS_MGR_ALTERNATE_GET6_PREFIX
- HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_ADDRESS6
- HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER
- HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS4
- HOSTS_MGR_ALTERNATE_GET_ALL_SUBNET_ID_ADDRESS6
- HOST_CACHE_GET_ONE_PREFIX
- HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4
- HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6
- HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER
- HTTP_BAD_CLIENT_REQUEST_RECEIVED
- HTTP_BAD_SERVER_RESPONSE_RECEIVED
- HTTP_CLIENT_MT_STARTED
- HTTP_CLIENT_REQUEST_RECEIVED
- HTTP_CONNECTION_SHUTDOWN
- HTTP_CONNECTION_STOP
- HTTP_SERVER_RESPONSE_RECEIVED
- HTTP_SERVER_RESPONSE_SEND
- KEY_LOOKUP_DISABLED
- KEY_LOOKUP_FOUND
- KEY_LOOKUP_NONE
- LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_CLIENT_CLASS
- LIMITS_CONFIGURED_ADDRESS_LIMIT_BY_SUBNET
- LIMITS_CONFIGURED_PREFIX_LIMIT_BY_CLIENT_CLASS
- LIMITS_CONFIGURED_PREFIX_LIMIT_BY_SUBNET
- LIMITS_CONFIGURED_RATE_LIMIT_BY_CLIENT_CLASS
- LIMITS_CONFIGURED_RATE_LIMIT_BY_SUBNET
- LIMITS_LEASE_LIMIT_EXCEEDED
- LIMITS_LEASE_WITHIN_LIMITS
- LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_DROPPED

- LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_DROPPED
- MT_TCP_LISTENER_MGR_STARTED
- MT_TCP_LISTENER_MGR_STOPPED
- MT_TCP_LISTENER_MGR_STOPPING
- MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4
- MYSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6
- MYSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6
- MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4
- MYSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6
- MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS4
- MYSQL_CB_CREATE_UPDATE_CLIENT_CLASS6
- MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4
- MYSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6
- MYSQL_CB_CREATE_UPDATE_OPTION4
- MYSQL_CB_CREATE_UPDATE_OPTION6
- MYSQL_CB_CREATE_UPDATE_OPTION_DEF4
- MYSQL_CB_CREATE_UPDATE_OPTION_DEF6
- MYSQL_CB_CREATE_UPDATE_SERVER4
- MYSQL_CB_CREATE_UPDATE_SERVER6
- MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK4
- MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK6
- MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4
- MYSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6
- MYSQL_CB_CREATE_UPDATE_SUBNET4
- MYSQL_CB_CREATE_UPDATE_SUBNET6
- MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4
- MYSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT
- MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6
- MYSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT
- MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4
- MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT
- MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6
- MYSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT
- MYSQL_CB_DELETE_ALL_OPTION_DEFS4
- MYSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT
- MYSQL_CB_DELETE_ALL_OPTION_DEFS6

- MYSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT
- MYSQL_CB_DELETE_ALL_SERVERS4
- MYSQL_CB_DELETE_ALL_SERVERS4_RESULT
- MYSQL_CB_DELETE_ALL_SERVERS6
- MYSQL_CB_DELETE_ALL_SERVERS6_RESULT
- MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4
- MYSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT
- MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6
- MYSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT
- MYSQL_CB_DELETE_ALL_SUBNETS4
- MYSQL_CB_DELETE_ALL_SUBNETS4_RESULT
- MYSQL_CB_DELETE_ALL_SUBNETS6
- MYSQL_CB_DELETE_ALL_SUBNETS6_RESULT
- MYSQL_CB_DELETE_BY_POOL_OPTION4
- MYSQL_CB_DELETE_BY_POOL_OPTION4_RESULT
- MYSQL_CB_DELETE_BY_POOL_OPTION6
- MYSQL_CB_DELETE_BY_POOL_OPTION6_RESULT
- MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6
- MYSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT
- MYSQL_CB_DELETE_BY_PREFIX_SUBNET4
- MYSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT
- MYSQL_CB_DELETE_BY_PREFIX_SUBNET6
- MYSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT
- MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4
- MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT
- MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6
- MYSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT
- MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4
- MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT
- MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6
- MYSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT
- MYSQL_CB_DELETE_CLIENT_CLASS4
- MYSQL_CB_DELETE_CLIENT_CLASS4_RESULT
- MYSQL_CB_DELETE_CLIENT_CLASS6
- MYSQL_CB_DELETE_CLIENT_CLASS6_RESULT
- MYSQL_CB_DELETE_GLOBAL_PARAMETER4

- MYSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT
- MYSQL_CB_DELETE_GLOBAL_PARAMETER6
- MYSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT
- MYSQL_CB_DELETE_OPTION4
- MYSQL_CB_DELETE_OPTION4_RESULT
- MYSQL_CB_DELETE_OPTION6
- MYSQL_CB_DELETE_OPTION6_RESULT
- MYSQL_CB_DELETE_OPTION_DEF4
- MYSQL_CB_DELETE_OPTION_DEF4_RESULT
- MYSQL_CB_DELETE_OPTION_DEF6
- MYSQL_CB_DELETE_OPTION_DEF6_RESULT
- MYSQL_CB_DELETE_SERVER4
- MYSQL_CB_DELETE_SERVER4_RESULT
- MYSQL_CB_DELETE_SERVER6
- MYSQL_CB_DELETE_SERVER6_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK4
- MYSQL_CB_DELETE_SHARED_NETWORK4_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK6
- MYSQL_CB_DELETE_SHARED_NETWORK6_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4
- MYSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6
- MYSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4
- MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT
- MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6
- MYSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT
- MYSQL_CB_GET_ALL_CLIENT_CLASSES4
- MYSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT
- MYSQL_CB_GET_ALL_CLIENT_CLASSES6
- MYSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT
- MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4
- MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT
- MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6
- MYSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT
- MYSQL_CB_GET_ALL_OPTIONS4

- MYSQL_CB_GET_ALL_OPTIONS4_RESULT
- MYSQL_CB_GET_ALL_OPTIONS6
- MYSQL_CB_GET_ALL_OPTIONS6_RESULT
- MYSQL_CB_GET_ALL_OPTION_DEFS4
- MYSQL_CB_GET_ALL_OPTION_DEFS4_RESULT
- MYSQL_CB_GET_ALL_OPTION_DEFS6
- MYSQL_CB_GET_ALL_OPTION_DEFS6_RESULT
- MYSQL_CB_GET_ALL_SERVERS4
- MYSQL_CB_GET_ALL_SERVERS4_RESULT
- MYSQL_CB_GET_ALL_SERVERS6
- MYSQL_CB_GET_ALL_SERVERS6_RESULT
- MYSQL_CB_GET_ALL_SHARED_NETWORKS4
- MYSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT
- MYSQL_CB_GET_ALL_SHARED_NETWORKS6
- MYSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT
- MYSQL_CB_GET_ALL_SUBNETS4
- MYSQL_CB_GET_ALL_SUBNETS4_RESULT
- MYSQL_CB_GET_ALL_SUBNETS6
- MYSQL_CB_GET_ALL_SUBNETS6_RESULT
- MYSQL_CB_GET_CLIENT_CLASS4
- MYSQL_CB_GET_CLIENT_CLASS6
- MYSQL_CB_GET_GLOBAL_PARAMETER4
- MYSQL_CB_GET_GLOBAL_PARAMETER6
- MYSQL_CB_GET_HOST4
- MYSQL_CB_GET_HOST6
- MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4
- MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT
- MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6
- MYSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT
- MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4
- MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT
- MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6
- MYSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT
- MYSQL_CB_GET_MODIFIED_OPTIONS4
- MYSQL_CB_GET_MODIFIED_OPTIONS4_RESULT
- MYSQL_CB_GET_MODIFIED_OPTIONS6

- MYSQL_CB_GET_MODIFIED_OPTIONS6_RESULT
- MYSQL_CB_GET_MODIFIED_OPTION_DEFS4
- MYSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT
- MYSQL_CB_GET_MODIFIED_OPTION_DEFS6
- MYSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT
- MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4
- MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT
- MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6
- MYSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT
- MYSQL_CB_GET_MODIFIED_SUBNETS4
- MYSQL_CB_GET_MODIFIED_SUBNETS4_RESULT
- MYSQL_CB_GET_MODIFIED_SUBNETS6
- MYSQL_CB_GET_MODIFIED_SUBNETS6_RESULT
- MYSQL_CB_GET_OPTION4
- MYSQL_CB_GET_OPTION6
- MYSQL_CB_GET_OPTION_DEF4
- MYSQL_CB_GET_OPTION_DEF6
- MYSQL_CB_GET_PORT4
- MYSQL_CB_GET_PORT6
- MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4
- MYSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT
- MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6
- MYSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT
- MYSQL_CB_GET_SERVER4
- MYSQL_CB_GET_SERVER6
- MYSQL_CB_GET_SHARED_NETWORK4
- MYSQL_CB_GET_SHARED_NETWORK6
- MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4
- MYSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT
- MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6
- MYSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT
- MYSQL_CB_GET_SUBNET4_BY_PREFIX
- MYSQL_CB_GET_SUBNET4_BY_SUBNET_ID
- MYSQL_CB_GET_SUBNET6_BY_PREFIX
- MYSQL_CB_GET_SUBNET6_BY_SUBNET_ID
- MYSQL_CB_GET_TYPE4

- MYSQL_CB_GET_TYPE6
- MYSQL_CB_REGISTER_BACKEND_TYPE4
- MYSQL_CB_REGISTER_BACKEND_TYPE6
- MYSQL_CB_TLS_CIPHER
- MYSQL_CB_UNREGISTER_BACKEND_TYPE4
- MYSQL_CB_UNREGISTER_BACKEND_TYPE6
- MYSQL_HB_TLS_CIPHER
- MYSQL_LB_TLS_CIPHER
- MYSQL_LB_UPGRADE_EXTENDED_INFO4
- MYSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR
- MYSQL_LB_UPGRADE_EXTENDED_INFO6
- MYSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR
- PERFMON_DHCP4_SOCKET_RECEIVED_TIME_SUPPORT
- PERFMON_DHCP6_SOCKET_RECEIVED_TIME_SUPPORT
- PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION4
- PGSQL_CB_CREATE_UPDATE_BY_POOL_OPTION6
- PGSQL_CB_CREATE_UPDATE_BY_PREFIX_OPTION6
- PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION4
- PGSQL_CB_CREATE_UPDATE_BY_SUBNET_ID_OPTION6
- PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS4
- PGSQL_CB_CREATE_UPDATE_CLIENT_CLASS6
- PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER4
- PGSQL_CB_CREATE_UPDATE_GLOBAL_PARAMETER6
- PGSQL_CB_CREATE_UPDATE_OPTION4
- PGSQL_CB_CREATE_UPDATE_OPTION6
- PGSQL_CB_CREATE_UPDATE_OPTION_DEF4
- PGSQL_CB_CREATE_UPDATE_OPTION_DEF6
- PGSQL_CB_CREATE_UPDATE_SERVER4
- PGSQL_CB_CREATE_UPDATE_SERVER6
- PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK4
- PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK6
- PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION4
- PGSQL_CB_CREATE_UPDATE_SHARED_NETWORK_OPTION6
- PGSQL_CB_CREATE_UPDATE_SUBNET4
- PGSQL_CB_CREATE_UPDATE_SUBNET6
- PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4

- PGSQL_CB_DELETE_ALL_CLIENT_CLASSES4_RESULT
- PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6
- PGSQL_CB_DELETE_ALL_CLIENT_CLASSES6_RESULT
- PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4
- PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS4_RESULT
- PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6
- PGSQL_CB_DELETE_ALL_GLOBAL_PARAMETERS6_RESULT
- PGSQL_CB_DELETE_ALL_OPTION_DEFS4
- PGSQL_CB_DELETE_ALL_OPTION_DEFS4_RESULT
- PGSQL_CB_DELETE_ALL_OPTION_DEFS6
- PGSQL_CB_DELETE_ALL_OPTION_DEFS6_RESULT
- PGSQL_CB_DELETE_ALL_SERVERS4
- PGSQL_CB_DELETE_ALL_SERVERS4_RESULT
- PGSQL_CB_DELETE_ALL_SERVERS6
- PGSQL_CB_DELETE_ALL_SERVERS6_RESULT
- PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4
- PGSQL_CB_DELETE_ALL_SHARED_NETWORKS4_RESULT
- PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6
- PGSQL_CB_DELETE_ALL_SHARED_NETWORKS6_RESULT
- PGSQL_CB_DELETE_ALL_SUBNETS4
- PGSQL_CB_DELETE_ALL_SUBNETS4_RESULT
- PGSQL_CB_DELETE_ALL_SUBNETS6
- PGSQL_CB_DELETE_ALL_SUBNETS6_RESULT
- PGSQL_CB_DELETE_BY_POOL_OPTION4
- PGSQL_CB_DELETE_BY_POOL_OPTION4_RESULT
- PGSQL_CB_DELETE_BY_POOL_OPTION6
- PGSQL_CB_DELETE_BY_POOL_OPTION6_RESULT
- PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6
- PGSQL_CB_DELETE_BY_POOL_PREFIX_OPTION6_RESULT
- PGSQL_CB_DELETE_BY_PREFIX_SUBNET4
- PGSQL_CB_DELETE_BY_PREFIX_SUBNET4_RESULT
- PGSQL_CB_DELETE_BY_PREFIX_SUBNET6
- PGSQL_CB_DELETE_BY_PREFIX_SUBNET6_RESULT
- PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4
- PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION4_RESULT
- PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6

- PGSQL_CB_DELETE_BY_SUBNET_ID_OPTION6_RESULT
- PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4
- PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET4_RESULT
- PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6
- PGSQL_CB_DELETE_BY_SUBNET_ID_SUBNET6_RESULT
- PGSQL_CB_DELETE_CLIENT_CLASS4
- PGSQL_CB_DELETE_CLIENT_CLASS4_RESULT
- PGSQL_CB_DELETE_CLIENT_CLASS6
- PGSQL_CB_DELETE_CLIENT_CLASS6_RESULT
- PGSQL_CB_DELETE_GLOBAL_PARAMETER4
- PGSQL_CB_DELETE_GLOBAL_PARAMETER4_RESULT
- PGSQL_CB_DELETE_GLOBAL_PARAMETER6
- PGSQL_CB_DELETE_GLOBAL_PARAMETER6_RESULT
- PGSQL_CB_DELETE_OPTION4
- PGSQL_CB_DELETE_OPTION4_RESULT
- PGSQL_CB_DELETE_OPTION6
- PGSQL_CB_DELETE_OPTION6_RESULT
- PGSQL_CB_DELETE_OPTION_DEF4
- PGSQL_CB_DELETE_OPTION_DEF4_RESULT
- PGSQL_CB_DELETE_OPTION_DEF6
- PGSQL_CB_DELETE_OPTION_DEF6_RESULT
- PGSQL_CB_DELETE_SERVER4
- PGSQL_CB_DELETE_SERVER4_RESULT
- PGSQL_CB_DELETE_SERVER6
- PGSQL_CB_DELETE_SERVER6_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK4
- PGSQL_CB_DELETE_SHARED_NETWORK4_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK6
- PGSQL_CB_DELETE_SHARED_NETWORK6_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4
- PGSQL_CB_DELETE_SHARED_NETWORK_OPTION4_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6
- PGSQL_CB_DELETE_SHARED_NETWORK_OPTION6_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4
- PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS4_RESULT
- PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6

- PGSQL_CB_DELETE_SHARED_NETWORK_SUBNETS6_RESULT
- PGSQL_CB_GET_ALL_CLIENT_CLASSES4
- PGSQL_CB_GET_ALL_CLIENT_CLASSES4_RESULT
- PGSQL_CB_GET_ALL_CLIENT_CLASSES6
- PGSQL_CB_GET_ALL_CLIENT_CLASSES6_RESULT
- PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4
- PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS4_RESULT
- PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6
- PGSQL_CB_GET_ALL_GLOBAL_PARAMETERS6_RESULT
- PGSQL_CB_GET_ALL_OPTIONS4
- PGSQL_CB_GET_ALL_OPTIONS4_RESULT
- PGSQL_CB_GET_ALL_OPTIONS6
- PGSQL_CB_GET_ALL_OPTIONS6_RESULT
- PGSQL_CB_GET_ALL_OPTION_DEFS4
- PGSQL_CB_GET_ALL_OPTION_DEFS4_RESULT
- PGSQL_CB_GET_ALL_OPTION_DEFS6
- PGSQL_CB_GET_ALL_OPTION_DEFS6_RESULT
- PGSQL_CB_GET_ALL_SERVERS4
- PGSQL_CB_GET_ALL_SERVERS4_RESULT
- PGSQL_CB_GET_ALL_SERVERS6
- PGSQL_CB_GET_ALL_SERVERS6_RESULT
- PGSQL_CB_GET_ALL_SHARED_NETWORKS4
- PGSQL_CB_GET_ALL_SHARED_NETWORKS4_RESULT
- PGSQL_CB_GET_ALL_SHARED_NETWORKS6
- PGSQL_CB_GET_ALL_SHARED_NETWORKS6_RESULT
- PGSQL_CB_GET_ALL_SUBNETS4
- PGSQL_CB_GET_ALL_SUBNETS4_RESULT
- PGSQL_CB_GET_ALL_SUBNETS6
- PGSQL_CB_GET_ALL_SUBNETS6_RESULT
- PGSQL_CB_GET_CLIENT_CLASS4
- PGSQL_CB_GET_CLIENT_CLASS6
- PGSQL_CB_GET_GLOBAL_PARAMETER4
- PGSQL_CB_GET_GLOBAL_PARAMETER6
- PGSQL_CB_GET_HOST4
- PGSQL_CB_GET_HOST6
- PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4

- PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES4_RESULT
- PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6
- PGSQL_CB_GET_MODIFIED_CLIENT_CLASSES6_RESULT
- PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4
- PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS4_RESULT
- PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6
- PGSQL_CB_GET_MODIFIED_GLOBAL_PARAMETERS6_RESULT
- PGSQL_CB_GET_MODIFIED_OPTIONS4
- PGSQL_CB_GET_MODIFIED_OPTIONS4_RESULT
- PGSQL_CB_GET_MODIFIED_OPTIONS6
- PGSQL_CB_GET_MODIFIED_OPTIONS6_RESULT
- PGSQL_CB_GET_MODIFIED_OPTION_DEFS4
- PGSQL_CB_GET_MODIFIED_OPTION_DEFS4_RESULT
- PGSQL_CB_GET_MODIFIED_OPTION_DEFS6
- PGSQL_CB_GET_MODIFIED_OPTION_DEFS6_RESULT
- PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4
- PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS4_RESULT
- PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6
- PGSQL_CB_GET_MODIFIED_SHARED_NETWORKS6_RESULT
- PGSQL_CB_GET_MODIFIED_SUBNETS4
- PGSQL_CB_GET_MODIFIED_SUBNETS4_RESULT
- PGSQL_CB_GET_MODIFIED_SUBNETS6
- PGSQL_CB_GET_MODIFIED_SUBNETS6_RESULT
- PGSQL_CB_GET_OPTION4
- PGSQL_CB_GET_OPTION6
- PGSQL_CB_GET_OPTION_DEF4
- PGSQL_CB_GET_OPTION_DEF6
- PGSQL_CB_GET_PORT4
- PGSQL_CB_GET_PORT6
- PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4
- PGSQL_CB_GET_RECENT_AUDIT_ENTRIES4_RESULT
- PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6
- PGSQL_CB_GET_RECENT_AUDIT_ENTRIES6_RESULT
- PGSQL_CB_GET_SERVER4
- PGSQL_CB_GET_SERVER6
- PGSQL_CB_GET_SHARED_NETWORK4

- PGSQL_CB_GET_SHARED_NETWORK6
- PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4
- PGSQL_CB_GET_SHARED_NETWORK_SUBNETS4_RESULT
- PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6
- PGSQL_CB_GET_SHARED_NETWORK_SUBNETS6_RESULT
- PGSQL_CB_GET_SUBNET4_BY_PREFIX
- PGSQL_CB_GET_SUBNET4_BY_SUBNET_ID
- PGSQL_CB_GET_SUBNET6_BY_PREFIX
- PGSQL_CB_GET_SUBNET6_BY_SUBNET_ID
- PGSQL_CB_GET_TYPE4
- PGSQL_CB_GET_TYPE6
- PGSQL_CB_REGISTER_BACKEND_TYPE4
- PGSQL_CB_REGISTER_BACKEND_TYPE6
- PGSQL_CB_UNREGISTER_BACKEND_TYPE4
- PGSQL_CB_UNREGISTER_BACKEND_TYPE6
- PGSQL_LB_UPGRADE_EXTENDED_INFO4
- PGSQL_LB_UPGRADE_EXTENDED_INFO4_ERROR
- PGSQL_LB_UPGRADE_EXTENDED_INFO6
- PGSQL_LB_UPGRADE_EXTENDED_INFO6_ERROR
- PING_CHECK_CHANNEL_MALFORMED_PACKET_RECEIVED
- PING_CHECK_CHANNEL_SOCKET_CLOSED
- PING_CHECK_CHANNEL_SOCKET_OPENED
- PING_CHECK_CHANNEL_STOP
- PING_CHECK_DUPLICATE_CHECK
- PING_CHECK_MGR_LEASE_FREE_TO_USE
- PING_CHECK_MGR_RECEIVED_ECHO_REPLY
- PING_CHECK_MGR_START_PING_CHECK
- PING_CHECK_MGR_STOPPING
- RADIUS_ACCESS_CACHE_GET
- RADIUS_ACCESS_CACHE_INSERT
- RADIUS_ACCESS_CONFLICT
- RADIUS_ACCESS_DROP_PARKED_QUERY
- RADIUS_ACCESS_GET_IDENTIFIER
- RADIUS_ACCESS_MAX_PENDING_REQUESTS
- RADIUS_ACCESS_RESUME_PARKED_QUERY
- RADIUS_ACCESS_SUBNET_RESELECT

- RADIUS_ACCOUNTING_ASYNC
- RADIUS_ACCOUNTING_ASYNC_FAILED
- RADIUS_ACCOUNTING_ASYNC_SUCCEED
- RADIUS_ACCOUNTING_NO_HISTORY
- RADIUS_ACCOUNTING_SYNC
- RADIUS_ACCOUNTING_SYNC_FAILED
- RADIUS_ACCOUNTING_SYNC_SUCCEED
- RADIUS_AUTHENTICATION_ASYNC
- RADIUS_AUTHENTICATION_ASYNC_ACCEPTED
- RADIUS_AUTHENTICATION_ASYNC_FAILED
- RADIUS_AUTHENTICATION_ASYNC_REJECTED
- RADIUS_AUTHENTICATION_SYNC
- RADIUS_AUTHENTICATION_SYNC_ACCEPTED
- RADIUS_AUTHENTICATION_SYNC_FAILED
- RADIUS_AUTHENTICATION_SYNC_REJECTED
- RADIUS_BACKEND_GET4
- RADIUS_BACKEND_GET6
- RADIUS_DECODE_MESSAGE
- RADIUS_ENCODE_MESSAGE
- RADIUS_EXCHANGE_RECEIVED
- RADIUS_EXCHANGE_RECEIVED_ACCESS_ACCEPT
- RADIUS_EXCHANGE_RECEIVED_ACCESS_REJECT
- RADIUS_EXCHANGE_RECEIVED_ACCOUNTING_RESPONSE
- RADIUS_EXCHANGE_RECEIVED_RESPONSE
- RADIUS_EXCHANGE_SEND_NEW
- RADIUS_EXCHANGE_SEND_RETRY
- RADIUS_EXCHANGE_SENT
- RADIUS_EXCHANGE_START
- RADIUS_EXCHANGE_SYNC_RETURN
- RADIUS_EXCHANGE_TERMINATE
- RBAC_CONFIGURED_ACLS
- RBAC_CONFIGURED_COMMANDS
- RBAC_CONFIGURED_ROLES
- RBAC_TRACE_HTTP_AUTH_ACCEPT
- RBAC_TRACE_HTTP_AUTH_BAD_BODY_TYPE
- RBAC_TRACE_HTTP_AUTH_BAD_COMMAND_TYPE

- RBAC_TRACE_HTTP_AUTH_COMMAND
- RBAC_TRACE_HTTP_AUTH_DISABLED
- RBAC_TRACE_HTTP_AUTH_EMPTY_BODY
- RBAC_TRACE_HTTP_AUTH_NO_COMMAND
- RBAC_TRACE_HTTP_AUTH_NO_JSON
- RBAC_TRACE_HTTP_AUTH_NO_REQUEST
- RBAC_TRACE_HTTP_AUTH_NO_TLS_REJECT
- RBAC_TRACE_HTTP_AUTH_REJECT
- RBAC_TRACE_HTTP_AUTH_RESPONSE
- RBAC_TRACE_HTTP_AUTH_ROLE
- RBAC_TRACE_HTTP_RESPONSE_BAD_BODY_TYPE
- RBAC_TRACE_HTTP_RESPONSE_CONTEXT
- RBAC_TRACE_HTTP_RESPONSE_DISABLED
- RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY
- RBAC_TRACE_HTTP_RESPONSE_EMPTY_BODY_LIST
- RBAC_TRACE_HTTP_RESPONSE_MODIFIED
- RBAC_TRACE_HTTP_RESPONSE_NO_ARGUMENTS
- START_REKEY_TIMER
- START_RETRY_TIMER
- STAT_CMDS_LEASE4_ORPHANED_STATS
- STAT_CMDS_LEASE6_ORPHANED_STATS
- TCP_CLIENT_REQUEST_RECEIVED
- TCP_CONNECTION_SHUTDOWN
- TCP_CONNECTION_STOP
- TCP_SERVER_RESPONSE_SEND
- TKEY_EXCHANGE_ANSWER_CLASS
- TKEY_EXCHANGE_NOT_A_RESPONSE
- TKEY_EXCHANGE_OUT_TOKEN_NOT_EMPTY
- TKEY_EXCHANGE_RDATA_COUNT
- TKEY_EXCHANGE_RECEIVE_MESSAGE
- TKEY_EXCHANGE_RESPONSE_TTL
- TKEY_EXCHANGE_SEND_MESSAGE
- TKEY_EXCHANGE_VALID
- TKEY_EXCHANGE_VERIFIED
- TLS_SERVER_RESPONSE_SEND

26.51.6 Messages printed on debuglevel 45

- DHCP4_DHCP406_SUBNET_SELECTED
- DHCP4_SUBNET_DYNAMICALLY_CHANGED
- DHCP4_SUBNET_SELECTED
- DHCP6_SUBNET_DYNAMICALLY_CHANGED
- DHCP6_SUBNET_SELECTED
- HOOKS_CALLOUTS_BEGIN
- HOOKS_CALLOUTS_COMPLETE
- HOOKS_CALLOUTS_REMOVED
- HOOKS_CALLOUT_REGISTRATION
- HOOKS_LIBRARY_MULTI_THREADING_COMPATIBLE
- HOOKS_LIBRARY_VERSION
- HOOKS_STD_CALLOUT_REGISTERED
- HOSTS_CFG_GET_ALL_ADDRESS4_COUNT
- HOSTS_CFG_GET_ALL_ADDRESS6_COUNT
- HOSTS_CFG_GET_ALL_COUNT
- HOSTS_CFG_GET_ALL_HOSTNAME_COUNT
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_COUNT
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_COUNT
- HOSTS_CFG_GET_ALL_IDENTIFIER_COUNT
- HOSTS_CFG_GET_ALL_SUBNET_ID4_COUNT
- HOSTS_CFG_GET_ALL_SUBNET_ID6_COUNT
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_COUNT
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_COUNT
- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_HOST
- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS4_NULL
- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_HOST
- HOSTS_CFG_GET_ONE_SUBNET_ID_ADDRESS6_NULL
- HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_HOST
- HOSTS_CFG_GET_ONE_SUBNET_ID_IDENTIFIER_NULL
- HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_HOST
- HOSTS_MGR_ALTERNATE_GET4_SUBNET_ID_IDENTIFIER_NULL
- HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_HOST
- HOSTS_MGR_ALTERNATE_GET6_SUBNET_ID_IDENTIFIER_NULL
- HOST_CACHE_ADD
- HOST_CACHE_ADD_DUPLICATE

- HOST_CACHE_DEL_SUBNET_ID_ADDRESS4
- HOST_CACHE_DEL_SUBNET_ID_ADDRESS6
- HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER4
- HOST_CACHE_DEL_SUBNET_ID_IDENTIFIER6
- HOST_CACHE_GET_ONE_PREFIX_HOST
- HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS4_HOST
- HOST_CACHE_GET_ONE_SUBNET_ID_ADDRESS6_HOST
- HOST_CACHE_GET_ONE_SUBNET_ID_IDENTIFIER_HOST
- HTTP_BAD_CLIENT_REQUEST_RECEIVED_DETAILS
- HTTP_BAD_SERVER_RESPONSE_RECEIVED_DETAILS
- HTTP_CLIENT_REQUEST_RECEIVED_DETAILS
- HTTP_SERVER_RESPONSE_RECEIVED_DETAILS
- HTTP_SERVER_RESPONSE_SEND_DETAILS

26.51.7 Messages printed on debuglevel 50

- ALLOC_ENGINE_V6_EXTEND_LEASE
- ASIODNS_READ_TIMEOUT
- DHCP4_ADDITIONAL_CLASS_EVAL_RESULT
- DHCP4_BUFFER_UNPACK
- DHCP4_BUFFER_WAIT_SIGNAL
- DHCP4_CLIENTID_IGNORED_FOR_LEASES
- DHCP4_CLIENT_FQDN_PROCESS
- DHCP4_CLIENT_FQDN_SCRUBBED_EMPTY
- DHCP4_CLIENT_HOSTNAME_MALFORMED
- DHCP4_CLIENT_HOSTNAME_PROCESS
- DHCP4_CLIENT_HOSTNAME_SCRUBBED_EMPTY
- DHCP4_DEFERRED_OPTION_MISSING
- DHCP4_DEFERRED_OPTION_UNPACK_FAIL
- DHCP4_DHCP4O6_BAD_PACKET
- DHCP4_DHCP4O6_RECEIVE_FAIL
- DHCP4_DHCP4O6_RECEIVING
- DHCP4_DHCP4O6_SUBNET_SELECTION_FAILED
- DHCP4_DISCOVER
- DHCP4_EMPTY_HOSTNAME
- DHCP4_HOOK_BUFFER_RCVD_SKIP
- DHCP4_INFORM_DIRECT_REPLY
- DHCP4_NO_LEASE_INIT_REBOOT

- DHCP4_PACKET_NAK_0002
- DHCP4_PACKET_NAK_0003
- DHCP4_PACKET_NAK_0004
- DHCP4_PACKET_OPTIONS_SKIPPED
- DHCP4_PACKET_PACK
- DHCP4_RELEASE
- DHCP4_RELEASE_FAIL
- DHCP4_RELEASE_FAIL_NO_LEASE
- DHCP4_RELEASE_FAIL_WRONG_CLIENT
- DHCP4_REQUEST
- DHCP4_RESPONSE_HOSTNAME_GENERATE
- DHCP4_SUBNET_SELECTION_FAILED
- DHCP4_UNKNOWN_ADDRESS_REQUESTED
- DHCP6_ADDITIONAL_CLASS_EVAL_RESULT
- DHCP6_ADD_GLOBAL_STATUS_CODE
- DHCP6_ADD_STATUS_CODE_FOR_IA
- DHCP6_BUFFER_UNPACK
- DHCP6_BUFFER_WAIT_SIGNAL
- DHCP6_CLIENT_FQDN_SCRUBBED_EMPTY
- DHCP6_DDNS_CREATE_ADD_NAME_CHANGE_REQUEST
- DHCP6_DDNS_GENERATE_FQDN
- DHCP6_DDNS_RECEIVE_FQDN
- DHCP6_DDNS_REMOVE_OLD_LEASE_FQDN
- DHCP6_DDNS_RESPONSE_FQDN_DATA
- DHCP6_DECLINE_PROCESS_IA
- DHCP6_DHCP4O6_RECEIVE_FAIL
- DHCP6_DHCP4O6_RECEIVING
- DHCP6_HOOK_BUFFER_RCVD_SKIP
- DHCP6_HOOK_DECLINE_SKIP
- DHCP6_LEASE_ADVERT
- DHCP6_LEASE_ADVERT_FAIL
- DHCP6_LEASE_ALLOC
- DHCP6_LEASE_ALLOC_FAIL
- DHCP6_PACKET_OPTIONS_SKIPPED
- DHCP6_PD_LEASE_ADVERT
- DHCP6_PD_LEASE_ADVERT_FAIL

- DHCP6_PD_LEASE_ALLOC
- DHCP6_PD_LEASE_ALLOC_FAIL
- DHCP6_PROCESS_IA_NA_EXTEND
- DHCP6_PROCESS_IA_NA_RELEASE
- DHCP6_PROCESS_IA_NA_REQUEST
- DHCP6_PROCESS_IA_NA_SOLICIT
- DHCP6_PROCESS_IA_PD_EXTEND
- DHCP6_PROCESS_IA_PD_REQUEST
- DHCP6_PROCESS_IA_PD_SOLICIT
- DHCP6_RAPID_COMMIT
- DHCP6_SUBNET_SELECTION_FAILED
- DHCPSRV_DHCP_DDNS_NCR_SENT
- DHCPSRV_EVAL_RESULT
- DHCPSRV_MEMFILE_ADD_ADDR4
- DHCPSRV_MEMFILE_ADD_ADDR6
- DHCPSRV_MEMFILE_COMMIT
- DHCPSRV_MEMFILE_DELETE_ADDR4
- DHCPSRV_MEMFILE_DELETE_ADDR6
- DHCPSRV_MEMFILE_DELETE_EXPIRED_RECLAIMED4
- DHCPSRV_MEMFILE_DELETE_EXPIRED_RECLAIMED6
- DHCPSRV_MEMFILE_DELETE_EXPIRED_RECLAIMED_START
- DHCPSRV_MEMFILE_GET4
- DHCPSRV_MEMFILE_GET6
- DHCPSRV_MEMFILE_GET6_DUID
- DHCPSRV_MEMFILE_GET_ADDR4
- DHCPSRV_MEMFILE_GET_ADDR6
- DHCPSRV_MEMFILE_GET_CLIENTID
- DHCPSRV_MEMFILE_GET_EXPIRED4
- DHCPSRV_MEMFILE_GET_EXPIRED6
- DHCPSRV_MEMFILE_GET_HOSTNAME4
- DHCPSRV_MEMFILE_GET_HOSTNAME6
- DHCPSRV_MEMFILE_GET_HWADDR
- DHCPSRV_MEMFILE_GET_IAID_DUID
- DHCPSRV_MEMFILE_GET_IAID_SUBID_DUID
- DHCPSRV_MEMFILE_GET_PAGE4
- DHCPSRV_MEMFILE_GET_PAGE6

- DHCPSRV_MEMFILE_GET_RELAYID4
- DHCPSRV_MEMFILE_GET_RELAYID6
- DHCPSRV_MEMFILE_GET_REMOTEID4
- DHCPSRV_MEMFILE_GET_REMOTEID6
- DHCPSRV_MEMFILE_GET_SUBID4
- DHCPSRV_MEMFILE_GET_SUBID6
- DHCPSRV_MEMFILE_GET_SUBID_CLIENTID
- DHCPSRV_MEMFILE_GET_SUBID_HWADDR
- DHCPSRV_MEMFILE_GET_SUBID_PAGE6
- DHCPSRV_MEMFILE_ROLLBACK
- DHCPSRV_MEMFILE_UPDATE_ADDR4
- DHCPSRV_MEMFILE_UPDATE_ADDR6
- DHCPSRV_QUEUE_NCR_SKIP
- DHCPSRV_TEMPLATE_EVAL_RESULT
- DHCPSRV_TIMERMGR_RUN_TIMER_OPERATION
- DHCP_DDNS_INVALID_RESPONSE
- DHCP_DDNS_STARTING_TRANSACTION
- DHCP_DDNS_UPDATE_REQUEST_SENT
- DHCP_DDNS_UPDATE_RESPONSE_RECEIVED
- FUZZ_DATA_READ
- FUZZ_SEND
- HTTPS_REQUEST_RECEIVE_START
- HTTP_CLIENT_REQUEST_SEND
- HTTP_CLIENT_REQUEST_TIMEOUT_OCCURRED
- HTTP_CONNECTION_HANDSHAKE_START
- HTTP_IDLE_CONNECTION_TIMEOUT_OCCURRED
- HTTP_REQUEST_RECEIVE_START
- LEGAL_LOG_MYSQL_GET_VERSION
- LEGAL_LOG_MYSQL_INSERT_LOG
- LEGAL_LOG_MYSQL_TLS_CIPHER
- LEGAL_LOG_PGSQL_GET_VERSION
- LEGAL_LOG_PGSQL_INSERT_LOG
- MYSQL_HB_DB_GET_VERSION
- MYSQL_LB_ADD_ADDR4
- MYSQL_LB_ADD_ADDR6
- MYSQL_LB_COMMIT

- MYSQL_LB_DELETED_EXPIRED_RECLAIMED
- MYSQL_LB_DELETED_SUBNET4_ID
- MYSQL_LB_DELETED_SUBNET6_ID
- MYSQL_LB_DELETE_ADDR4
- MYSQL_LB_DELETE_ADDR6
- MYSQL_LB_DELETE_EXPIRED_RECLAIMED4
- MYSQL_LB_DELETE_EXPIRED_RECLAIMED6
- MYSQL_LB_GET4
- MYSQL_LB_GET6
- MYSQL_LB_GET_ADDR4
- MYSQL_LB_GET_ADDR6
- MYSQL_LB_GET_CLIENTID
- MYSQL_LB_GET_DUID
- MYSQL_LB_GET_EXPIRED4
- MYSQL_LB_GET_EXPIRED6
- MYSQL_LB_GET_HOSTNAME4
- MYSQL_LB_GET_HOSTNAME6
- MYSQL_LB_GET_HWADDR
- MYSQL_LB_GET_IAID_DUID
- MYSQL_LB_GET_IAID_SUBID_DUID
- MYSQL_LB_GET_PAGE4
- MYSQL_LB_GET_PAGE6
- MYSQL_LB_GET_RELAYID4
- MYSQL_LB_GET_RELAYID6
- MYSQL_LB_GET_REMOTEID4
- MYSQL_LB_GET_REMOTEID6
- MYSQL_LB_GET_SUBID4
- MYSQL_LB_GET_SUBID6
- MYSQL_LB_GET_SUBID_CLIENTID
- MYSQL_LB_GET_SUBID_HWADDR
- MYSQL_LB_GET_SUBID_PAGE6
- MYSQL_LB_GET_VERSION
- MYSQL_LB_ROLLBACK
- MYSQL_LB_UPDATE_ADDR4
- MYSQL_LB_UPDATE_ADDR6
- MYSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE

- MYSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE
- PERFMON_DHCP4_PKT_EVENTS
- PERFMON_DHCP4_PKT_PROCESS_ERROR
- PERFMON_DHCP6_PKT_EVENTS
- PERFMON_DHCP6_PKT_PROCESS_ERROR
- PGSQL_HB_DB_GET_VERSION
- PGSQL_LB_ADD_ADDR4
- PGSQL_LB_ADD_ADDR6
- PGSQL_LB_COMMIT
- PGSQL_LB_DELETED_SUBNET4_ID
- PGSQL_LB_DELETED_SUBNET6_ID
- PGSQL_LB_DELETE_ADDR4
- PGSQL_LB_DELETE_ADDR6
- PGSQL_LB_DELETE_EXPIRED_RECLAIMED4
- PGSQL_LB_DELETE_EXPIRED_RECLAIMED6
- PGSQL_LB_GET4
- PGSQL_LB_GET6
- PGSQL_LB_GET_ADDR4
- PGSQL_LB_GET_ADDR6
- PGSQL_LB_GET_CLIENTID
- PGSQL_LB_GET_DUID
- PGSQL_LB_GET_EXPIRED4
- PGSQL_LB_GET_EXPIRED6
- PGSQL_LB_GET_HOSTNAME4
- PGSQL_LB_GET_HOSTNAME6
- PGSQL_LB_GET_HWADDR
- PGSQL_LB_GET_IAID_DUID
- PGSQL_LB_GET_IAID_SUBID_DUID
- PGSQL_LB_GET_PAGE4
- PGSQL_LB_GET_PAGE6
- PGSQL_LB_GET_RELAYID4
- PGSQL_LB_GET_RELAYID6
- PGSQL_LB_GET_REMOTEID4
- PGSQL_LB_GET_REMOTEID6
- PGSQL_LB_GET_SUBID4
- PGSQL_LB_GET_SUBID6

- PGSQL_LB_GET_SUBID_CLIENTID
- PGSQL_LB_GET_SUBID_HWADDR
- PGSQL_LB_GET_SUBID_PAGE6
- PGSQL_LB_GET_VERSION
- PGSQL_LB_ROLLBACK
- PGSQL_LB_UPDATE_ADDR4
- PGSQL_LB_UPDATE_ADDR6
- PGSQL_LB_UPGRADE_EXTENDED_INFO4_PAGE
- PGSQL_LB_UPGRADE_EXTENDED_INFO6_PAGE
- PING_CHECK_CHANNEL_ECHO_REPLY_RECEIVED
- PING_CHECK_CHANNEL_ECHO_REQUEST_SENT
- PING_CHECK_MGR_NEXT_ECHO_SCHEDULED
- PING_CHECK_MGR_RECEIVED_UNEXPECTED_ECHO_REPLY
- PING_CHECK_MGR_RECEIVED_UNEXPECTED_UNREACHABLE_MSG
- PING_CHECK_MGR_RECEIVED_UNREACHABLE_MSG
- PING_CHECK_MGR_REPLY_TIMEOUT_EXPIRED
- PING_CHECK_NO_LEASE_OR_LEASE_REUSED
- TCP_CONNECTION_REJECTED_BY_FILTER
- TCP_IDLE_CONNECTION_TIMEOUT_OCCURRED
- TCP_REQUEST_RECEIVE_START
- TLS_CONNECTION_HANDSHAKE_START
- TLS_REQUEST_RECEIVE_START

26.51.8 Messages printed on debuglevel 55

- ALLOC_ENGINE_V4_REUSE_EXPIRED_LEASE_DATA
- ALLOC_ENGINE_V6_EXTEND_LEASE_DATA
- ALLOC_ENGINE_V6_EXTEND_NEW_LEASE_DATA
- ALLOC_ENGINE_V6_REUSE_EXPIRED_LEASE_DATA
- DHCP4_CLIENT_FQDN_DATA
- DHCP4_CLIENT_HOSTNAME_DATA
- DHCP4_CLIENT_NAME_PROC_FAIL
- DHCP4_DHCP4O6_RESPONSE_DATA
- DHCP4_DHCP4O6_SUBNET_DATA
- DHCP4_GENERATE_FQDN
- DHCP4_QUERY_DATA
- DHCP4_RECOVERED_STASHED_RELAY_AGENT_INFO
- DHCP4_RESERVED_HOSTNAME_ASSIGNED

- DHCP4_RESPONSE_DATA
- DHCP4_RESPONSE_FQDN_DATA
- DHCP4_RESPONSE_HOSTNAME_DATA
- DHCP4_SUBNET_DATA
- DHCP6_DDNS_FQDN_GENERATED
- DHCP6_DHCP4O6_RESPONSE_DATA
- DHCP6_LEASE_DATA
- DHCP6_QUERY_DATA
- DHCP6_RESPONSE_DATA
- DHCP6_SUBNET_DATA
- DHCP_SRV_DDNS_TTL_TOO_LARGE
- DHCP_SRV_DDNS_TTL_TOO_SMALL
- DHCP_SRV_MEMFILE_LEASE_LOAD
- DHCP_SRV_QUEUE_NCR
- DHCP_DDNS_AT_MAX_TRANSACTIONS
- DHCP_DDNS_FWD_REQUEST_IGNORED
- DHCP_DDNS_NO_ELIGIBLE_JOBS
- DHCP_DDNS_QUEUE_MGR_QUEUE_RECEIVE
- DHCP_DDNS_REQUEST_DROPPED
- DHCP_DDNS_REV_REQUEST_IGNORED
- EVAL_DEBUG_AND
- EVAL_DEBUG_BRANCH
- EVAL_DEBUG_CONCAT
- EVAL_DEBUG_EQUAL
- EVAL_DEBUG_HEXSTRING
- EVAL_DEBUG_IFELSE_FALSE
- EVAL_DEBUG_IFELSE_TRUE
- EVAL_DEBUG_INT16TOTEXT
- EVAL_DEBUG_INT32TOTEXT
- EVAL_DEBUG_INT8TOTEXT
- EVAL_DEBUG_IPADDRESS
- EVAL_DEBUG_IPADDRESSTOTEXT
- EVAL_DEBUG_LCASE
- EVAL_DEBUG_MATCH
- EVAL_DEBUG_MEMBER
- EVAL_DEBUG_NOT

- EVAL_DEBUG_OPTION
- EVAL_DEBUG_OR
- EVAL_DEBUG_PKT
- EVAL_DEBUG_PKT4
- EVAL_DEBUG_PKT6
- EVAL_DEBUG_POP_AND_BRANCH_FALSE
- EVAL_DEBUG_POP_OR_BRANCH_FALSE
- EVAL_DEBUG_POP_OR_BRANCH_TRUE
- EVAL_DEBUG_RELAY6
- EVAL_DEBUG_RELAY6_RANGE
- EVAL_DEBUG_SPLIT
- EVAL_DEBUG_SPLIT_DELIM_EMPTY
- EVAL_DEBUG_SPLIT_EMPTY
- EVAL_DEBUG_SPLIT_FIELD_OUT_OF_RANGE
- EVAL_DEBUG_STRING
- EVAL_DEBUG_SUBSTRING
- EVAL_DEBUG_SUBSTRING_EMPTY
- EVAL_DEBUG_SUBSTRING_RANGE
- EVAL_DEBUG_TOHEXSTRING
- EVAL_DEBUG_UCASE
- EVAL_DEBUG_UINT16TOTEXT
- EVAL_DEBUG_UINT32TOTEXT
- EVAL_DEBUG_UINT8TOTEXT
- EVAL_DEBUG_VENDOR_CLASS_DATA
- EVAL_DEBUG_VENDOR_CLASS_DATA_NOT_FOUND
- EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID
- EVAL_DEBUG_VENDOR_CLASS_ENTERPRISE_ID_MISMATCH
- EVAL_DEBUG_VENDOR_CLASS_EXISTS
- EVAL_DEBUG_VENDOR_CLASS_NO_OPTION
- EVAL_DEBUG_VENDOR_ENTERPRISE_ID
- EVAL_DEBUG_VENDOR_ENTERPRISE_ID_MISMATCH
- EVAL_DEBUG_VENDOR_EXISTS
- EVAL_DEBUG_VENDOR_NO_OPTION
- HOOKS_ALL_CALLOUTS_DEREGISTERED
- HOOKS_CALLOUT_CALLED
- HOOKS_CALLOUT_DEREGISTERED

- HOSTS_CFG_GET_ALL_ADDRESS4_HOST
- HOSTS_CFG_GET_ALL_ADDRESS6_HOST
- HOSTS_CFG_GET_ALL_HOST
- HOSTS_CFG_GET_ALL_HOSTNAME_HOST
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID4_HOST
- HOSTS_CFG_GET_ALL_HOSTNAME_SUBNET_ID6_HOST
- HOSTS_CFG_GET_ALL_IDENTIFIER_HOST
- HOSTS_CFG_GET_ALL_SUBNET_ID4_HOST
- HOSTS_CFG_GET_ALL_SUBNET_ID6_HOST
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS4_HOST
- HOSTS_CFG_GET_ALL_SUBNET_ID_ADDRESS6_HOST
- HOSTS_CFG_GET_ONE_PREFIX_HOST
- HOSTS_CFG_GET_ONE_PREFIX_NULL
- HTTP_CLIENT_REQUEST_SEND_DETAILS
- HTTP_DATA_RECEIVED
- LIMITS_PACKET_WITH_SUBNET_ID_RATE_NO_SUBNET
- LIMITS_PACKET_WITH_CLIENT_CLASSES_RATE_LIMIT_HONORED
- LIMITS_PACKET_WITH_SUBNET_ID_RATE_LIMIT_HONORED
- NETCONF_CONFIG_CHANGED_DETAIL
- NETCONF_GET_CONFIG
- NETCONF_SET_CONFIG
- NETCONF_UPDATE_CONFIG
- NETCONF_VALIDATE_CONFIG
- TCP_DATA_RECEIVED
- TCP_DATA_SENT

26.51.9 Messages printed on debuglevel 70

- ASIODNS_FETCH_COMPLETED

CONFIGURATION TEMPLATES

The following sections include configuration templates for certain deployment types. The example configuration files are also available in the Kea sources, in the `doc/examples` directory.

27.1 Template: Home Network of a Power User

Below are some templates to assist in configuring the home network of a power user; they may also be appropriate for a small office. These templates make the following assumptions:

- The administrator wants to use a single /24 class of IPv4 addresses.
- High Availability (HA) is desired, so there are two DHCP servers.
- There are a handful of devices, and some of them (e.g. a printer or NAS) require static addresses or extra options.
- The administrator does not want to be bothered with database management.
- The setup is optimized for minimal to zero maintenance.
- Performance is not an issue; hundreds of queries per second are not expected.
- IPv6 is not used.
- DNS updates will not be performed by Kea.

The logical setup consists of two hosts, each running a Kea DHCPv4 server and a Control Agent (CA). The server connects with the CA using UNIX sockets. Each DHCPv4+CA acts as one partner of the HA pair.

```
+--host-1--+      +-host-2--+
|           |      |           |
|   CA <===\    /===> CA   |   ===== - HTTP connection
|   #       | \ /       |   #   |
|   #       |  X       |   #   |   ##### - UNIX socket
|   #       | / \       |   #   |
| DHCPv4 ==/   \== DHCPv4 |
|           |      |           |
+-----+      +-----+
```

The CAs on host-1 and host-2 both listen on port 8000. The DHCP servers communicate with each other via the CAs, which forward control commands to the DHCP servers over the UNIX domain sockets.

27.1.1 Deployment Considerations

This setup is not expected to be very performant. Most modest hardware will do; Kea has been successfully deployed on Raspberry Pi platforms, for example. If it is running on a VM, 2GB of RAM with one CPU core should be sufficient. Ubuntu LTS is a choice that is easy to set up and is low maintenance; however, any Linux or FreeBSD operating system

is fine. Less popular systems, such as OpenBSD or NetBSD, should also work in principle, but they are not regularly tested.

In this example, there are two hosts running Kea:

- 192.168.1.2 - primary HA server (active, handles all the traffic)
- 192.168.1.3 - secondary HA server (passive, ready to take over if the primary fails)

The network is 192.168.1.0/24. It is assumed that 192.168.1.1 is the default router.

The whole subnet is split into dynamic and static pools:

- 192.168.1.100 - 192.168.1.199 - this is the dynamic pool. When new devices appear in the network, they are assigned dynamic addresses from this pool.
- The reservations are done outside of this dynamic range (depending on the addressing preference, either 192.168.1.1-192.168.1.99 or 192.168.1.200-192.168.1.254).

To deploy this setup, perform the following steps:

1. Install the CA and the DHCPv4 server on host-1, and copy the configuration files to their typical locations. They are usually in `/etc/kea` on Linux and `/usr/local/etc/kea` on FreeBSD, and the files are typically called `kea-ctrl-agent.conf` and `kea-dhcp4.conf`. Please consult the startup scripts for any specific system.
2. Alter the following to match the local setup:
 - The interface name that Kea should listen on (`interfaces` in `interfaces-config`).
 - The interface name that is used to access the subnet (`interface` in `subnet4`).
 - The addressing, if using something other than 192.168.1.0/24. Make sure the CA port configuration (`http-host` and `http-port` in `kea-ca.conf`) matches the DHCPv4 server configuration (`url` in `hook-libraries/parameters/high-availability/peers` in `kea-dhcp4.conf`).
 - The router option, to match the actual network.
 - The DNS option, to match the actual network.
 - The path to the hook libraries. This is a very OS-specific parameter; the library names are generally the same everywhere, but the path varies. See [Introduction](#) for details.
3. If using a firewall, make sure host-1 can reach host-2. An easy way to ensure that is to try to retrieve host-2's config from host-1:

```
curl -X POST -H "Content-Type: application/json" -d '{ "command": "config-get",  
"service": [ "dhcp4" ] }' http://192.168.1.3:8000/
```

The DHCPv4 running configuration should be returned, in JSON format.
4. Verify that communication between the hosts works in the opposite direction as well (host-2 can connect to host-1), by repeating step 3 from host-2 using host-1's IP address and port.
5. Install the CA and the DHCPv4 server on host-2, as in steps 1 and 2. The config file for the standby server is very similar to the one on the primary server, other than the definition of the `this-server-name` field (and possibly the interface names).

27.1.2 Possible Extensions

This sample configuration is basic but functional. Once it is set up and running, administrators may wish to consider the following changes:

- If there is a local DNS server, DNS updates can be configured via Kea. This requires running a DHCP-DDNS update server (`kea-dhcp-ddns`). See [Overview](#) for details.

- To run Stateful DHCP for IPv6, a `kea-dhcp6` server is necessary. Its configuration is very similar to `kea-dhcp4`, but there are some notable differences: the default gateway is not configured via the DHCPv6 protocol, but via router advertisements sent by the local router. Also, the DHCPv6 concept of prefix delegation does not exist in DHCPv4. See *The DHCPv6 Server* for details.
- To expand the local network, adding a MySQL or PostgreSQL database is a popular solution. Users can choose to store leases, host reservations, and even most of the configuration in a database. See *Kea Database Administration* and the `lease-database`, `hosts-database`, and `config-control` parameters in *The DHCPv4 Server*.
- To provide more insight into how the DHCP server operates, Kea's RESTful API can query for many runtime statistics or even change the configuration during runtime. Users may also consider deploying Stork, which is a rapidly developing dashboard for Kea. See *Monitoring Kea With Stork* for more information.
- All Kea users should read *Kea Security*: to learn about various trade-offs between convenience and security in Kea.

Some tweaking of these templates may be required to match specific system needs: at a minimum, the lines highlighted in yellow must be adjusted to match the actual deployment.

Server1's Control Agent configuration file:

```

1 // This is an example of a configuration for Control-Agent (CA) listening
2 // for incoming HTTP traffic. This is necessary for handling API commands,
3 // in particular lease update commands needed for HA setup.
4 {
5     "Control-agent":
6     {
7         // We need to specify where the agent should listen to incoming HTTP
8         // queries.
9         "http-host": "192.168.1.2",
10
11         // This specifies the port CA will listen on.
12         "http-port": 8000,
13
14         "control-sockets":
15         {
16             // This is how the Agent can communicate with the DHCPv4 server.
17             "dhcp4":
18             {
19                 "comment": "socket to DHCPv4 server",
20                 "socket-type": "unix",
21                 "socket-name": "kea4-ctrl-socket"
22             },
23
24             // Location of the DHCPv6 command channel socket.
25             "dhcp6":
26             {
27                 "socket-type": "unix",
28                 "socket-name": "kea6-ctrl-socket"
29             },
30
31             // Location of the D2 command channel socket.
32             "d2":
33             {
34                 "socket-type": "unix",
35                 "socket-name": "kea-ddns-ctrl-socket",

```

(continues on next page)

(continued from previous page)

```

36     "user-context": { "in-use": false }
37   },
38 },
39
40 // Similar to other Kea components, CA also uses logging.
41 "loggers": [
42   {
43     "name": "kea-ctrl-agent",
44     "output-options": [
45       {
46         "output": "/var/log/kea/kea-ctrl-agent.log",
47
48         // Several additional parameters are possible in addition
49         // to the typical output. Flush determines whether logger
50         // flushes output to a file. Maxsize determines maximum
51         // filesize before the file is rotated. maxver
52         // specifies the maximum number of rotated files being
53         // kept.
54         "flush": true,
55         "maxsize": 204800,
56         "maxver": 4,
57         // We use pattern to specify custom log message layout
58         "pattern": "%d{%y.%m.%d %H:%M:%S.%q} %-5p [%c/%i] %m\n"
59       }
60     ],
61     "severity": "INFO",
62     "debuglevel": 0 // debug level only applies when severity is set to
63     ↪ DEBUG.
64   }
65 ]
66 }
```

Server1's DHCPv4 configuration file:

```

1 // This is an example configuration of the Kea DHCPv4 server 1:
2 //
3 // - uses High Availability hook library and Lease Commands hook library
4 //   to enable High Availability function for the DHCP server. This config
5 //   file is for the primary (the active) server.
6 // - uses memfile, which stores lease data in a local CSV file
7 // - it assumes a single /24 addressing over a link that is directly reachable
8 //   (no DHCP relays)
9 // - there is a handful of IP reservations
10 //
11 // It is expected to run with a standby (the passive) server, which has a very similar
12 // configuration. The only difference is that "this-server-name" must be set to "server2
13 ↪" on the
14 // other server. Also, the interface configuration depends on the network settings of the
15 // particular machine.
16 {
```

(continues on next page)

(continued from previous page)

```

17 "Dhcp4": {
18
19
20     // Add names of your network interfaces to listen on.
21     "interfaces-config": {
22         // The DHCPv4 server listens on this interface. When changing this to
23         // the actual name of your interface, make sure to also update the
24         // interface parameter in the subnet definition below.
25         "interfaces": [ "enp0s8" ]
26     },
27
28     // Control socket is required for communication between the Control
29     // Agent and the DHCP server. High Availability requires Control Agent
30     // to be running because lease updates are sent over the RESTful
31     // API between the HA peers.
32     "control-socket": {
33         "socket-type": "unix",
34         "socket-name": "kea4-ctrl-socket"
35     },
36
37     // Use Memfile lease database backend to store leases in a CSV file.
38     // Depending on how Kea was compiled, it may also support SQL databases
39     // (MySQL and/or PostgreSQL). Those database backends require more
40     // parameters, like name, host and possibly user and password.
41     // There are dedicated examples for each backend. See Section 7.2.2 "Lease
42     // Storage" for details.
43     "lease-database": {
44         // Memfile is the simplest and easiest backend to use. It's an in-memory
45         // database with data being written to a CSV file. It is very similar to
46         // what ISC DHCP does.
47         "type": "memfile"
48     },
49
50     // Let's configure some global parameters. The home network is not very dynamic
51     // and there's no shortage of addresses, so no need to recycle aggressively.
52     "valid-lifetime": 43200, // leases will be valid for 12h
53     "renew-timer": 21600, // clients should renew every 6h
54     "rebind-timer": 32400, // clients should start looking for other servers after 9h
55
56     // Kea will clean up its database of expired leases once per hour. However, it
57     // will keep the leases in expired state for 2 days. This greatly increases the
58     // chances for returning devices to get the same address again. To guarantee that,
59     // use host reservation.
60     // If both "flush-reclaimed-timer-wait-time" and "hold-reclaimed-time" are
61     // not 0, when the client sends a release message the lease is expired
62     // instead of being deleted from lease storage.
63     "expired-leases-processing": {
64         "reclaim-timer-wait-time": 3600,
65         "hold-reclaimed-time": 172800,
66         "max-reclaim-leases": 0,
67         "max-reclaim-time": 0
68     },

```

(continues on next page)

(continued from previous page)

```

69 // HA requires two hook libraries to be loaded: libdhcp_lease_cmds.so and
70 // libdhcp_ha.so. The former handles incoming lease updates from the HA peers.
71 // The latter implements high availability feature for Kea. Note the library name
72 // should be the same, but the path is OS specific.
73 "hooks-libraries": [
74     // The lease_cmds library must be loaded because HA makes use of it to
75     // deliver lease updates to the server as well as synchronize the
76     // lease database after failure.
77     {
78         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_lease_cmds.so"
79     },
80     {
81         // The HA hook library should be loaded.
82         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_ha.so",
83         "parameters": {
84             // Each server should have the same HA configuration, except for the
85             // "this-server-name" parameter.
86             "high-availability": [ {
87                 // This parameter points to this server instance. The respective
88                 // HA peers must have this parameter set to their own names.
89                 "this-server-name": "server1",
90                 // The HA mode is set to hot-standby. In this mode, the active_
91                 ↪ server handles
92                 // all the traffic. The standby takes over if the primary becomes_
93                 ↪ unavailable.
94                 "mode": "hot-standby",
95                 // Heartbeat is to be sent every 10 seconds if no other control
96                 // commands are transmitted.
97                 "heartbeat-delay": 10000,
98                 // Maximum time for partner's response to a heartbeat, after which
99                 // failure detection is started. This is specified in milliseconds.
100                // If we don't hear from the partner in 60 seconds, it's time to
101                // start worrying.
102                "max-response-delay": 60000,
103                // The following parameters control how the server detects the
104                // partner's failure. The ACK delay sets the threshold for the
105                // 'secs' field of the received discovers. This is specified in
106                // milliseconds.
107                "max-ack-delay": 5000,
108                // This specifies the number of clients which send messages to
109                // the partner but appear to not receive any response.
110                "max-unacked-clients": 5,
111                // This specifies the maximum timeout (in milliseconds) for the_
112                ↪ server
113                // to complete sync. If you have a large deployment (high tens or
114                // hundreds of thousands of clients), you may need to increase it
115                // further. The default value is 60000ms (60 seconds).
116                "sync-timeout": 60000,
117                "peers": [
118                    // This is the configuration of this server instance.

```

(continues on next page)

(continued from previous page)

```

118         {
119             "name": "server1",
120             // This specifies the URL of this server instance. The
121             // Control Agent must run along with this DHCPv4 server
122             // instance and the "http-host" and "http-port" must be
123             // set to the corresponding values.
124             "url": "http://192.168.1.2:8005/",
125             // This server is primary. The other one must be
126             // secondary.
127             "role": "primary"
128         },
129         // This is the configuration of the secondary server.
130         {
131             "name": "server2",
132             // Specifies the URL on which the partner's control
133             // channel can be reached. The Control Agent is required
134             // to run on the partner's machine with "http-host" and
135             // "http-port" values set to the corresponding values.
136             "url": "http://192.168.1.3:8005/",
137             // The other server is secondary. This one must be
138             // primary.
139             "role": "standby"
140         }
141     ]
142 }
143
144 ],
145
146
147 // This example contains a single subnet declaration.
148 "subnet4": [
149     {
150         // Subnet prefix.
151         "subnet": "192.168.1.0/24",
152
153         // There are no relays in this network, so we need to tell Kea that this
154         ↪ subnet // is reachable directly via the specified interface.
155         "interface": "enp0s8",
156
157         // Specify a dynamic address pool.
158         "pools": [
159             {
160                 "pool": "192.168.1.100-192.168.1.199"
161             }
162         ],
163
164         // These are options that are subnet specific. In most cases, you need to
165         ↪ define at // least routers option, as without this option your clients will not be
166         ↪ able to reach // their default gateway and will not have Internet connectivity. If you

```

(continues on next page)

(continued from previous page)

```

167 ↪have many           // subnets and they share the same options (e.g. DNS servers typically is
168 ↪the same            // everywhere), you may define options at the global scope, so you don't
169 ↪repeat them         // for every network.
170                     "option-data": [
171                         {
172 ↪router.              // For each IPv4 subnet you typically need to specify at least one
173                         "name": "routers",
174                         "data": "192.168.1.1"
175                         },
176                         {
177                             // Using cloudflare or Quad9 is a reasonable option. Change this
178                             // to your own DNS servers is you have them. Another popular
179                             // choice is 8.8.8.8, owned by Google. Using third party DNS
180                             // service raises some privacy concerns.
181                             "name": "domain-name-servers",
182                             "data": "1.1.1.1,9.9.9.9"
183                         }
184                     ],
185
186 ↪dynamic,            // Some devices should get a static address. Since the .100 - .199 range is
187 ↪reservation         // let's use the lower address space for this. There are many ways how
188 ↪popular one.        // can be defined, but using MAC address (hw-address) is by far the most
189 ↪whatever            // You can use client-id, duid and even custom defined flex-id that may use
190 ↪more things         // parts of the packet you want to use as identifiers. Also, there are many
191 ↪server, hostname,   // you can specify in addition to just an IP address: extra options, next-
192 ↪details.            // assign device to client classes etc. See the Kea ARM, Section 8.3 for
193                     // The reservations are subnet specific.
194                     "reservations": [
195                         {
196                             "hw-address": "1a:1b:1c:1d:1e:1f",
197                             "ip-address": "192.168.1.10"
198                         },
199                         {
200                             "client-id": "01:11:22:33:44:55:66",
201                             "ip-address": "192.168.1.11"
202                         }
203                     ]
204                 }
205             ],
206
207             // Logging configuration starts here.

```

(continues on next page)

(continued from previous page)

```

208     "loggers": [
209     {
210         // This section affects kea-dhcp4, which is the base logger for DHCPv4 component.
211         ↪ It tells
212         // DHCPv4 server to write all log messages (on severity INFO or higher) to a
213         ↪ file. The file
214         // will be rotated once it grows to 2MB and up to 4 files will be kept. The
215         ↪ debuglevel
216         // (range 0 to 99) is used only when logging on DEBUG level.
217         "name": "kea-dhcp4",
218         "output-options": [
219             {
220                 "output": "/var/log/kea/kea-dhcp4.log",
221                 "maxsize": 2048000,
222                 "maxver": 4
223             }
224         ],
225         "severity": "INFO",
226         "debuglevel": 0
227     }
228 ]
229 }
230 }

```

Server2's Control Agent configuration file:

```

1  // This is an example of a configuration for Control-Agent (CA) listening
2  // for incoming HTTP traffic. This is necessary for handling API commands,
3  // in particular lease update commands needed for HA setup.
4  {
5      "Control-agent":
6      {
7          // We need to specify where the agent should listen to incoming HTTP
8          // queries.
9          "http-host": "192.168.1.3",
10
11         // This specifies the port CA will listen on.
12         "http-port": 8000,
13
14         "control-sockets":
15         {
16             // This is how the Agent can communicate with the DHCPv4 server.
17             "dhcp4":
18             {
19                 "comment": "socket to DHCPv4 server",
20                 "socket-type": "unix",
21                 "socket-name": "kea4-ctrl-socket"
22             },
23
24             // Location of the DHCPv6 command channel socket.
25             "dhcp6":
26             {

```

(continues on next page)

(continued from previous page)

```

27         "socket-type": "unix",
28         "socket-name": "kea6-ctrl-socket"
29     },
30
31     // Location of the D2 command channel socket.
32     "d2":
33     {
34         "socket-type": "unix",
35         "socket-name": "kea-ddns-ctrl-socket",
36         "user-context": { "in-use": false }
37     }
38 },
39
40 // Similar to other Kea components, CA also uses logging.
41 "loggers": [
42     {
43         "name": "kea-ctrl-agent",
44         "output-options": [
45             {
46                 "output": "/var/log/kea/kea-ctrl-agent.log",
47
48                 // Several additional parameters are possible in addition
49                 // to the typical output. Flush determines whether logger
50                 // flushes output to a file. Maxsize determines maximum
51                 // filesize before the file is rotated. maxver
52                 // specifies the maximum number of rotated files being
53                 // kept.
54                 "flush": true,
55                 "maxsize": 204800,
56                 "maxver": 4,
57                 // We use pattern to specify custom log message layout
58                 "pattern": "%d{%y.%m.%d %H:%M:%S.%q} %-5p [%c/%i] %m\n"
59             }
60         ],
61         "severity": "INFO",
62         "debuglevel": 0 // debug level only applies when severity is set to
63         ↪ DEBUG.
64     }
65 ]
66 }
```

Server2's DHCPv4 configuration file:

```

1 // This is an example configuration of the Kea DHCPv4 server 2:
2 //
3 // - uses High Availability hook library and Lease Commands hook library
4 //   to enable High Availability function for the DHCP server. This config
5 //   file is for the primary (the active) server.
6 // - uses memfile, which stores lease data in a local CSV file
7 // - it assumes a single /24 addressing over a link that is directly reachable
8 //   (no DHCP relays)
```

(continues on next page)

(continued from previous page)

```

9 // - there is a handful of IP reservations
10 //
11 // It is expected to run with a primary (the active) server, which has a very similar
12 // configuration. The only difference is that "this-server-name" must be set to "server2
13 ↪" on the
14 // other server. Also, the interface configuration depends on the network settings of the
15 // particular machine.
16 {
17
18 "Dhcp4": {
19
20     // Add names of your network interfaces to listen on.
21     "interfaces-config": {
22         // The DHCPv4 server listens on this interface. When changing this to
23         // the actual name of your interface, make sure to also update the
24         // interface parameter in the subnet definition below.
25         "interfaces": [ "enp0s8" ]
26     },
27
28     // Control socket is required for communication between the Control
29     // Agent and the DHCP server. High Availability requires Control Agent
30     // to be running because lease updates are sent over the RESTful
31     // API between the HA peers.
32     "control-socket": {
33         "socket-type": "unix",
34         "socket-name": "kea4-ctrl-socket"
35     },
36
37     // Use Memfile lease database backend to store leases in a CSV file.
38     // Depending on how Kea was compiled, it may also support SQL databases
39     // (MySQL and/or PostgreSQL). Those database backends require more
40     // parameters, like name, host and possibly user and password.
41     // There are dedicated examples for each backend. See Section 7.2.2 "Lease
42     // Storage" for details.
43     "lease-database": {
44         // Memfile is the simplest and easiest backend to use. It's an in-memory
45         // database with data being written to a CSV file. It is very similar to
46         // what ISC DHCP does.
47         "type": "memfile"
48     },
49
50     // Let's configure some global parameters. The home network is not very dynamic
51     // and there's no shortage of addresses, so no need to recycle aggressively.
52     "valid-lifetime": 43200, // leases will be valid for 12h
53     "renew-timer": 21600, // clients should renew every 6h
54     "rebind-timer": 32400, // clients should start looking for other servers after 9h
55
56     // Kea will clean up its database of expired leases once per hour. However, it
57     // will keep the leases in expired state for 2 days. This greatly increases the
58     // chances for returning devices to get the same address again. To guarantee that,
59     // use host reservation.

```

(continues on next page)

(continued from previous page)

```

60 // If both "flush-reclaimed-timer-wait-time" and "hold-reclaimed-time" are
61 // not 0, when the client sends a release message the lease is expired
62 // instead of being deleted from lease storage.
63 "expired-leases-processing": {
64     "reclaim-timer-wait-time": 3600,
65     "hold-reclaimed-time": 172800,
66     "max-reclaim-leases": 0,
67     "max-reclaim-time": 0
68 },
69
70 // HA requires two hook libraries to be loaded: libdhcp_lease_cmds.so and
71 // libdhcp_ha.so. The former handles incoming lease updates from the HA peers.
72 // The latter implements high availability feature for Kea. Note the library name
73 // should be the same, but the path is OS specific.
74 "hooks-libraries": [
75     // The lease_cmds library must be loaded because HA makes use of it to
76     // deliver lease updates to the server as well as synchronize the
77     // lease database after failure.
78     {
79         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_lease_cmds.so"
80     },
81
82     {
83         // The HA hook library should be loaded.
84         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_ha.so",
85         "parameters": {
86             // Each server should have the same HA configuration, except for the
87             // "this-server-name" parameter.
88             "high-availability": [ {
89                 // This parameter points to this server instance. The respective
90                 // HA peers must have this parameter set to their own names.
91                 "this-server-name": "server2",
92                 // The HA mode is set to hot-standby. In this mode, the active_
93                 // server handles
94                 // all the traffic. The standby takes over if the primary becomes_
95                 // unavailable.
96                 "mode": "hot-standby",
97                 // Heartbeat is to be sent every 10 seconds if no other control
98                 // commands are transmitted.
99                 "heartbeat-delay": 10000,
100                // Maximum time for partner's response to a heartbeat, after which
101                // failure detection is started. This is specified in milliseconds.
102                // If we don't hear from the partner in 60 seconds, it's time to
103                // start worrying.
104                "max-response-delay": 60000,
105                // The following parameters control how the server detects the
106                // partner's failure. The ACK delay sets the threshold for the
107                // 'secs' field of the received discovers. This is specified in
108                // milliseconds.
109                "max-ack-delay": 5000,
110                // This specifies the number of clients which send messages to
111                // the partner but appear to not receive any response.

```

(continues on next page)

(continued from previous page)

```

110         "max-unacked-clients": 5,
111         // This specifies the maximum timeout (in milliseconds) for the
↪server
112         // to complete sync. If you have a large deployment (high tens or
113         // hundreds of thousands of clients), you may need to increase it
114         // further. The default value is 60000ms (60 seconds).
115         "sync-timeout": 60000,
116         "peers": [
117             // This is the configuration of the primary server.
118             {
119                 "name": "server1",
120                 // Specifies the URL on which the partner's control
121                 // channel can be reached. The Control Agent is required
122                 // to run on the partner's machine with "http-host" and
123                 // "http-port" values set to the corresponding values.
124                 "url": "http://192.168.1.2:8005/",
125                 // The other server is primary. This one must be
126                 // secondary.
127                 "role": "primary"
128             },
129             // This is the configuration of this server instance.
130             {
131                 "name": "server2",
132                 // This specifies the URL of this server instance. The
133                 // Control Agent must run along with this DHCPv4 server
134                 // instance and the "http-host" and "http-port" must be
135                 // set to the corresponding values.
136                 "url": "http://192.168.1.3:8005/",
137                 // This server is secondary. The other one must be
138                 // primary.
139                 "role": "standby"
140             }
141         ]
142     } ]
143 }
144 ],
145 ],
146
147 // This example contains a single subnet declaration.
148 "subnet4": [
149     {
150         // Subnet prefix.
151         "subnet": "192.168.1.0/24",
152
153         // There are no relays in this network, so we need to tell Kea that this
↪subnet
154         // is reachable directly via the specified interface.
155         "interface": "enp0s8",
156
157         // Specify a dynamic address pool.
158         "pools": [
159             {

```

(continues on next page)

(continued from previous page)

```

160         "pool": "192.168.1.100-192.168.1.199"
161     }
162 ],
163
164     // These are options that are subnet specific. In most cases, you need to
165     ↪define at
166     // least routers option, as without this option your clients will not be
167     ↪able to reach
168     // their default gateway and will not have Internet connectivity. If you
169     ↪have many
170     // subnets and they share the same options (e.g. DNS servers typically is
171     ↪the same
172     // everywhere), you may define options at the global scope, so you don't
173     ↪repeat them
174     // for every network.
175     "option-data": [
176     {
177         // For each IPv4 subnet you typically need to specify at least one
178         ↪router.
179         "name": "routers",
180         "data": "192.168.1.1"
181     },
182     {
183         // Using cloudflare or Quad9 is a reasonable option. Change this
184         // to your own DNS servers is you have them. Another popular
185         // choice is 8.8.8.8, owned by Google. Using third party DNS
186         // service raises some privacy concerns.
187         "name": "domain-name-servers",
188         "data": "1.1.1.1,9.9.9.9"
189     }
190 ],
191
192     // Some devices should get a static address. Since the .100 - .199 range is
193     ↪dynamic,
194     // let's use the lower address space for this. There are many ways how
195     ↪reservation
196     // can be defined, but using MAC address (hw-address) is by far the most
197     ↪popular one.
198     // You can use client-id, duid and even custom defined flex-id that may use
199     ↪whatever
200     // parts of the packet you want to use as identifiers. Also, there are many
201     ↪more things
202     // you can specify in addition to just an IP address: extra options, next-
203     ↪server, hostname,
204     // assign device to client classes etc. See the Kea ARM, Section 8.3 for
205     ↪details.
206     // The reservations are subnet specific.
207     "reservations": [
208     {
209         "hw-address": "1a:1b:1c:1d:1e:1f",
210         "ip-address": "192.168.1.10"
211     },

```

(continues on next page)

(continued from previous page)

```

199         {
200             "client-id": "01:11:22:33:44:55:66",
201             "ip-address": "192.168.1.11"
202         }
203     ]
204 }
205 ],
206
207 // Logging configuration starts here.
208 "loggers": [
209     {
210         // This section affects kea-dhcp4, which is the base logger for DHCPv4 component.
211         ↪ It tells
212         // DHCPv4 server to write all log messages (on severity INFO or higher) to a
213         ↪ file. The file
214         // will be rotated once it grows to 2MB and up to 4 files will be kept. The
215         ↪ debuglevel
216         // (range 0 to 99) is used only when logging on DEBUG level.
217         "name": "kea-dhcp4",
218         "output-options": [
219             {
220                 "output": "/var/log/kea/kea-dhcp4.log",
221                 "maxsize": 2048000,
222                 "maxver": 4
223             }
224         ],
225         "severity": "INFO",
226         "debuglevel": 0
227     }
228 ]
229 }
230 }

```

27.2 Template: Secure High Availability Kea DHCP with Multi-Threading

Below are some templates to assist in configuring a secure Kea DHCP server with multi-threading. These templates make the following assumptions:

- The administrator wants to set up High Availability (HA) with multi-threading.
- The machines running Kea with multi-threading have at least four CPU cores.
- The connection to the peer is secured using TLS.

The logical setup consists of two hosts, each running a Kea DHCPv4 server and a Control Agent (CA). In the multi-threading setup, the CA is not required, as the server is using its own dedicated HTTP listener to communicate with the peer. However, the CA can still be used to handle user commands.

```

+-host-1-+      +-host-2-+
|         |      |         |
|  CA  |      |  CA  |      ===== - HTTPS connection

```

(continues on next page)

(continued from previous page)

```
|  #  |      |  #  |
|  #  |      |  #  |      ##### - UNIX socket
|  #  |      |  #  |
| DHCPv4 ===== DHCPv4 |
|-----+      +-----+
```

The CAs on host-1 and host-2 both listen on port 8001, and the server's dedicated HTTP listener uses port 8000. The DHCP servers communicate with each other via the dedicated HTTP listener, which forwards only the lease-update commands to the peer server.

27.2.1 Deployment Considerations

The setup is not expected to scale automatically. This example uses four threads for processing DHCP traffic, four threads for listening and handling HA peer HTTP requests, and four threads for sending lease updates to the HA peer. The thread queue used to store incoming DHCP requests is set to 64, but proper testing and benchmarks are required to determine the appropriate values for best performance on the deployment setup.

In this example, there are two hosts running Kea:

- 192.168.1.2 - primary HA server (active, handles all the traffic)
- 192.168.1.3 - secondary HA server (passive, ready to take over if the primary fails)

The network is 192.168.1.0/24. It is assumed that 192.168.1.1 is the default router.

The whole subnet is split into dynamic pools:

- 192.168.1.100 - 192.168.1.199 - this is the dynamic pool. When new devices appear in the network, they are assigned dynamic addresses from this pool.

To deploy this setup, follow the steps provided in the power user home setup with the following distinctions:

1. Install the CA only if the administrator is planning to manage Kea using the RESTful API. Otherwise, the CA is not required for the High Availability Kea server with multi-threading.
2. Alter the following to match the local setup:
 - The paths to `trust-anchor`, `cert-file`, and `key-file` must be set to the respective values corresponding to the deployment machine.
 - The addressing must be updated, if using something other than 192.168.1.0/24. Make sure the CA port configuration (`http-host` and `http-port` in `kea-ca.conf`) is different from the DHCPv4 server configuration (`url` in `hook-libraries/parameters/high-availability/peers` in `kea-dhcp4.conf`). The CA is used to handle only management commands, as the HA module sends lease updates using the dedicated HTTP listener to the peer.
3. Verify the communication between the HA peers by checking the Kea logs.
4. Verify that communication between the hosts works in the opposite direction as well (host-2 can connect to host-1), by repeating step 3 from host-2 using host-1's IP address and port.
5. Install the CA and DHCPv4 on host-2, as in steps 1 and 2. The config file for the standby server is very similar to the one on the primary server, other than the definition of the `this-server-name` field (and possibly the interface names).

27.2.2 Possible Extensions

This sample configuration is basic but functional. Once it is set up and running, administrators may wish to consider the following changes:

- If using a database, it is also possible to configure TLS for the database backend (for lease, host, configuration backend, or forensic logging). See *Database Connectivity* for more information.

Some tweaking of these templates may be required to match specific system needs: at a minimum, the lines highlighted in yellow must be adjusted to match the actual deployment.

Server1's Control Agent configuration file:

```

1 // This is an example of a configuration for Control-Agent (CA) listening
2 // for incoming HTTPS traffic. This is necessary for handling API commands.
3 // For a High Availability setup with multi-threading enabled the CA is not
4 // needed as the peers communicate using a dedicated HTTP listener.
5
6 // It is expected to run with a standby (the passive) server, which has a very similar
7 // configuration. The only difference is that the location of TLS specific files
8 // depend on the configuration of the particular machine.
9 {
10     "Control-agent":
11     {
12         // We need to specify where the agent should listen to incoming HTTP
13         // queries.
14         "http-host": "192.168.1.2",
15
16         // TLS trust anchor (Certificate Authority). This is a file name or
17         // (for OpenSSL only) a directory path.
18         "trust-anchor": "/usr/lib/kea/CA.pem",
19
20         // TLS server certificate file name.
21         "cert-file": "/usr/lib/kea/ca1_cert.pem",
22
23         // TLS server private key file name.
24         "key-file": "/usr/lib/kea/ca1_key.pem",
25
26         // TLS require client certificates flag.
27         "cert-required": true,
28
29         // This specifies the port CA will listen on.
30         // If enabling HA and multi-threading, the 8000 port is used by the HA
31         // hook library http listener. When using HA hook library with
32         // multi-threading to function, make sure the port used by dedicated
33         // listener is different (e.g. 8001) than the one used by CA. Note
34         // the commands should still be sent via CA. The dedicated listener
35         // is specifically for HA updates only.
36         "http-port": 8001,
37
38         "control-sockets":
39         {
40             // This is how the Agent can communicate with the DHCPv4 server.
41             "dhcp4":
42             {

```

(continues on next page)

(continued from previous page)

```

43         "comment": "socket to DHCPv4 server",
44         "socket-type": "unix",
45         "socket-name": "kea4-ctrl-socket"
46     },
47
48     // Location of the DHCPv6 command channel socket.
49     "dhcp6":
50     {
51         "socket-type": "unix",
52         "socket-name": "kea6-ctrl-socket"
53     },
54
55     // Location of the D2 command channel socket.
56     "d2":
57     {
58         "socket-type": "unix",
59         "socket-name": "kea-ddns-ctrl-socket",
60         "user-context": { "in-use": false }
61     }
62 },
63
64 // Similar to other Kea components, CA also uses logging.
65 "loggers": [
66     {
67         "name": "kea-ctrl-agent",
68         "output-options": [
69             {
70                 "output": "/var/log/kea/kea-ctrl-agent.log",
71
72                 // Several additional parameters are possible in addition
73                 // to the typical output. Flush determines whether logger
74                 // flushes output to a file. Maxsize determines maximum
75                 // filesize before the file is rotated. maxver
76                 // specifies the maximum number of rotated files being
77                 // kept.
78                 "flush": true,
79                 "maxsize": 204800,
80                 "maxver": 4,
81                 // We use pattern to specify custom log message layout
82                 "pattern": "%d{%y.%m.%d %H:%M:%S.%q} %-5p [%c/%i] %m\n"
83             }
84         ],
85         "severity": "INFO",
86         "debuglevel": 0 // debug level only applies when severity is set to
87         ↪ DEBUG.
88     }
89 ]
90 }
```

Server1's DHCPv4 configuration file:

```

1 // This is an example configuration of the Kea DHCPv4 server 1:
2 //
3 // - uses High Availability hook library and Lease Commands hook library
4 //   to enable High Availability function for the DHCP server. This config
5 //   file is for the primary (the active) server.
6 // - uses memfile, which stores lease data in a local CSV file
7 // - it assumes a single /24 addressing over a link that is directly reachable
8 //   (no DHCP relays)
9 // - there is a handful of IP reservations
10 //
11 // It is expected to run with a standby (the passive) server, which has a very similar
12 // configuration. The only difference is that "this-server-name" must be set to "server2
13 //   " on the
14 // other server. Also, the interface configuration and location of TLS specific files
15 // depend on the network settings and configuration of the particular machine.
16
17 {
18   "Dhcp4": {
19     // Add names of your network interfaces to listen on.
20     "interfaces-config": {
21       // The DHCPv4 server listens on this interface. When changing this to
22       // the actual name of your interface, make sure to also update the
23       // interface parameter in the subnet definition below.
24       "interfaces": [ "enp0s8" ]
25     },
26
27     // Control socket is required for communication between the Control
28     // Agent and the DHCP server. High Availability requires Control Agent
29     // to be running because lease updates are sent over the RESTful
30     // API between the HA peers.
31     "control-socket": {
32       "socket-type": "unix",
33       "socket-name": "kea4-ctrl-socket"
34     },
35
36     // Multi-threading parameters.
37     "multi-threading": {
38       // By default, Kea processes packets on multiple threads if the hardware permits.
39       "enable-multi-threading": true,
40
41       // When multi-threading is enabled, Kea will process packets on a
42       // number of multiple threads configurable through this option. The
43       // value must be a positive integer (0 means auto-detect).
44       "thread-pool-size": 4,
45
46       // When multi-threading is enabled, Kea will read packets from the
47       // interface and append a working item to the thread pool. This
48       // option configures the maximum number of items that can be queued.
49       // The value must be a positive integer (0 means unlimited).
50       "packet-queue-size": 64
51     },
52   },

```

(continues on next page)

(continued from previous page)

```

53 // Use Memfile lease database backend to store leases in a CSV file.
54 // Depending on how Kea was compiled, it may also support SQL databases
55 // (MySQL and/or PostgreSQL). Those database backends require more
56 // parameters, like name, host and possibly user and password.
57 // There are dedicated examples for each backend. See Section 7.2.2 "Lease
58 // Storage" for details.
59 "lease-database": {
60     // Memfile is the simplest and easiest backend to use. It's an in-memory
61     // database with data being written to a CSV file. It is very similar to
62     // what ISC DHCP does.
63     "type": "memfile"
64 },
65
66 // Let's configure some global parameters. The home network is not very dynamic
67 // and there's no shortage of addresses, so no need to recycle aggressively.
68 "valid-lifetime": 43200, // leases will be valid for 12h
69 "renew-timer": 21600, // clients should renew every 6h
70 "rebind-timer": 32400, // clients should start looking for other servers after 9h
71
72 // Kea will clean up its database of expired leases once per hour. However, it
73 // will keep the leases in expired state for 2 days. This greatly increases the
74 // chances for returning devices to get the same address again. To guarantee that,
75 // use host reservation.
76 // If both "flush-reclaimed-timer-wait-time" and "hold-reclaimed-time" are
77 // not 0, when the client sends a release message the lease is expired
78 // instead of being deleted from lease storage.
79 "expired-leases-processing": {
80     "reclaim-timer-wait-time": 3600,
81     "hold-reclaimed-time": 172800,
82     "max-reclaim-leases": 0,
83     "max-reclaim-time": 0
84 },
85
86 // HA requires two hook libraries to be loaded: libdhcp_lease_cmds.so and
87 // libdhcp_ha.so. The former handles incoming lease updates from the HA peers.
88 // The latter implements high availability feature for Kea. Note the library name
89 // should be the same, but the path is OS specific.
90 "hooks-libraries": [
91     // The lease_cmds library must be loaded because HA makes use of it to
92     // deliver lease updates to the server as well as synchronize the
93     // lease database after failure.
94     {
95         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_lease_cmds.so"
96     },
97
98     {
99         // The HA hook library should be loaded.
100         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_ha.so",
101         "parameters": {
102             // Each server should have the same HA configuration, except for the
103             // "this-server-name" parameter.

```

(continues on next page)

(continued from previous page)

```

105     "high-availability": [ {
106         // This parameter points to this server instance. The respective
107         // HA peers must have this parameter set to their own names.
108         "this-server-name": "server1",
109         // The HA mode is set to hot-standby. In this mode, the active
↪server handles
110         // all the traffic. The standby takes over if the primary becomes
↪unavailable.
111         "mode": "hot-standby",
112         // Heartbeat is to be sent every 10 seconds if no other control
113         // commands are transmitted.
114         "heartbeat-delay": 10000,
115         // Maximum time for partner's response to a heartbeat, after which
116         // failure detection is started. This is specified in milliseconds.
117         // If we don't hear from the partner in 60 seconds, it's time to
118         // start worrying.
119         "max-response-delay": 60000,
120         // The following parameters control how the server detects the
121         // partner's failure. The ACK delay sets the threshold for the
122         // 'secs' field of the received discovers. This is specified in
123         // milliseconds.
124         "max-ack-delay": 5000,
125         // This specifies the number of clients which send messages to
126         // the partner but appear to not receive any response.
127         "max-unacked-clients": 5,
128         // This specifies the maximum timeout (in milliseconds) for the
↪server
129         // to complete sync. If you have a large deployment (high tens or
130         // hundreds of thousands of clients), you may need to increase it
131         // further. The default value is 60000ms (60 seconds).
132         "sync-timeout": 60000,
133         // Multi-threading parameters.
134         // To not experience performance degradation when the Kea server is
135         // processing packets on multiple threads, the High Availability
↪module
136         // must have multi-threading enabled.
137         "multi-threading": {
138             // Enable High Availability to benefit from multi-threading.
↪Default: true.
139             "enable-multi-threading": true,
140             // When running in MT mode, the dedicated listener is used to
↪handle
141             // lease updates.
142             "http-dedicated-listener": true,
143             // The number of threads used to handle incoming requests.
144             // A value of 0 instructs the server to use the same number of
145             // threads that the Kea core is using for DHCP multi-threading.
146             "http-listener-threads": 0,
147             // The number of threads used to handle outgoing requests.
148             // A value of 0 instructs the server to use the same number of
149             // threads that the Kea core is using for DHCP multi-threading.
150             "http-client-threads": 0

```

(continues on next page)

(continued from previous page)

```

151         },
152         "peers": [
153             // This is the configuration of this server instance.
154             {
155                 "name": "server1",
156                 // This specifies the URL of this server dedicated HTTP
157                 ↪ listener.
158                 // The Control Agent is not needed for the High Availability
159                 // with multi-threading, but if it is used, it must use
160                 // different values for "http-host" and "http-port".
161                 "url": "http://192.168.1.2:8005/",
162                 // Trust anchor aka certificate authority file or directory.
163                 "trust-anchor": "/usr/lib/kea/CA.pem",
164                 // Client certificate file name.
165                 "cert-file": "/usr/lib/kea/server1_cert.pem",
166                 // Private key file name.
167                 "key-file": "/usr/lib/kea/server1_key.pem",
168                 // Client certificates are required and verified.
169                 "require-client-certs": true,
170                 // This server is primary. The other one must be
171                 // secondary.
172                 "role": "primary"
173             },
174             // This is the configuration of the secondary server.
175             {
176                 "name": "server2",
177                 // This specifies the URL of the other server's dedicated
178                 ↪ HTTP listener.
179                 // The Control Agent is not needed for the High Availability
180                 // with multi-threading, but if it is used, it must use
181                 // different values for "http-host" and "http-port".
182                 "url": "http://192.168.1.3:8005/",
183                 // Trust anchor aka certificate authority file or directory.
184                 "trust-anchor": "/usr/lib/kea/CA.pem",
185                 // Client certificate file name.
186                 "cert-file": "/usr/lib/kea/server2_cert.pem",
187                 // Private key file name.
188                 "key-file": "/usr/lib/kea/server2_key.pem",
189                 // Client certificates are required and verified.
190                 "require-client-certs": true,
191                 // The other server is secondary. This one must be
192                 // primary.
193                 "role": "standby"
194             }
195         ]
196     }
197 ],
198
199 // This example contains a single subnet declaration.
200 "subnet4": [

```

(continues on next page)

(continued from previous page)

```

201     {
202         // Subnet prefix.
203         "subnet": "192.168.1.0/24",
204
205         // There are no relays in this network, so we need to tell Kea that this
↪ subnet
206         // is reachable directly via the specified interface.
207         "interface": "enp0s8",
208
209         // Specify a dynamic address pool.
210         "pools": [
211             {
212                 "pool": "192.168.1.100-192.168.1.199"
213             }
214         ]
215     }
216 ],
217
218 // Logging configuration starts here.
219 "loggers": [
220     {
221         // This section affects kea-dhcp4, which is the base logger for DHCPv4 component.
↪ It tells
222         // DHCPv4 server to write all log messages (on severity INFO or higher) to a
↪ file. The file
223         // will be rotated once it grows to 2MB and up to 4 files will be kept. The
↪ debuglevel
224         // (range 0 to 99) is used only when logging on DEBUG level.
225         "name": "kea-dhcp4",
226         "output-options": [
227             {
228                 "output": "/var/log/kea/kea-dhcp4.log",
229                 "maxsize": 2048000,
230                 "maxver": 4
231             }
232         ],
233         "severity": "INFO",
234         "debuglevel": 0
235     }
236 ]
237 }
238 }
```

Server2's Control Agent configuration file:

```

1 // This is an example of a configuration for Control-Agent (CA) listening
2 // for incoming HTTPS traffic. This is necessary for handling API commands.
3 // For a High Availability setup with multi-threading enabled the CA is not
4 // needed as the peers communicate using a dedicated HTTP listener.
5
6 // It is expected to run with a primary (the active) server, which has a very similar
7 // configuration. The only difference is that the location of TLS specific files
```

(continues on next page)

(continued from previous page)

```

8 // depend on the configuration of the particular machine.
9 {
10     "Control-agent":
11     {
12         // We need to specify where the agent should listen to incoming HTTP
13         // queries.
14         "http-host": "192.168.1.3",
15
16         // TLS trust anchor (Certificate Authority). This is a file name or
17         // (for OpenSSL only) a directory path.
18         "trust-anchor": "/usr/lib/kea/CA.pem",
19
20         // TLS server certificate file name.
21         "cert-file": "/usr/lib/kea/ca2_cert.pem",
22
23         // TLS server private key file name.
24         "key-file": "/usr/lib/kea/ca2_key.pem",
25
26         // TLS require client certificates flag.
27         "cert-required": true,
28
29         // This specifies the port CA will listen on.
30         // If enabling HA and multi-threading, the 8000 port is used by the HA
31         // hook library http listener. When using HA hook library with
32         // multi-threading to function, make sure the port used by dedicated
33         // listener is different (e.g. 8001) than the one used by CA. Note
34         // the commands should still be sent via CA. The dedicated listener
35         // is specifically for HA updates only.
36         "http-port": 8001,
37
38         "control-sockets":
39         {
40             // This is how the Agent can communicate with the DHCPv4 server.
41             "dhcp4":
42             {
43                 "comment": "socket to DHCPv4 server",
44                 "socket-type": "unix",
45                 "socket-name": "kea4-ctrl-socket"
46             },
47
48             // Location of the DHCPv6 command channel socket.
49             "dhcp6":
50             {
51                 "socket-type": "unix",
52                 "socket-name": "kea6-ctrl-socket"
53             },
54
55             // Location of the D2 command channel socket.
56             "d2":
57             {
58                 "socket-type": "unix",
59                 "socket-name": "kea-ddns-ctrl-socket",

```

(continues on next page)

(continued from previous page)

```

60     "user-context": { "in-use": false }
61   },
62 },
63
64 // Similar to other Kea components, CA also uses logging.
65 "loggers": [
66   {
67     "name": "kea-ctrl-agent",
68     "output-options": [
69       {
70         "output": "/var/log/kea/kea-ctrl-agent.log",
71
72         // Several additional parameters are possible in addition
73         // to the typical output. Flush determines whether logger
74         // flushes output to a file. Maxsize determines maximum
75         // filesize before the file is rotated. maxver
76         // specifies the maximum number of rotated files being
77         // kept.
78         "flush": true,
79         "maxsize": 204800,
80         "maxver": 4,
81         // We use pattern to specify custom log message layout
82         "pattern": "%d{%y.%m.%d %H:%M:%S.%q} %-5p [%c/%i] %m\n"
83       }
84     ],
85     "severity": "INFO",
86     "debuglevel": 0 // debug level only applies when severity is set to
87     ↪ DEBUG.
88   }
89 ]
90 }

```

Server2's DHCPv4 configuration file:

```

1 // This is an example configuration of the Kea DHCPv4 server 2:
2 //
3 // - uses High Availability hook library and Lease Commands hook library
4 //   to enable High Availability function for the DHCP server. This config
5 //   file is for the secondary (the standby) server.
6 // - uses memfile, which stores lease data in a local CSV file
7 // - it assumes a single /24 addressing over a link that is directly reachable
8 //   (no DHCP relays)
9 // - there is a handful of IP reservations
10 //
11 // It is expected to run with a primary (the active) server, which has a very similar
12 // configuration. The only difference is that "this-server-name" must be set to "server2
13 ↪" on the
14 // other server. Also, the interface configuration and location of TLS specific files
15 // depend on the network settings and configuration of the particular machine.
16 {

```

(continues on next page)

(continued from previous page)

```

17
18 "Dhcp4": {
19
20     // Add names of your network interfaces to listen on.
21     "interfaces-config": {
22         // The DHCPv4 server listens on this interface. When changing this to
23         // the actual name of your interface, make sure to also update the
24         // interface parameter in the subnet definition below.
25         "interfaces": [ "enp0s8" ]
26     },
27
28     // Control socket is required for communication between the Control
29     // Agent and the DHCP server. High Availability requires Control Agent
30     // to be running because lease updates are sent over the RESTful
31     // API between the HA peers.
32     "control-socket": {
33         "socket-type": "unix",
34         "socket-name": "kea4-ctrl-socket"
35     },
36
37     // Multi-threading parameters.
38     "multi-threading": {
39         // By default, Kea processes packets on multiple threads if the hardware permits.
40         "enable-multi-threading": true,
41
42         // When multi-threading is enabled, Kea will process packets on a
43         // number of multiple threads configurable through this option. The
44         // value must be a positive integer (0 means auto-detect).
45         "thread-pool-size": 4,
46
47         // When multi-threading is enabled, Kea will read packets from the
48         // interface and append a working item to the thread pool. This
49         // option configures the maximum number of items that can be queued.
50         // The value must be a positive integer (0 means unlimited).
51         "packet-queue-size": 64
52     },
53
54     // Use Memfile lease database backend to store leases in a CSV file.
55     // Depending on how Kea was compiled, it may also support SQL databases
56     // (MySQL and/or PostgreSQL). Those database backends require more
57     // parameters, like name, host and possibly user and password.
58     // There are dedicated examples for each backend. See Section 7.2.2 "Lease
59     // Storage" for details.
60     "lease-database": {
61         // Memfile is the simplest and easiest backend to use. It's an in-memory
62         // database with data being written to a CSV file. It is very similar to
63         // what ISC DHCP does.
64         "type": "memfile"
65     },
66
67     // Let's configure some global parameters. The home network is not very dynamic
68     // and there's no shortage of addresses, so no need to recycle aggressively.

```

(continues on next page)

(continued from previous page)

```

69 "valid-lifetime": 43200, // leases will be valid for 12h
70 "renew-timer": 21600, // clients should renew every 6h
71 "rebind-timer": 32400, // clients should start looking for other servers after 9h
72
73 // Kea will clean up its database of expired leases once per hour. However, it
74 // will keep the leases in expired state for 2 days. This greatly increases the
75 // chances for returning devices to get the same address again. To guarantee that,
76 // use host reservation.
77 // If both "flush-reclaimed-timer-wait-time" and "hold-reclaimed-time" are
78 // not 0, when the client sends a release message the lease is expired
79 // instead of being deleted from lease storage.
80 "expired-leases-processing": {
81     "reclaim-timer-wait-time": 3600,
82     "hold-reclaimed-time": 172800,
83     "max-reclaim-leases": 0,
84     "max-reclaim-time": 0
85 },
86
87 // HA requires two hook libraries to be loaded: libdhcp_lease_cmds.so and
88 // libdhcp_ha.so. The former handles incoming lease updates from the HA peers.
89 // The latter implements high availability feature for Kea. Note the library name
90 // should be the same, but the path is OS specific.
91 "hooks-libraries": [
92     // The lease_cmds library must be loaded because HA makes use of it to
93     // deliver lease updates to the server as well as synchronize the
94     // lease database after failure.
95     {
96         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_lease_cmds.so"
97     },
98
99     {
100         // The HA hook library should be loaded.
101         "library": "/usr/lib/x86_64-linux-gnu/kea/hooks/libdhcp_ha.so",
102         "parameters": {
103             // Each server should have the same HA configuration, except for the
104             // "this-server-name" parameter.
105             "high-availability": [ {
106                 // This parameter points to this server instance. The respective
107                 // HA peers must have this parameter set to their own names.
108                 "this-server-name": "server2",
109                 // The HA mode is set to hot-standby. In this mode, the active_
110                 ↪ server handles
111                 // all the traffic. The standby takes over if the primary becomes_
112                 ↪ unavailable.
113                 "mode": "hot-standby",
114                 // Heartbeat is to be sent every 10 seconds if no other control
115                 // commands are transmitted.
116                 "heartbeat-delay": 10000,
117                 // Maximum time for partner's response to a heartbeat, after which
118                 // failure detection is started. This is specified in milliseconds.
119                 // If we don't hear from the partner in 60 seconds, it's time to
120                 // start worrying.

```

(continues on next page)

(continued from previous page)

```

119         "max-response-delay": 60000,
120         // The following parameters control how the server detects the
121         // partner's failure. The ACK delay sets the threshold for the
122         // 'secs' field of the received discovers. This is specified in
123         // milliseconds.
124         "max-ack-delay": 5000,
125         // This specifies the number of clients which send messages to
126         // the partner but appear to not receive any response.
127         "max-unacked-clients": 5,
128         // This specifies the maximum timeout (in milliseconds) for the
129         ↪server
130         // to complete sync. If you have a large deployment (high tens or
131         // hundreds of thousands of clients), you may need to increase it
132         // further. The default value is 60000ms (60 seconds).
133         "sync-timeout": 60000,
134         // Multi-threading parameters.
135         // To not experience performance degradation when the Kea server is
136         // processing packets on multiple threads, the High Availability
137         ↪module
138         // must have multi-threading enabled.
139         "multi-threading": {
140             // Enable High Availability to benefit from multi-threading.
141         ↪Default: true.
142             "enable-multi-threading": true,
143             // When running in MT mode, the dedicated listener is used to
144         ↪handle
145             // lease updates.
146             "http-dedicated-listener": true,
147             // The number of threads used to handle incoming requests.
148             // A value of 0 instructs the server to use the same number of
149             // threads that the Kea core is using for DHCP multi-threading.
150             "http-listener-threads": 0,
151             // The number of threads used to handle outgoing requests.
152             // A value of 0 instructs the server to use the same number of
153             // threads that the Kea core is using for DHCP multi-threading.
154             "http-client-threads": 0
155         },
156         "peers": [
157             // This is the configuration of the primary server.
158             {
159                 "name": "server1",
160                 // This specifies the URL of the other server's dedicated
161         ↪HTTP listener.
162                 // The Control Agent is not needed for the High Availability
163                 // with multi-threading, but if it is used, it must use
164                 // different values for "http-host" and "http-port".
165                 "url": "http://192.168.1.2:8005/",
166                 // Trust anchor aka certificate authority file or directory.
167                 "trust-anchor": "/usr/lib/kea/CA.pem",
168                 // Client certificate file name.
169                 "cert-file": "/usr/lib/kea/server1_cert.pem",
170                 // Private key file name.

```

(continues on next page)

(continued from previous page)

```

166         "key-file": "/usr/lib/kea/server1_key.pem",
167         // Client certificates are required and verified.
168         "require-client-certs": true,
169         // The other server is primary. This one must be
170         // secondary.
171         "role": "primary"
172     },
173     // This is the configuration of this server instance.
174     {
175         "name": "server2",
176         // This specifies the URL of this server dedicated HTTP
177         // listener.
178         // The Control Agent is not needed for the High Availability
179         // with multi-threading, but if it is used, it must use
180         // different values for "http-host" and "http-port".
181         "url": "http://192.168.1.3:8005/",
182         // Trust anchor aka certificate authority file or directory.
183         "trust-anchor": "/usr/lib/kea/CA.pem",
184         // Client certificate file name.
185         "cert-file": "/usr/lib/kea/server2_cert.pem",
186         // Private key file name.
187         "key-file": "/usr/lib/kea/server2_key.pem",
188         // Client certificates are required and verified.
189         "require-client-certs": true,
190         // This server is secondary. The other one must be
191         // primary.
192         "role": "standby"
193     }
194 ]
195 }
196 }
197 ],
198
199 // This example contains a single subnet declaration.
200 "subnet4": [
201     {
202         // Subnet prefix.
203         "subnet": "192.168.1.0/24",
204
205         // There are no relays in this network, so we need to tell Kea that this
206         // subnet
207         // is reachable directly via the specified interface.
208         "interface": "enp0s8",
209
210         // Specify a dynamic address pool.
211         "pools": [
212             {
213                 "pool": "192.168.1.100-192.168.1.199"
214             }
215         ]
216     }
217 ]

```

(continues on next page)

(continued from previous page)

```
216 ],
217
218 // Logging configuration starts here.
219 "loggers": [
220 {
221     // This section affects kea-dhcp4, which is the base logger for DHCPv4 component.
222     ↪ It tells
223     // DHCPv4 server to write all log messages (on severity INFO or higher) to a
224     ↪ file. The file
225     // will be rotated once it grows to 2MB and up to 4 files will be kept. The
226     ↪ debuglevel
227     // (range 0 to 99) is used only when logging on DEBUG level.
228     "name": "kea-dhcp4",
229     "output-options": [
230     {
231         "output": "/var/log/kea/kea-dhcp4.log",
232         "maxsize": 2048000,
233         "maxver": 4
234     }
235     ],
236     "severity": "INFO",
237     "debuglevel": 0
238 }
239 ]
240 }
241 }
```

KEA FLOW DIAGRAMS

These flow diagrams describe Kea's DHCPv4 server implementation, and they may be useful for system administrators. To design a configuration that results in clients getting the intended addresses and options, it is important to understand the sequence of request-processing steps. For example, Kea iterates looking for a suitable address, and conditionally accepts the first available address, so the order in which addresses are evaluated matters.

It is also useful to understand Kea's processing logic because there are configuration choices which can make the process far more efficient. Kea is very flexible, so it can be applied to very different use cases and in different environments. In an environment where throughput and efficiency are a priority, the administrator can choose to limit some of the processing steps. For example, it is possible to limit the number of different client identifiers Kea evaluates in looking for a host reservation, or even to skip the step of checking for host reservations.

These diagrams are focused on those aspects of Kea processing that will be most useful to operators. The diagrams illustrate DHCPv4 request processing, but most of the logic applies equally to DHCPv6. Following the title of each diagram is a Kea version number. Kea behavior has evolved over time, and the diagrams document the behavior as of the Kea version indicated. These diagrams are provided in the Kea source tree in UML (source), PNG, and SVG formats.

28.1 Main Loop

The main loop is common to both DHCPv4 and DHCPv6 servers.

28.2 DHCPv4 Packet Processing

DHCPv4 packet processing evaluates the type DHCP message: Discover, Request, Release, Decline, or Inform. This diagram shows the general, high-level flow for processing an inbound client DHCP packet from receipt to the server's response.

28.3 DHCPREQUEST Processing

The following diagrams focus on DHCPREQUEST processing. This chart gives an overview of the process, from subnet selection to checking for host reservations to evaluating client classes. Finally, before acknowledging the lease, the options are evaluated and added to the message.

28.4 DHCPv4 Subnet Selection

Subnet selection is the process of choosing a subnet that is topologically appropriate for the client. When the selected subnet is a member of a shared network, the whole shared network is selected. During subnet selection the client class may be checked more than once while iterating through subnets, to determine whether it is permitted in the selected subnet.




Fig. 1: The DHCP server main loop




Fig. 2: DHCPv4 packet processing




Fig. 3: DHCPREQUEST processing




Fig. 4: DHCPv4 subnet selection

28.5 DHCPv4 Special Case of Double-Booting

After subnet selection and before lease allocation, the DHCPv4 server handles the special case of clients restarting with an image provided by PXE boot or bootp. The Lease Request box is expanded below.




Fig. 5: DHCPv4 lease assignment

28.6 DHCPv4 Lease Allocation

The first diagram below illustrates the details of processing the client request, showing the renewal of an existing lease, the assignment of a reserved lease, and the allocation of an unreserved lease.

The second diagram shows the algorithm used to validate a requested lease or select a new address to offer. The right-




Fig. 6: DHCPREQUEST lease allocation

hand side of the diagram shows how a new address is selected, when a new lease is required and the client has neither a requested address nor a reservation. When a new lease is required and Kea iterates over pools and subnets, it starts with the subnet selected above in the subnet selection process.

Note

Declined addresses are included in the statistic for assigned addresses, so the *assigned + free = total* equation is true.

28.7 Lease States

This diagram illustrates the different lease states, including the `free` one, where no lease object exists.

28.8 Checking for Host Reservations

The allocation engine checks for host reservations after selecting a subnet. The following diagram shows the details of that operation. Subnet selection is based on network topology. Host reservations are primarily for assigning options, and options are evaluated after subnet selection. However, if client classes are added in the host reservation, those are also evaluated against the selected subnet in a further check. Kea includes several options to skip checking for host reservations, which can make this process much more efficient if reservations are not being used.

Note

To find a free lease, the allocation engine begins by evaluating the most recently used subnet. The current subnet depends on the history of prior queries.

28.9 Building the Options List

Before sending a response, options are added:

- evaluate required client classes
- build the configured option list
- append requested options
- append requested vendor options
- append basic options

28.10 How Kea Recognizes the Same Client In Different DHCP Messages

28.11 RADIUS workflows for lease allocation




Fig. 8: Lease states





Fig. 9: Host reservation evaluation

buildCfgOptionList: build configured option list algorithm (Kea 1.8.0)



Append requested options algorithm (Kea 1.8.0)








Fig. 11: The appendRequestedOptions (append requested options) algorithm

Append vendor requested options algorithm (Kea 1.8.0)







KEA CONFIGURATION FILE SYNTAX (BNF)

Kea consists of several daemons, each with its own configuration syntax. The following sections provide a complete syntax of all possible parameters, written in Backus-Naur Form (BNF). See this [Wikipedia article on BNF](#) for more information.

29.1 BNF Grammar for DHCPv4

This grammar is generated from `dhcp4_parser.yy`. See *The DHCPv4 Server* for more details.

```
1  Grammar
2
3  $accept ::= start EOF
4
5  start ::= TOPLEVEL_JSON sub_json
6
7  start ::= TOPLEVEL_DHCP4 syntax_map
8
9  start ::= SUB_DHCP4 sub_dhcp4
10
11 start ::= SUB_INTERFACES4 sub_interfaces4
12
13 start ::= SUB_SUBNET4 sub_subnet4
14
15 start ::= SUB_POOL4 sub_pool4
16
17 start ::= SUB_RESERVATION sub_reservation
18
19 start ::= SUB_OPTION_DEFS sub_option_def_list
20
21 start ::= SUB_OPTION_DEF sub_option_def
22
23 start ::= SUB_OPTION_DATA sub_option_data
24
25 start ::= SUB_HOOKS_LIBRARY sub_hooks_library
26
27 start ::= SUB_DHCP_DDNS sub_dhcp_ddns
28
29 start ::= SUB_CONFIG_CONTROL sub_config_control
30
31 value ::= INTEGER
32         | FLOAT
```

(continues on next page)

(continued from previous page)

```

33     | BOOLEAN
34     | STRING
35     | NULL
36     | map2
37     | list_generic
38
39 sub_json ::= value
40
41 map2 ::= "{" map_content "}"
42
43 map_value ::= map2
44
45 map_content ::=
46     | not_empty_map
47
48 not_empty_map ::= STRING ":" value
49     | not_empty_map "," STRING ":" value
50     | not_empty_map ","
51
52 list_generic ::= "[" list_content "]"
53
54 list_content ::=
55     | not_empty_list
56
57 not_empty_list ::= value
58     | not_empty_list "," value
59     | not_empty_list ","
60
61 list_strings ::= "[" list_strings_content "]"
62
63 list_strings_content ::=
64     | not_empty_list_strings
65
66 not_empty_list_strings ::= STRING
67     | not_empty_list_strings "," STRING
68     | not_empty_list_strings ","
69
70 unknown_map_entry ::= STRING ":"
71
72 syntax_map ::= "{" global_object "}"
73
74 global_object ::= "Dhcp4" ":" "{" global_params "}"
75     | global_object_comma
76
77 global_object_comma ::= global_object ","
78
79 sub_dhcp4 ::= "{" global_params "}"
80
81 global_params ::= global_param
82     | global_params "," global_param
83     | global_params ","
84

```

(continues on next page)

(continued from previous page)

```
85 global_param ::= valid_lifetime
86                | min_valid_lifetime
87                | max_valid_lifetime
88                | renew_timer
89                | rebind_timer
90                | decline_probation_period
91                | subnet4_list
92                | shared_networks
93                | interfaces_config
94                | lease_database
95                | hosts_database
96                | hosts_databases
97                | host_reservation_identifiers
98                | client_classes
99                | option_def_list
100               | option_data_list
101               | hooks_libraries
102               | expired_leases_processing
103               | dhcp4o6_port
104               | control_socket
105               | control_sockets
106               | dhcp_queue_control
107               | dhcp_ddns
108               | echo_client_id
109               | match_client_id
110               | authoritative
111               | next_server
112               | server_hostname
113               | boot_file_name
114               | user_context
115               | comment
116               | sanity_checks
117               | reservations
118               | config_control
119               | server_tag
120               | reservations_global
121               | reservations_in_subnet
122               | reservations_out_of_pool
123               | calculate_tee_times
124               | t1_percent
125               | t2_percent
126               | cache_threshold
127               | cache_max_age
128               | loggers
129               | hostname_char_set
130               | hostname_char_replacement
131               | ddns_send_updates
132               | ddns_override_no_update
133               | ddns_override_client_update
134               | ddns_replace_client_name
135               | ddns_generated_prefix
136               | ddns_qualifying_suffix
```

(continues on next page)

(continued from previous page)

```

137         | ddns_update_on_renew
138         | ddns_use_conflict_resolution
139         | ddns_conflict_resolution_mode
140         | ddns_ttl_percent
141         | ddns_ttl
142         | ddns_ttl_min
143         | ddns_ttl_max
144         | store_extended_info
145         | statistic_default_sample_count
146         | statistic_default_sample_age
147         | dhcp_multi_threading
148         | early_global_reservations_lookup
149         | ip_reservations_unique
150         | reservations_lookup_first
151         | compatibility
152         | parked_packet_limit
153         | allocator
154         | offer_lifetime
155         | stash_agent_options
156         | unknown_map_entry
157
158 valid_lifetime ::= "valid-lifetime" ":" INTEGER
159
160 min_valid_lifetime ::= "min-valid-lifetime" ":" INTEGER
161
162 max_valid_lifetime ::= "max-valid-lifetime" ":" INTEGER
163
164 renew_timer ::= "renew-timer" ":" INTEGER
165
166 rebind_timer ::= "rebind-timer" ":" INTEGER
167
168 calculate_tee_times ::= "calculate-tee-times" ":" BOOLEAN
169
170 t1_percent ::= "t1-percent" ":" FLOAT
171
172 t2_percent ::= "t2-percent" ":" FLOAT
173
174 cache_threshold ::= "cache-threshold" ":" FLOAT
175
176 cache_max_age ::= "cache-max-age" ":" INTEGER
177
178 decline_probation_period ::= "decline-probation-period" ":" INTEGER
179
180 server_tag ::= "server-tag" ":" STRING
181
182 parked_packet_limit ::= "parked-packet-limit" ":" INTEGER
183
184 allocator ::= "allocator" ":" STRING
185
186 echo_client_id ::= "echo-client-id" ":" BOOLEAN
187
188 match_client_id ::= "match-client-id" ":" BOOLEAN

```

(continues on next page)

(continued from previous page)

```

189 authoritative ::= "authoritative" ":" BOOLEAN
190
191 ddns_send_updates ::= "ddns-send-updates" ":" BOOLEAN
192
193 ddns_override_no_update ::= "ddns-override-no-update" ":" BOOLEAN
194
195 ddns_override_client_update ::= "ddns-override-client-update" ":" BOOLEAN
196
197 ddns_replace_client_name ::= "ddns-replace-client-name" ":" ddns_replace_client_name_
198 ↪value
199
200 ddns_replace_client_name_value ::= "when-present"
201                                | "never"
202                                | "always"
203                                | "when-not-present"
204                                | BOOLEAN
205
206 ddns_generated_prefix ::= "ddns-generated-prefix" ":" STRING
207
208 ddns_qualifying_suffix ::= "ddns-qualifying-suffix" ":" STRING
209
210 ddns_update_on_renew ::= "ddns-update-on-renew" ":" BOOLEAN
211
212 ddns_use_conflict_resolution ::= "ddns-use-conflict-resolution" ":" BOOLEAN
213
214 ddns_conflict_resolution_mode ::= "ddns-conflict-resolution-mode" ":" ddns_conflict_
215 ↪resolution_mode_value
216
217 ddns_conflict_resolution_mode_value ::= "check-with-dhcid"
218                                     | "no-check-with-dhcid"
219                                     | "check-exists-with-dhcid"
220                                     | "no-check-without-dhcid"
221
222 ddns_ttl_percent ::= "ddns-ttl-percent" ":" FLOAT
223
224 ddns_ttl ::= "ddns-ttl" ":" INTEGER
225
226 ddns_ttl_min ::= "ddns-ttl-min" ":" INTEGER
227
228 ddns_ttl_max ::= "ddns-ttl-mix" ":" INTEGER
229
230 hostname_char_set ::= "hostname-char-set" ":" STRING
231
232 hostname_char_replacement ::= "hostname-char-replacement" ":" STRING
233
234 store_extended_info ::= "store-extended-info" ":" BOOLEAN
235
236 statistic_default_sample_count ::= "statistic-default-sample-count" ":" INTEGER
237
238 statistic_default_sample_age ::= "statistic-default-sample-age" ":" INTEGER

```

(continues on next page)

(continued from previous page)

```

239 early_global_reservations_lookup ::= "early-global-reservations-lookup" ":" BOOLEAN
240
241 ip_reservations_unique ::= "ip-reservations-unique" ":" BOOLEAN
242
243 reservations_lookup_first ::= "reservations-lookup-first" ":" BOOLEAN
244
245 offer_lifetime ::= "offer-lifetime" ":" INTEGER
246
247 stash_agent_options ::= "stash-agent-options" ":" BOOLEAN
248
249 interfaces_config ::= "interfaces-config" ":" "{" interfaces_config_params "}"
250
251 interfaces_config_params ::= interfaces_config_param
252                             | interfaces_config_params "," interfaces_config_param
253                             | interfaces_config_params ","
254
255 interfaces_config_param ::= interfaces_list
256                           | dhcp_socket_type
257                           | outbound_interface
258                           | re_detect
259                           | service_sockets_require_all
260                           | service_sockets_retry_wait_time
261                           | service_sockets_max_retries
262                           | user_context
263                           | comment
264                           | unknown_map_entry
265
266 sub_interfaces4 ::= "{" interfaces_config_params "}"
267
268 interfaces_list ::= "interfaces" ":" list_strings
269
270 dhcp_socket_type ::= "dhcp-socket-type" ":" socket_type
271
272 socket_type ::= "raw"
273               | "udp"
274
275 outbound_interface ::= "outbound-interface" ":" outbound_interface_value
276
277 outbound_interface_value ::= "same-as-inbound"
278                             | "use-routing"
279
280 re_detect ::= "re-detect" ":" BOOLEAN
281
282 service_sockets_require_all ::= "service-sockets-require-all" ":" BOOLEAN
283
284 service_sockets_retry_wait_time ::= "service-sockets-retry-wait-time" ":" INTEGER
285
286 service_sockets_max_retries ::= "service-sockets-max-retries" ":" INTEGER
287
288 lease_database ::= "lease-database" ":" "{" database_map_params "}"
289
290 sanity_checks ::= "sanity-checks" ":" "{" sanity_checks_params "}"

```

(continues on next page)

(continued from previous page)

```

291 sanity_checks_params ::= sanity_checks_param
292                         | sanity_checks_params "," sanity_checks_param
293                         | sanity_checks_params ","
294
295
296 sanity_checks_param ::= lease_checks
297                      | extended_info_checks
298
299
300 lease_checks ::= "lease-checks" ":" STRING
301
302 extended_info_checks ::= "extended-info-checks" ":" STRING
303
304 hosts_database ::= "hosts-database" ":" "{" database_map_params "}"
305
306 hosts_databases ::= "hosts-databases" ":" "[" database_list "]"
307
308 database_list ::=
309                | not_empty_database_list
310
311 not_empty_database_list ::= database
312                          | not_empty_database_list "," database
313                          | not_empty_database_list ","
314
315 database ::= "{" database_map_params "}"
316
317 database_map_params ::= database_map_param
318                    | database_map_params "," database_map_param
319                    | database_map_params ","
320
321 database_map_param ::= database_type
322                   | user
323                   | password
324                   | host
325                   | port
326                   | name
327                   | persist
328                   | lfc_interval
329                   | readonly
330                   | connect_timeout
331                   | read_timeout
332                   | write_timeout
333                   | tcp_user_timeout
334                   | max_reconnect_tries
335                   | reconnect_wait_time
336                   | on_fail
337                   | retry_on_startup
338                   | max_row_errors
339                   | trust_anchor
340                   | cert_file
341                   | key_file
342                   | cipher_list
343                   | unknown_map_entry

```

(continues on next page)

(continued from previous page)

```

343 database_type ::= "type" ":" STRING
344
345 user ::= "user" ":" STRING
346
347 password ::= "password" ":" STRING
348
349 host ::= "host" ":" STRING
350
351 port ::= "port" ":" INTEGER
352
353 name ::= "name" ":" STRING
354
355 persist ::= "persist" ":" BOOLEAN
356
357 lfc_interval ::= "lfc-interval" ":" INTEGER
358
359 readonly ::= "readonly" ":" BOOLEAN
360
361 connect_timeout ::= "connect-timeout" ":" INTEGER
362
363 read_timeout ::= "read-timeout" ":" INTEGER
364
365 write_timeout ::= "write-timeout" ":" INTEGER
366
367 tcp_user_timeout ::= "tcp-user-timeout" ":" INTEGER
368
369 max_reconnect_tries ::= "max-reconnect-tries" ":" INTEGER
370
371 reconnect_wait_time ::= "reconnect-wait-time" ":" INTEGER
372
373 on_fail ::= "on-fail" ":" on_fail_mode
374
375 on_fail_mode ::= "stop-retry-exit"
376                | "serve-retry-exit"
377                | "serve-retry-continue"
378
379 retry_on_startup ::= "retry-on-startup" ":" BOOLEAN
380
381 max_row_errors ::= "max-row-errors" ":" INTEGER
382
383 trust_anchor ::= "trust-anchor" ":" STRING
384
385 cert_file ::= "cert-file" ":" STRING
386
387 key_file ::= "key-file" ":" STRING
388
389 cipher_list ::= "cipher-list" ":" STRING
390
391 host_reservation_identifiers ::= "host-reservation-identifiers" ":" "[" host_
392 ↪ reservation_identifiers_list "]"
393

```

(continues on next page)

(continued from previous page)

```

394 host_reservation_identifiers_list ::= host_reservation_identifier
395                                     | host_reservation_identifiers_list "," host_
↪ reservation_identifier
396                                     | host_reservation_identifiers_list ","
397
398 host_reservation_identifier ::= duid_id
399                             | hw_address_id
400                             | circuit_id
401                             | client_id
402                             | flex_id
403
404 duid_id ::= "duid"
405
406 hw_address_id ::= "hw-address"
407
408 circuit_id ::= "circuit-id"
409
410 client_id ::= "client-id"
411
412 flex_id ::= "flex-id"
413
414 dhcp_multi_threading ::= "multi-threading" ":" "{" multi_threading_params "}"
415
416 multi_threading_params ::= multi_threading_param
417                         | multi_threading_params "," multi_threading_param
418                         | multi_threading_params ","
419
420 multi_threading_param ::= enable_multi_threading
421                       | thread_pool_size
422                       | packet_queue_size
423                       | user_context
424                       | comment
425                       | unknown_map_entry
426
427 enable_multi_threading ::= "enable-multi-threading" ":" BOOLEAN
428
429 thread_pool_size ::= "thread-pool-size" ":" INTEGER
430
431 packet_queue_size ::= "packet-queue-size" ":" INTEGER
432
433 hooks_libraries ::= "hooks-libraries" ":" "[" hooks_libraries_list "]"
434
435 hooks_libraries_list ::=
436                       | not_empty_hooks_libraries_list
437
438 not_empty_hooks_libraries_list ::= hooks_library
439                                | not_empty_hooks_libraries_list "," hooks_library
440                                | not_empty_hooks_libraries_list ","
441
442 hooks_library ::= "{" hooks_params "}"
443
444 sub_hooks_library ::= "{" hooks_params "}"

```

(continues on next page)

(continued from previous page)

```

445 hooks_params ::= hooks_param
446             | hooks_params "," hooks_param
447             | hooks_params ","
448             | unknown_map_entry
449
450
451 hooks_param ::= library
452             | parameters
453
454 library ::= "library" ":" STRING
455
456 parameters ::= "parameters" ":" map_value
457
458 expired_leases_processing ::= "expired-leases-processing" ":" "{" expired_leases_
459 ↪ params "}"
460
461 expired_leases_params ::= expired_leases_param
462                       | expired_leases_params "," expired_leases_param
463                       | expired_leases_params ","
464
465 expired_leases_param ::= reclaim_timer_wait_time
466                     | flush_reclaimed_timer_wait_time
467                     | hold_reclaimed_time
468                     | max_reclaim_leases
469                     | max_reclaim_time
470                     | unwarned_reclaim_cycles
471
472 reclaim_timer_wait_time ::= "reclaim-timer-wait-time" ":" INTEGER
473
474 flush_reclaimed_timer_wait_time ::= "flush-reclaimed-timer-wait-time" ":" INTEGER
475
476 hold_reclaimed_time ::= "hold-reclaimed-time" ":" INTEGER
477
478 max_reclaim_leases ::= "max-reclaim-leases" ":" INTEGER
479
480 max_reclaim_time ::= "max-reclaim-time" ":" INTEGER
481
482 unwarned_reclaim_cycles ::= "unwarned-reclaim-cycles" ":" INTEGER
483
484 subnet4_list ::= "subnet4" ":" "[" subnet4_list_content "]"
485
486 subnet4_list_content ::=
487                     | not_empty_subnet4_list
488
489 not_empty_subnet4_list ::= subnet4
490                       | not_empty_subnet4_list "," subnet4
491                       | not_empty_subnet4_list ","
492
493 subnet4 ::= "{" subnet4_params "}"
494
495 sub_subnet4 ::= "{" subnet4_params "}"

```

(continues on next page)

(continued from previous page)

```

496 subnet4_params ::= subnet4_param
497                 | subnet4_params "," subnet4_param
498                 | subnet4_params ","
499
500 subnet4_param ::= valid_lifetime
501                | min_valid_lifetime
502                | max_valid_lifetime
503                | renew_timer
504                | rebind_timer
505                | option_data_list
506                | pools_list
507                | subnet
508                | interface
509                | id
510                | client_class
511                | network_client_classes
512                | require_client_classes
513                | evaluate_additional_classes
514                | reservations
515                | reservations_global
516                | reservations_in_subnet
517                | reservations_out_of_pool
518                | relay
519                | match_client_id
520                | authoritative
521                | next_server
522                | server_hostname
523                | boot_file_name
524                | subnet_4o6_interface
525                | subnet_4o6_interface_id
526                | subnet_4o6_subnet
527                | user_context
528                | comment
529                | calculate_tee_times
530                | t1_percent
531                | t2_percent
532                | cache_threshold
533                | cache_max_age
534                | ddns_send_updates
535                | ddns_override_no_update
536                | ddns_override_client_update
537                | ddns_replace_client_name
538                | ddns_generated_prefix
539                | ddns_qualifying_suffix
540                | ddns_update_on_renew
541                | ddns_use_conflict_resolution
542                | ddns_conflict_resolution_mode
543                | ddns_ttl_percent
544                | ddns_ttl
545                | ddns_ttl_min
546                | ddns_ttl_max
547                | hostname_char_set

```

(continues on next page)

(continued from previous page)

```

548         | hostname_char_replacement
549         | store_extended_info
550         | allocator
551         | offer_lifetime
552         | unknown_map_entry
553
554 subnet ::= "subnet" ":" STRING
555
556 subnet_4o6_interface ::= "4o6-interface" ":" STRING
557
558 subnet_4o6_interface_id ::= "4o6-interface-id" ":" STRING
559
560 subnet_4o6_subnet ::= "4o6-subnet" ":" STRING
561
562 interface ::= "interface" ":" STRING
563
564 client_class ::= "client-class" ":" STRING
565
566 network_client_classes ::= "client-classes" ":" list_strings
567
568 require_client_classes ::= "require-client-classes" ":" list_strings
569
570 evaluate_additional_classes ::= "evaluate-additional-classes" ":" list_strings
571
572 reservations_global ::= "reservations-global" ":" BOOLEAN
573
574 reservations_in_subnet ::= "reservations-in-subnet" ":" BOOLEAN
575
576 reservations_out_of_pool ::= "reservations-out-of-pool" ":" BOOLEAN
577
578 id ::= "id" ":" INTEGER
579
580 shared_networks ::= "shared-networks" ":" "[" shared_networks_content "]"
581
582 shared_networks_content ::=
583         | shared_networks_list
584
585 shared_networks_list ::= shared_network
586         | shared_networks_list "," shared_network
587         | shared_networks_list ","
588
589 shared_network ::= "{" shared_network_params "}"
590
591 shared_network_params ::= shared_network_param
592         | shared_network_params "," shared_network_param
593         | shared_network_params ","
594
595 shared_network_param ::= name
596         | subnet4_list
597         | interface
598         | renew_timer
599         | rebind_timer

```

(continues on next page)

(continued from previous page)

```

600         | option_data_list
601         | match_client_id
602         | authoritative
603         | next_server
604         | server_hostname
605         | boot_file_name
606         | relay
607         | reservations_global
608         | reservations_in_subnet
609         | reservations_out_of_pool
610         | client_class
611         | network_client_classes
612         | require_client_classes
613         | evaluate_additional_classes
614         | valid_lifetime
615         | min_valid_lifetime
616         | max_valid_lifetime
617         | user_context
618         | comment
619         | calculate_tee_times
620         | t1_percent
621         | t2_percent
622         | cache_threshold
623         | cache_max_age
624         | ddns_send_updates
625         | ddns_override_no_update
626         | ddns_override_client_update
627         | ddns_replace_client_name
628         | ddns_generated_prefix
629         | ddns_qualifying_suffix
630         | ddns_update_on_renew
631         | ddns_use_conflict_resolution
632         | ddns_conflict_resolution_mode
633         | ddns_ttl_percent
634         | ddns_ttl
635         | ddns_ttl_min
636         | ddns_ttl_max
637         | hostname_char_set
638         | hostname_char_replacement
639         | store_extended_info
640         | allocator
641         | offer_lifetime
642         | unknown_map_entry
643
644 option_def_list ::= "option-def" ":" "[" option_def_list_content "]"
645
646 sub_option_def_list ::= "{" option_def_list "}"
647
648 option_def_list_content ::=
649         | not_empty_option_def_list
650
651 not_empty_option_def_list ::= option_def_entry

```

(continues on next page)

(continued from previous page)

```

652         | not_empty_option_def_list "," option_def_entry
653         | not_empty_option_def_list ","
654
655 option_def_entry ::= "{" option_def_params "}"
656
657 sub_option_def ::= "{" option_def_params "}"
658
659 option_def_params ::=
660     | not_empty_option_def_params
661
662 not_empty_option_def_params ::= option_def_param
663     | not_empty_option_def_params "," option_def_param
664     | not_empty_option_def_params ","
665
666 option_def_param ::= option_def_name
667     | option_def_code
668     | option_def_type
669     | option_def_record_types
670     | option_def_space
671     | option_def_encapsulate
672     | option_def_array
673     | user_context
674     | comment
675     | unknown_map_entry
676
677 option_def_name ::= name
678
679 code ::= "code" ":" INTEGER
680
681 option_def_code ::= code
682
683 option_def_type ::= "type" ":" STRING
684
685 option_def_record_types ::= "record-types" ":" STRING
686
687 space ::= "space" ":" STRING
688
689 option_def_space ::= space
690
691 option_def_encapsulate ::= "encapsulate" ":" STRING
692
693 option_def_array ::= "array" ":" BOOLEAN
694
695 option_data_list ::= "option-data" ":" "[" option_data_list_content "]"
696
697 option_data_list_content ::=
698     | not_empty_option_data_list
699
700 not_empty_option_data_list ::= option_data_entry
701     | not_empty_option_data_list "," option_data_entry
702     | not_empty_option_data_list ","
703

```

(continues on next page)

(continued from previous page)

```

704 option_data_entry ::= "{" option_data_params "}"
705
706 sub_option_data ::= "{" option_data_params "}"
707
708 option_data_params ::=
709     | not_empty_option_data_params
710
711 not_empty_option_data_params ::= option_data_param
712     | not_empty_option_data_params "," option_data_param
713     | not_empty_option_data_params ","
714
715 option_data_param ::= option_data_name
716     | option_data_data
717     | option_data_code
718     | option_data_space
719     | option_data_csv_format
720     | option_data_always_send
721     | option_data_never_send
722     | user_context
723     | comment
724     | option_data_client_classes
725     | unknown_map_entry
726
727 option_data_name ::= name
728
729 option_data_data ::= "data" ":" STRING
730
731 option_data_code ::= code
732
733 option_data_space ::= space
734
735 option_data_csv_format ::= "csv-format" ":" BOOLEAN
736
737 option_data_always_send ::= "always-send" ":" BOOLEAN
738
739 option_data_never_send ::= "never-send" ":" BOOLEAN
740
741 option_data_client_classes ::= "client-classes" ":" list_strings
742
743 pools_list ::= "pools" ":" "[" pools_list_content "]"
744
745 pools_list_content ::=
746     | not_empty_pools_list
747
748 not_empty_pools_list ::= pool_list_entry
749     | not_empty_pools_list "," pool_list_entry
750     | not_empty_pools_list ","
751
752 pool_list_entry ::= "{" pool_params "}"
753
754 sub_pool4 ::= "{" pool_params "}"
755

```

(continues on next page)

(continued from previous page)

```

756 pool_params ::= pool_param
757             | pool_params "," pool_param
758             | pool_params ","
759
760 pool_param ::= pool_entry
761            | pool_id
762            | option_data_list
763            | client_class
764            | network_client_classes
765            | require_client_classes
766            | evaluate_additional_classes
767            | ddns_send_updates
768            | ddns_override_no_update
769            | ddns_override_client_update
770            | ddns_replace_client_name
771            | ddns_generated_prefix
772            | ddns_qualifying_suffix
773            | ddns_update_on_renew
774            | ddns_conflict_resolution_mode
775            | ddns_ttl_percent
776            | ddns_ttl
777            | ddns_ttl_min
778            | ddns_ttl_max
779            | hostname_char_set
780            | hostname_char_replacement
781            | user_context
782            | comment
783            | unknown_map_entry
784
785 pool_entry ::= "pool" ":" STRING
786
787 pool_id ::= "pool-id" ":" INTEGER
788
789 user_context ::= "user-context" ":" map_value
790
791 comment ::= "comment" ":" STRING
792
793 reservations ::= "reservations" ":" "[" reservations_list "]"
794
795 reservations_list ::=
796                 | not_empty_reservations_list
797
798 not_empty_reservations_list ::= reservation
799                             | not_empty_reservations_list "," reservation
800                             | not_empty_reservations_list ","
801
802 reservation ::= "{" reservation_params "}"
803
804 sub_reservation ::= "{" reservation_params "}"
805
806 reservation_params ::=
807                 | not_empty_reservation_params

```

(continues on next page)

(continued from previous page)

```

808
809 not_empty_reservation_params ::= reservation_param
810                               | not_empty_reservation_params "," reservation_param
811                               | not_empty_reservation_params ","
812
813 reservation_param ::= duid
814                   | reservation_client_classes
815                   | client_id_value
816                   | circuit_id_value
817                   | flex_id_value
818                   | ip_address
819                   | hw_address
820                   | hostname
821                   | option_data_list
822                   | next_server
823                   | server_hostname
824                   | boot_file_name
825                   | user_context
826                   | comment
827                   | unknown_map_entry
828
829 next_server ::= "next-server" ":" STRING
830
831 server_hostname ::= "server-hostname" ":" STRING
832
833 boot_file_name ::= "boot-file-name" ":" STRING
834
835 ip_address ::= "ip-address" ":" STRING
836
837 duid ::= "duid" ":" STRING
838
839 hw_address ::= "hw-address" ":" STRING
840
841 client_id_value ::= "client-id" ":" STRING
842
843 circuit_id_value ::= "circuit-id" ":" STRING
844
845 flex_id_value ::= "flex-id" ":" STRING
846
847 hostname ::= "hostname" ":" STRING
848
849 reservation_client_classes ::= "client-classes" ":" list_strings
850
851 relay ::= "relay" ":" "{" relay_map "}"
852
853 relay_map ::= ip_addresses
854
855 ip_addresses ::= "ip-addresses" ":" list_strings
856
857 client_classes ::= "client-classes" ":" "[" client_classes_list "]"
858
859 client_classes_list ::= client_class_entry

```

(continues on next page)

(continued from previous page)

```

860         | client_classes_list "," client_class_entry
861         | client_classes_list ","
862
863 client_class_entry ::= "{" client_class_params "}"
864
865 client_class_params ::=
866         | not_empty_client_class_params
867
868 not_empty_client_class_params ::= client_class_param
869         | not_empty_client_class_params "," client_class_param
870         | not_empty_client_class_params ","
871
872 client_class_param ::= client_class_name
873         | client_class_test
874         | client_class_template_test
875         | only_if_required
876         | only_in_additional_list
877         | option_def_list
878         | option_data_list
879         | next_server
880         | server_hostname
881         | boot_file_name
882         | user_context
883         | comment
884         | unknown_map_entry
885         | valid_lifetime
886         | min_valid_lifetime
887         | max_valid_lifetime
888         | offer_lifetime
889
890 client_class_name ::= name
891
892 client_class_test ::= "test" ":" STRING
893
894 client_class_template_test ::= "template-test" ":" STRING
895
896 only_if_required ::= "only-if-required" ":" BOOLEAN
897
898 only_in_additional_list ::= "only-in-additional-list" ":" BOOLEAN
899
900 dhcp4o6_port ::= "dhcp4o6-port" ":" INTEGER
901
902 control_socket ::= "control-socket" ":" "{" control_socket_params "}"
903
904 control_sockets ::= "control-sockets" ":" "[" control_socket_list "]"
905
906 control_socket_list ::=
907         | not_empty_control_socket_list
908
909 not_empty_control_socket_list ::= control_socket_entry
910         | not_empty_control_socket_list "," control_socket_entry
911         | not_empty_control_socket_list ","

```

(continues on next page)

(continued from previous page)

```

912 control_socket_entry ::= "{" control_socket_params "}"
913
914 control_socket_params ::= control_socket_param
915                          | control_socket_params "," control_socket_param
916                          | control_socket_params ","
917
918 control_socket_param ::= control_socket_type
919                        | control_socket_name
920                        | control_socket_address
921                        | control_socket_port
922                        | authentication
923                        | trust_anchor
924                        | cert_file
925                        | key_file
926                        | cert_required
927                        | http_headers
928                        | user_context
929                        | comment
930                        | unknown_map_entry
931
932 control_socket_type ::= "socket-type" ":" control_socket_type_value
933
934 control_socket_type_value ::= "unix"
935                             | "http"
936                             | "https"
937
938 control_socket_name ::= "socket-name" ":" STRING
939
940 control_socket_address ::= "socket-address" ":" STRING
941
942 control_socket_port ::= "socket-port" ":" INTEGER
943
944 cert_required ::= "cert-required" ":" BOOLEAN
945
946 http_headers ::= "http-headers" ":" "[" http_header_list "]"
947
948 http_header_list ::=
949                   | not_empty_http_header_list
950
951 not_empty_http_header_list ::= http_header
952                               | not_empty_http_header_list "," http_header
953                               | not_empty_http_header_list ","
954
955 http_header ::= "{" http_header_params "}"
956
957 http_header_params ::= http_header_param
958                     | http_header_params "," http_header_param
959                     | http_header_params ","
960
961 http_header_param ::= name
962                   | header_value

```

(continues on next page)

(continued from previous page)

```

964         | user_context
965         | comment
966         | unknown_map_entry
967
968 header_value ::= "value" ":" STRING
969
970 authentication ::= "authentication" ":" "{" auth_params "}"
971
972 auth_params ::= auth_param
973             | auth_params "," auth_param
974             | auth_params ","
975
976 auth_param ::= auth_type
977             | realm
978             | directory
979             | clients
980             | comment
981             | user_context
982             | unknown_map_entry
983
984 auth_type ::= "type" ":" auth_type_value
985
986 auth_type_value ::= "basic"
987
988 realm ::= "realm" ":" STRING
989
990 directory ::= "directory" ":" STRING
991
992 clients ::= "clients" ":" "[" clients_list "]"
993
994 clients_list ::=
995             | not_empty_clients_list
996
997 not_empty_clients_list ::= basic_auth
998                         | not_empty_clients_list "," basic_auth
999                         | not_empty_clients_list ","
1000
1001 basic_auth ::= "{" clients_params "}"
1002
1003 clients_params ::= clients_param
1004                 | clients_params "," clients_param
1005                 | clients_params ","
1006
1007 clients_param ::= user
1008                | user_file
1009                | password
1010                | password_file
1011                | user_context
1012                | comment
1013                | unknown_map_entry
1014
1015 user_file ::= "user-file" ":" STRING

```

(continues on next page)

(continued from previous page)

```

1016 password_file ::= "password-file" ":" STRING
1017
1018 dhcp_queue_control ::= "dhcp-queue-control" ":" "{" queue_control_params "}"
1019
1020 queue_control_params ::= queue_control_param
1021                        | queue_control_params "," queue_control_param
1022                        | queue_control_params ","
1023
1024 queue_control_param ::= enable_queue
1025                      | queue_type
1026                      | capacity
1027                      | user_context
1028                      | comment
1029                      | arbitrary_map_entry
1030
1031 enable_queue ::= "enable-queue" ":" BOOLEAN
1032
1033 queue_type ::= "queue-type" ":" STRING
1034
1035 capacity ::= "capacity" ":" INTEGER
1036
1037 arbitrary_map_entry ::= STRING ":" value
1038
1039 dhcp_ddns ::= "dhcp-ddns" ":" "{" dhcp_ddns_params "}"
1040
1041 sub_dhcp_ddns ::= "{" dhcp_ddns_params "}"
1042
1043 dhcp_ddns_params ::= dhcp_ddns_param
1044                   | dhcp_ddns_params "," dhcp_ddns_param
1045                   | dhcp_ddns_params ","
1046
1047 dhcp_ddns_param ::= enable_updates
1048                  | server_ip
1049                  | server_port
1050                  | sender_ip
1051                  | sender_port
1052                  | max_queue_size
1053                  | ncr_protocol
1054                  | ncr_format
1055                  | user_context
1056                  | comment
1057                  | unknown_map_entry
1058
1059 enable_updates ::= "enable-updates" ":" BOOLEAN
1060
1061 server_ip ::= "server-ip" ":" STRING
1062
1063 server_port ::= "server-port" ":" INTEGER
1064
1065 sender_ip ::= "sender-ip" ":" STRING
1066
1067

```

(continues on next page)

(continued from previous page)

```

1068 sender_port ::= "sender-port" ":" INTEGER
1069
1070 max_queue_size ::= "max-queue-size" ":" INTEGER
1071
1072 ncr_protocol ::= "ncr-protocol" ":" ncr_protocol_value
1073
1074 ncr_protocol_value ::= "udp"
1075                        | "tcp"
1076
1077 ncr_format ::= "ncr-format" ":" "JSON"
1078
1079 config_control ::= "config-control" ":" "{" config_control_params "}"
1080
1081 sub_config_control ::= "{" config_control_params "}"
1082
1083 config_control_params ::= config_control_param
1084                        | config_control_params "," config_control_param
1085                        | config_control_params ","
1086
1087 config_control_param ::= config_databases
1088                        | config_fetch_wait_time
1089
1090 config_databases ::= "config-databases" ":" "[" database_list "]"
1091
1092 config_fetch_wait_time ::= "config-fetch-wait-time" ":" INTEGER
1093
1094 loggers ::= "loggers" ":" "[" loggers_entries "]"
1095
1096 loggers_entries ::= logger_entry
1097                  | loggers_entries "," logger_entry
1098                  | loggers_entries ","
1099
1100 logger_entry ::= "{" logger_params "}"
1101
1102 logger_params ::= logger_param
1103                | logger_params "," logger_param
1104                | logger_params ","
1105
1106 logger_param ::= name
1107                | output_options_list
1108                | debuglevel
1109                | severity
1110                | user_context
1111                | comment
1112                | unknown_map_entry
1113
1114 debuglevel ::= "debuglevel" ":" INTEGER
1115
1116 severity ::= "severity" ":" STRING
1117
1118 output_options_list ::= "output-options" ":" "[" output_options_list_content "]"
1119

```

(continues on next page)

(continued from previous page)

```

1120 output_options_list_content ::= output_entry
1121                               | output_options_list_content "," output_entry
1122                               | output_options_list_content ","
1123
1124 output_entry ::= "{" output_params_list "}"
1125
1126 output_params_list ::= output_params
1127                      | output_params_list "," output_params
1128                      | output_params_list ","
1129
1130 output_params ::= output
1131                | flush
1132                | maxsize
1133                | maxver
1134                | pattern
1135
1136 output ::= "output" ":" STRING
1137
1138 flush ::= "flush" ":" BOOLEAN
1139
1140 maxsize ::= "maxsize" ":" INTEGER
1141
1142 maxver ::= "maxver" ":" INTEGER
1143
1144 pattern ::= "pattern" ":" STRING
1145
1146 compatibility ::= "compatibility" ":" "{" compatibility_params "}"
1147
1148 compatibility_params ::= compatibility_param
1149                       | compatibility_params "," compatibility_param
1150                       | compatibility_params ","
1151
1152 compatibility_param ::= lenient_option_parsing
1153                      | ignore_dhcp_server_identifier
1154                      | ignore_rai_link_selection
1155                      | exclude_first_last_24
1156                      | unknown_map_entry
1157
1158 lenient_option_parsing ::= "lenient-option-parsing" ":" BOOLEAN
1159
1160 ignore_dhcp_server_identifier ::= "ignore-dhcp-server-identifier" ":" BOOLEAN
1161
1162 ignore_rai_link_selection ::= "ignore-rai-link-selection" ":" BOOLEAN
1163
1164 exclude_first_last_24 ::= "exclude-first-last-24" ":" BOOLEAN

```

29.2 BNF Grammar for DHCPv6

This grammar is generated from `dhcp6_parser.yy`. See *The DHCPv6 Server* for more details.

```
Grammar

$accept ::= start EOF

start ::= TOPLEVEL_JSON sub_json

start ::= TOPLEVEL_DHCP6 syntax_map

start ::= SUB_DHCP6 sub_dhcp6

start ::= SUB_INTERFACES6 sub_interfaces6

start ::= SUB_SUBNET6 sub_subnet6

start ::= SUB_POOL6 sub_pool6

start ::= SUB_PD_POOL sub_pd_pool

start ::= SUB_RESERVATION sub_reservation

start ::= SUB_OPTION_DEFS sub_option_def_list

start ::= SUB_OPTION_DEF sub_option_def

start ::= SUB_OPTION_DATA sub_option_data

start ::= SUB_HOOKS_LIBRARY sub_hooks_library

start ::= SUB_DHCP_DDNS sub_dhcp_ddns

start ::= SUB_CONFIG_CONTROL sub_config_control

value ::= INTEGER
        | FLOAT
        | BOOLEAN
        | STRING
        | NULL
        | map2
        | list_generic

sub_json ::= value

map2 ::= "{" map_content "}"

map_value ::= map2

map_content ::=
    | not_empty_map

not_empty_map ::= STRING ":" value
                | not_empty_map "," STRING ":" value
                | not_empty_map ","
```

(continues on next page)

(continued from previous page)

```

54 list_generic ::= "[" list_content "]"
55
56 list_content ::=
57     | not_empty_list
58
59 not_empty_list ::= value
60     | not_empty_list "," value
61     | not_empty_list ","
62
63 list_strings ::= "[" list_strings_content "]"
64
65 list_strings_content ::=
66     | not_empty_list_strings
67
68 not_empty_list_strings ::= STRING
69     | not_empty_list_strings "," STRING
70     | not_empty_list_strings ","
71
72 unknown_map_entry ::= STRING ":"
73
74 syntax_map ::= "{" global_object "}"
75
76 global_object ::= "Dhcp6" ":" "{" global_params "}"
77     | global_object_comma
78
79 global_object_comma ::= global_object ","
80
81 sub_dhcp6 ::= "{" global_params "}"
82
83 global_params ::= global_param
84     | global_params "," global_param
85     | global_params ","
86
87 global_param ::= data_directory
88     | preferred_lifetime
89     | min_preferred_lifetime
90     | max_preferred_lifetime
91     | valid_lifetime
92     | min_valid_lifetime
93     | max_valid_lifetime
94     | renew_timer
95     | rebind_timer
96     | decline_probation_period
97     | subnet6_list
98     | shared_networks
99     | interfaces_config
100     | lease_database
101     | hosts_database
102     | hosts_databases
103     | mac_sources
104     | relay_supplied_options
105     | host_reservation_identifiers

```

(continues on next page)

(continued from previous page)

```
106 | client_classes
107 | option_def_list
108 | option_data_list
109 | hooks_libraries
110 | expired_leases_processing
111 | server_id
112 | dhcp4o6_port
113 | control_socket
114 | control_sockets
115 | dhcp_queue_control
116 | dhcp_ddns
117 | user_context
118 | comment
119 | sanity_checks
120 | reservations
121 | config_control
122 | server_tag
123 | reservations_global
124 | reservations_in_subnet
125 | reservations_out_of_pool
126 | calculate_tee_times
127 | t1_percent
128 | t2_percent
129 | cache_threshold
130 | cache_max_age
131 | loggers
132 | hostname_char_set
133 | hostname_char_replacement
134 | ddns_send_updates
135 | ddns_override_no_update
136 | ddns_override_client_update
137 | ddns_replace_client_name
138 | ddns_generated_prefix
139 | ddns_qualifying_suffix
140 | ddns_update_on_renew
141 | ddns_use_conflict_resolution
142 | ddns_conflict_resolution_mode
143 | ddns_ttl_percent
144 | ddns_ttl
145 | ddns_ttl_min
146 | ddns_ttl_max
147 | store_extended_info
148 | statistic_default_sample_count
149 | statistic_default_sample_age
150 | dhcp_multi_threading
151 | early_global_reservations_lookup
152 | ip_reservations_unique
153 | reservations_lookup_first
154 | compatibility
155 | parked_packet_limit
156 | allocator
157 | pd_allocator
```

(continues on next page)

(continued from previous page)

```

158         | unknown_map_entry
159
160 data_directory ::= "data-directory" ":" STRING
161
162 preferred_lifetime ::= "preferred-lifetime" ":" INTEGER
163
164 min_preferred_lifetime ::= "min-preferred-lifetime" ":" INTEGER
165
166 max_preferred_lifetime ::= "max-preferred-lifetime" ":" INTEGER
167
168 valid_lifetime ::= "valid-lifetime" ":" INTEGER
169
170 min_valid_lifetime ::= "min-valid-lifetime" ":" INTEGER
171
172 max_valid_lifetime ::= "max-valid-lifetime" ":" INTEGER
173
174 renew_timer ::= "renew-timer" ":" INTEGER
175
176 rebind_timer ::= "rebind-timer" ":" INTEGER
177
178 calculate_tee_times ::= "calculate-tee-times" ":" BOOLEAN
179
180 t1_percent ::= "t1-percent" ":" FLOAT
181
182 t2_percent ::= "t2-percent" ":" FLOAT
183
184 cache_threshold ::= "cache-threshold" ":" FLOAT
185
186 cache_max_age ::= "cache-max-age" ":" INTEGER
187
188 decline_probation_period ::= "decline-probation-period" ":" INTEGER
189
190 ddns_send_updates ::= "ddns-send-updates" ":" BOOLEAN
191
192 ddns_override_no_update ::= "ddns-override-no-update" ":" BOOLEAN
193
194 ddns_override_client_update ::= "ddns-override-client-update" ":" BOOLEAN
195
196 ddns_replace_client_name ::= "ddns-replace-client-name" ":" ddns_replace_client_name_
↪ value
197
198 ddns_replace_client_name_value ::= "when-present"
199         | "never"
200         | "always"
201         | "when-not-present"
202         | BOOLEAN
203
204 ddns_generated_prefix ::= "ddns-generated-prefix" ":" STRING
205
206 ddns_qualifying_suffix ::= "ddns-qualifying-suffix" ":" STRING
207
208 ddns_update_on_renew ::= "ddns-update-on-renew" ":" BOOLEAN

```

(continues on next page)

(continued from previous page)

```

209 ddns_use_conflict_resolution ::= "ddns-use-conflict-resolution" ":" BOOLEAN
210
211 ddns_conflict_resolution_mode ::= "ddns-conflict-resolution-mode" ":" ddns_conflict_
212 ↪ resolution_mode_value
213
214 ddns_conflict_resolution_mode_value ::= "check-with-dhcid"
215                                     | "no-check-with-dhcid"
216                                     | "check-exists-with-dhcid"
217                                     | "no-check-without-dhcid"
218
219 ddns_ttl_percent ::= "ddns-ttl-percent" ":" FLOAT
220
221 ddns_ttl ::= "ddns-ttl" ":" INTEGER
222
223 ddns_ttl_min ::= "ddns-ttl-min" ":" INTEGER
224
225 ddns_ttl_max ::= "ddns-ttl-mix" ":" INTEGER
226
227 hostname_char_set ::= "hostname-char-set" ":" STRING
228
229 hostname_char_replacement ::= "hostname-char-replacement" ":" STRING
230
231 store_extended_info ::= "store-extended-info" ":" BOOLEAN
232
233 statistic_default_sample_count ::= "statistic-default-sample-count" ":" INTEGER
234
235 statistic_default_sample_age ::= "statistic-default-sample-age" ":" INTEGER
236
237 server_tag ::= "server-tag" ":" STRING
238
239 parked_packet_limit ::= "parked-packet-limit" ":" INTEGER
240
241 allocator ::= "allocator" ":" STRING
242
243 pd_allocator ::= "pd-allocator" ":" STRING
244
245 early_global_reservations_lookup ::= "early-global-reservations-lookup" ":" BOOLEAN
246
247 ip_reservations_unique ::= "ip-reservations-unique" ":" BOOLEAN
248
249 reservations_lookup_first ::= "reservations-lookup-first" ":" BOOLEAN
250
251 interfaces_config ::= "interfaces-config" ":" "{" interfaces_config_params "}"
252
253 sub_interfaces6 ::= "{" interfaces_config_params "}"
254
255 interfaces_config_params ::= interfaces_config_param
256                             | interfaces_config_params "," interfaces_config_param
257                             | interfaces_config_params ","
258
259 interfaces_config_param ::= interfaces_list

```

(continues on next page)

(continued from previous page)

```

260         | re_detect
261         | service_sockets_require_all
262         | service_sockets_retry_wait_time
263         | service_sockets_max_retries
264         | user_context
265         | comment
266         | unknown_map_entry
267
268 interfaces_list ::= "interfaces" ":" list_strings
269
270 re_detect ::= "re-detect" ":" BOOLEAN
271
272 service_sockets_require_all ::= "service-sockets-require-all" ":" BOOLEAN
273
274 service_sockets_retry_wait_time ::= "service-sockets-retry-wait-time" ":" INTEGER
275
276 service_sockets_max_retries ::= "service-sockets-max-retries" ":" INTEGER
277
278 lease_database ::= "lease-database" ":" "{" database_map_params "}"
279
280 hosts_database ::= "hosts-database" ":" "{" database_map_params "}"
281
282 hosts_databases ::= "hosts-databases" ":" "[" database_list "]"
283
284 database_list ::=
285     | not_empty_database_list
286
287 not_empty_database_list ::= database
288     | not_empty_database_list "," database
289     | not_empty_database_list ","
290
291 database ::= "{" database_map_params "}"
292
293 database_map_params ::= database_map_param
294     | database_map_params "," database_map_param
295     | database_map_params ","
296
297 database_map_param ::= database_type
298     | user
299     | password
300     | host
301     | port
302     | name
303     | persist
304     | lfc_interval
305     | readonly
306     | connect_timeout
307     | read_timeout
308     | write_timeout
309     | tcp_user_timeout
310     | max_reconnect_tries
311     | reconnect_wait_time

```

(continues on next page)

(continued from previous page)

```

312         | on_fail
313         | retry_on_startup
314         | max_row_errors
315         | trust_anchor
316         | cert_file
317         | key_file
318         | cipher_list
319         | unknown_map_entry
320
321 database_type ::= "type" ":" STRING
322
323 user ::= "user" ":" STRING
324
325 password ::= "password" ":" STRING
326
327 host ::= "host" ":" STRING
328
329 port ::= "port" ":" INTEGER
330
331 name ::= "name" ":" STRING
332
333 persist ::= "persist" ":" BOOLEAN
334
335 lfc_interval ::= "lfc-interval" ":" INTEGER
336
337 readonly ::= "readonly" ":" BOOLEAN
338
339 connect_timeout ::= "connect-timeout" ":" INTEGER
340
341 read_timeout ::= "read-timeout" ":" INTEGER
342
343 write_timeout ::= "write-timeout" ":" INTEGER
344
345 tcp_user_timeout ::= "tcp-user-timeout" ":" INTEGER
346
347 reconnect_wait_time ::= "reconnect-wait-time" ":" INTEGER
348
349 on_fail ::= "on-fail" ":" on_fail_mode
350
351 on_fail_mode ::= "stop-retry-exit"
352               | "serve-retry-exit"
353               | "serve-retry-continue"
354
355 retry_on_startup ::= "retry-on-startup" ":" BOOLEAN
356
357 max_row_errors ::= "max-row-errors" ":" INTEGER
358
359 max_reconnect_tries ::= "max-reconnect-tries" ":" INTEGER
360
361 trust_anchor ::= "trust-anchor" ":" STRING
362
363 cert_file ::= "cert-file" ":" STRING

```

(continues on next page)

(continued from previous page)

```

364 key_file ::= "key-file" ":" STRING
365
366 cipher_list ::= "cipher-list" ":" STRING
367
368 sanity_checks ::= "sanity-checks" ":" "{" sanity_checks_params "}"
369
370 sanity_checks_params ::= sanity_checks_param
371                        | sanity_checks_params "," sanity_checks_param
372                        | sanity_checks_params ","
373
374 sanity_checks_param ::= lease_checks
375                      | extended_info_checks
376
377 lease_checks ::= "lease-checks" ":" STRING
378
379 extended_info_checks ::= "extended-info-checks" ":" STRING
380
381 mac_sources ::= "mac-sources" ":" "[" mac_sources_list "]"
382
383 mac_sources_list ::= mac_sources_value
384                  | mac_sources_list "," mac_sources_value
385                  | mac_sources_list ","
386
387 mac_sources_value ::= duid_id
388                  | string_id
389
390 duid_id ::= "duid"
391
392 string_id ::= STRING
393
394 host_reservation_identifiers ::= "host-reservation-identifiers" ":" "[" host_
395 ↪reservation_identifiers_list "]"
396
397 host_reservation_identifiers_list ::= host_reservation_identifier
398                                   | host_reservation_identifiers_list "," host_
399 ↪reservation_identifier
400                                   | host_reservation_identifiers_list ","
401
402 host_reservation_identifier ::= duid_id
403                             | hw_address_id
404                             | flex_id
405
406 hw_address_id ::= "hw-address"
407
408 flex_id ::= "flex-id"
409
410 relay_supplied_options ::= "relay-supplied-options" ":" "[" list_content "]"
411
412 dhcp_multi_threading ::= "multi-threading" ":" "{" multi_threading_params "}"
413
414 multi_threading_params ::= multi_threading_param

```

(continues on next page)

(continued from previous page)

```

414         | multi_threading_params "," multi_threading_param
415         | multi_threading_params ","
416
417 multi_threading_param ::= enable_multi_threading
418                        | thread_pool_size
419                        | packet_queue_size
420                        | user_context
421                        | comment
422                        | unknown_map_entry
423
424 enable_multi_threading ::= "enable-multi-threading" ":" BOOLEAN
425
426 thread_pool_size ::= "thread-pool-size" ":" INTEGER
427
428 packet_queue_size ::= "packet-queue-size" ":" INTEGER
429
430 hooks_libraries ::= "hooks-libraries" ":" "[" hooks_libraries_list "]"
431
432 hooks_libraries_list ::=
433                        | not_empty_hooks_libraries_list
434
435 not_empty_hooks_libraries_list ::= hooks_library
436                                | not_empty_hooks_libraries_list "," hooks_library
437                                | not_empty_hooks_libraries_list ","
438
439 hooks_library ::= "{" hooks_params "}"
440
441 sub_hooks_library ::= "{" hooks_params "}"
442
443 hooks_params ::= hooks_param
444                | hooks_params "," hooks_param
445                | hooks_params ","
446                | unknown_map_entry
447
448 hooks_param ::= library
449              | parameters
450
451 library ::= "library" ":" STRING
452
453 parameters ::= "parameters" ":" map_value
454
455 expired_leases_processing ::= "expired-leases-processing" ":" "{" expired_leases_
456 ↪ params "}"
457
458 expired_leases_params ::= expired_leases_param
459                        | expired_leases_params "," expired_leases_param
460                        | expired_leases_params ","
461
462 expired_leases_param ::= reclaim_timer_wait_time
463                      | flush_reclaimed_timer_wait_time
464                      | hold_reclaimed_time
465                      | max_reclaim_leases

```

(continues on next page)

(continued from previous page)

```

465         | max_reclaim_time
466         | unwarned_reclaim_cycles
467
468 reclaim_timer_wait_time ::= "reclaim-timer-wait-time" ":" INTEGER
469
470 flush_reclaimed_timer_wait_time ::= "flush-reclaimed-timer-wait-time" ":" INTEGER
471
472 hold_reclaimed_time ::= "hold-reclaimed-time" ":" INTEGER
473
474 max_reclaim_leases ::= "max-reclaim-leases" ":" INTEGER
475
476 max_reclaim_time ::= "max-reclaim-time" ":" INTEGER
477
478 unwarned_reclaim_cycles ::= "unwarned-reclaim-cycles" ":" INTEGER
479
480 subnet6_list ::= "subnet6" ":" "[" subnet6_list_content "]"
481
482 subnet6_list_content ::=
483     | not_empty_subnet6_list
484
485 not_empty_subnet6_list ::= subnet6
486     | not_empty_subnet6_list "," subnet6
487     | not_empty_subnet6_list ","
488
489 subnet6 ::= "{" subnet6_params "}"
490
491 sub_subnet6 ::= "{" subnet6_params "}"
492
493 subnet6_params ::= subnet6_param
494     | subnet6_params "," subnet6_param
495     | subnet6_params ","
496
497 subnet6_param ::= preferred_lifetime
498     | min_preferred_lifetime
499     | max_preferred_lifetime
500     | valid_lifetime
501     | min_valid_lifetime
502     | max_valid_lifetime
503     | renew_timer
504     | rebind_timer
505     | option_data_list
506     | pools_list
507     | pd_pools_list
508     | subnet
509     | interface
510     | interface_id
511     | id
512     | rapid_commit
513     | client_class
514     | network_client_classes
515     | require_client_classes
516     | evaluate_additional_classes

```

(continues on next page)

(continued from previous page)

```

517         | reservations
518         | reservations_global
519         | reservations_in_subnet
520         | reservations_out_of_pool
521         | relay
522         | user_context
523         | comment
524         | calculate_tee_times
525         | t1_percent
526         | t2_percent
527         | cache_threshold
528         | cache_max_age
529         | hostname_char_set
530         | hostname_char_replacement
531         | ddns_send_updates
532         | ddns_override_no_update
533         | ddns_override_client_update
534         | ddns_replace_client_name
535         | ddns_generated_prefix
536         | ddns_qualifying_suffix
537         | ddns_update_on_renew
538         | ddns_use_conflict_resolution
539         | ddns_conflict_resolution_mode
540         | ddns_ttl_percent
541         | ddns_ttl
542         | ddns_ttl_min
543         | ddns_ttl_max
544         | store_extended_info
545         | allocator
546         | pd_allocator
547         | unknown_map_entry
548
549 subnet ::= "subnet" ":" STRING
550
551 interface ::= "interface" ":" STRING
552
553 interface_id ::= "interface-id" ":" STRING
554
555 client_class ::= "client-class" ":" STRING
556
557 network_client_classes ::= "client-classes" ":" list_strings
558
559 require_client_classes ::= "require-client-classes" ":" list_strings
560
561 evaluate_additional_classes ::= "evaluate-additional-classes" ":" list_strings
562
563 reservations_global ::= "reservations-global" ":" BOOLEAN
564
565 reservations_in_subnet ::= "reservations-in-subnet" ":" BOOLEAN
566
567 reservations_out_of_pool ::= "reservations-out-of-pool" ":" BOOLEAN
568

```

(continues on next page)

(continued from previous page)

```

569 id ::= "id" ":" INTEGER
570
571 rapid_commit ::= "rapid-commit" ":" BOOLEAN
572
573 shared_networks ::= "shared-networks" ":" "[" shared_networks_content "]"
574
575 shared_networks_content ::=
576     | shared_networks_list
577
578 shared_networks_list ::= shared_network
579     | shared_networks_list "," shared_network
580     | shared_networks_list ","
581
582 shared_network ::= "{" shared_network_params "}"
583
584 shared_network_params ::= shared_network_param
585     | shared_network_params "," shared_network_param
586     | shared_network_params ","
587
588 shared_network_param ::= name
589     | subnet6_list
590     | interface
591     | interface_id
592     | renew_timer
593     | rebind_timer
594     | option_data_list
595     | relay
596     | reservations_global
597     | reservations_in_subnet
598     | reservations_out_of_pool
599     | client_class
600     | network_client_classes
601     | require_client_classes
602     | evaluate_additional_classes
603     | preferred_lifetime
604     | min_preferred_lifetime
605     | max_preferred_lifetime
606     | rapid_commit
607     | valid_lifetime
608     | min_valid_lifetime
609     | max_valid_lifetime
610     | user_context
611     | comment
612     | calculate_tee_times
613     | t1_percent
614     | t2_percent
615     | cache_threshold
616     | cache_max_age
617     | hostname_char_set
618     | hostname_char_replacement
619     | ddns_send_updates
620     | ddns_override_no_update

```

(continues on next page)

(continued from previous page)

```

621         | ddns_override_client_update
622         | ddns_replace_client_name
623         | ddns_generated_prefix
624         | ddns_qualifying_suffix
625         | ddns_update_on_renew
626         | ddns_use_conflict_resolution
627         | ddns_conflict_resolution_mode
628         | ddns_ttl_percent
629         | ddns_ttl
630         | ddns_ttl_min
631         | ddns_ttl_max
632         | store_extended_info
633         | allocator
634         | pd_allocator
635         | unknown_map_entry
636
637 option_def_list ::= "option-def" ":" "[" option_def_list_content "]"
638
639 sub_option_def_list ::= "{" option_def_list "}"
640
641 option_def_list_content ::=
642         | not_empty_option_def_list
643
644 not_empty_option_def_list ::= option_def_entry
645         | not_empty_option_def_list "," option_def_entry
646         | not_empty_option_def_list ","
647
648 option_def_entry ::= "{" option_def_params "}"
649
650 sub_option_def ::= "{" option_def_params "}"
651
652 option_def_params ::=
653         | not_empty_option_def_params
654
655 not_empty_option_def_params ::= option_def_param
656         | not_empty_option_def_params "," option_def_param
657         | not_empty_option_def_params ","
658
659 option_def_param ::= option_def_name
660         | option_def_code
661         | option_def_type
662         | option_def_record_types
663         | option_def_space
664         | option_def_encapsulate
665         | option_def_array
666         | user_context
667         | comment
668         | unknown_map_entry
669
670 option_def_name ::= name
671
672 code ::= "code" ":" INTEGER

```

(continues on next page)

(continued from previous page)

```

673 option_def_code ::= code
674
675 option_def_type ::= "type" ":" STRING
676
677 option_def_record_types ::= "record-types" ":" STRING
678
679
680 space ::= "space" ":" STRING
681
682 option_def_space ::= space
683
684 option_def_encapsulate ::= "encapsulate" ":" STRING
685
686 option_def_array ::= "array" ":" BOOLEAN
687
688 option_data_list ::= "option-data" ":" "[" option_data_list_content "]"
689
690 option_data_list_content ::=
691     | not_empty_option_data_list
692
693 not_empty_option_data_list ::= option_data_entry
694     | not_empty_option_data_list "," option_data_entry
695     | not_empty_option_data_list ","
696
697 option_data_entry ::= "{" option_data_params "}"
698
699 sub_option_data ::= "{" option_data_params "}"
700
701 option_data_params ::=
702     | not_empty_option_data_params
703
704 not_empty_option_data_params ::= option_data_param
705     | not_empty_option_data_params "," option_data_param
706     | not_empty_option_data_params ","
707
708 option_data_param ::= option_data_name
709     | option_data_data
710     | option_data_code
711     | option_data_space
712     | option_data_csv_format
713     | option_data_always_send
714     | option_data_never_send
715     | user_context
716     | comment
717     | option_data_client_classes
718     | unknown_map_entry
719
720 option_data_name ::= name
721
722 option_data_data ::= "data" ":" STRING
723
724 option_data_code ::= code

```

(continues on next page)

(continued from previous page)

```

725 option_data_space ::= space
726
727 option_data_csv_format ::= "csv-format" ":" BOOLEAN
728
729 option_data_always_send ::= "always-send" ":" BOOLEAN
730
731 option_data_never_send ::= "never-send" ":" BOOLEAN
732
733 option_data_client_classes ::= "client-classes" ":" list_strings
734
735 pools_list ::= "pools" ":" "[" pools_list_content "]"
736
737 pools_list_content ::=
738     | not_empty_pools_list
739
740 not_empty_pools_list ::= pool_list_entry
741     | not_empty_pools_list "," pool_list_entry
742     | not_empty_pools_list ","
743
744 pool_list_entry ::= "{" pool_params "}"
745
746 sub_pool6 ::= "{" pool_params "}"
747
748 pool_params ::= pool_param
749     | pool_params "," pool_param
750     | pool_params ","
751
752 pool_param ::= pool_entry
753     | pool_id
754     | option_data_list
755     | client_class
756     | network_client_classes
757     | require_client_classes
758     | evaluate_additional_classes
759     | ddns_send_updates
760     | ddns_override_no_update
761     | ddns_override_client_update
762     | ddns_replace_client_name
763     | ddns_generated_prefix
764     | ddns_qualifying_suffix
765     | ddns_update_on_renew
766     | ddns_conflict_resolution_mode
767     | ddns_ttl_percent
768     | ddns_ttl
769     | ddns_ttl_min
770     | ddns_ttl_max
771     | hostname_char_set
772     | hostname_char_replacement
773     | user_context
774     | comment
775     | unknown_map_entry
776

```

(continues on next page)

(continued from previous page)

```

777 pool_entry ::= "pool" ":" STRING
778
779
780 pool_id ::= "pool-id" ":" INTEGER
781
782 user_context ::= "user-context" ":" map_value
783
784 comment ::= "comment" ":" STRING
785
786 pd_pools_list ::= "pd-pools" ":" "[" pd_pools_list_content "]"
787
788 pd_pools_list_content ::=
789     | not_empty_pd_pools_list
790
791 not_empty_pd_pools_list ::= pd_pool_entry
792     | not_empty_pd_pools_list "," pd_pool_entry
793     | not_empty_pd_pools_list ","
794
795 pd_pool_entry ::= "{" pd_pool_params "}"
796
797 sub_pd_pool ::= "{" pd_pool_params "}"
798
799 pd_pool_params ::= pd_pool_param
800     | pd_pool_params "," pd_pool_param
801     | pd_pool_params ","
802
803 pd_pool_param ::= pd_prefix
804     | pd_prefix_len
805     | pd_delegated_len
806     | pool_id
807     | option_data_list
808     | client_class
809     | network_client_classes
810     | require_client_classes
811     | evaluate_additional_classes
812     | excluded_prefix
813     | excluded_prefix_len
814     | user_context
815     | comment
816     | unknown_map_entry
817
818 pd_prefix ::= "prefix" ":" STRING
819
820 pd_prefix_len ::= "prefix-len" ":" INTEGER
821
822 excluded_prefix ::= "excluded-prefix" ":" STRING
823
824 excluded_prefix_len ::= "excluded-prefix-len" ":" INTEGER
825
826 pd_delegated_len ::= "delegated-len" ":" INTEGER
827
828 reservations ::= "reservations" ":" "[" reservations_list "]"

```

(continues on next page)

(continued from previous page)

```

829
830 reservations_list ::=
831     | not_empty_reservations_list
832
833 not_empty_reservations_list ::= reservation
834     | not_empty_reservations_list "," reservation
835     | not_empty_reservations_list ","
836
837 reservation ::= "{" reservation_params "}"
838
839 sub_reservation ::= "{" reservation_params "}"
840
841 reservation_params ::=
842     | not_empty_reservation_params
843
844 not_empty_reservation_params ::= reservation_param
845     | not_empty_reservation_params "," reservation_param
846     | not_empty_reservation_params ","
847
848 reservation_param ::= duid
849     | reservation_client_classes
850     | ip_addresses
851     | prefixes
852     | excluded_prefixes
853     | hw_address
854     | hostname
855     | flex_id_value
856     | option_data_list
857     | user_context
858     | comment
859     | unknown_map_entry
860
861 ip_addresses ::= "ip-addresses" ":" list_strings
862
863 prefixes ::= "prefixes" ":" list_strings
864
865 excluded_prefixes ::= "excluded-prefixes" ":" list_strings
866
867 duid ::= "duid" ":" STRING
868
869 hw_address ::= "hw-address" ":" STRING
870
871 hostname ::= "hostname" ":" STRING
872
873 flex_id_value ::= "flex-id" ":" STRING
874
875 reservation_client_classes ::= "client-classes" ":" list_strings
876
877 relay ::= "relay" ":" "{" relay_map "}"
878
879 relay_map ::= ip_addresses
880

```

(continues on next page)

(continued from previous page)

```

881 client_classes ::= "client-classes" ":" "[" client_classes_list "]"
882
883 client_classes_list ::= client_class_entry
884                       | client_classes_list "," client_class_entry
885                       | client_classes_list ","
886
887 client_class_entry ::= "{" client_class_params "}"
888
889 client_class_params ::=
890                       | not_empty_client_class_params
891
892 not_empty_client_class_params ::= client_class_param
893                                   | not_empty_client_class_params "," client_class_param
894                                   | not_empty_client_class_params ","
895
896 client_class_param ::= client_class_name
897                       | client_class_test
898                       | client_class_template_test
899                       | only_if_required
900                       | only_in_additional_list
901                       | option_data_list
902                       | user_context
903                       | comment
904                       | preferred_lifetime
905                       | min_preferred_lifetime
906                       | max_preferred_lifetime
907                       | valid_lifetime
908                       | min_valid_lifetime
909                       | max_valid_lifetime
910                       | unknown_map_entry
911
912 client_class_name ::= name
913
914 client_class_test ::= "test" ":" STRING
915
916 client_class_template_test ::= "template-test" ":" STRING
917
918 only_if_required ::= "only-if-required" ":" BOOLEAN
919
920 only_in_additional_list ::= "only-in-additional-list" ":" BOOLEAN
921
922 server_id ::= "server-id" ":" "{" server_id_params "}"
923
924 server_id_params ::= server_id_param
925                   | server_id_params "," server_id_param
926                   | server_id_params ","
927
928 server_id_param ::= server_id_type
929                  | identifier
930                  | time
931                  | htype
932                  | enterprise_id

```

(continues on next page)

(continued from previous page)

```

933         | persist
934         | user_context
935         | comment
936         | unknown_map_entry
937
938 server_id_type ::= "type" ":" duid_type
939
940 duid_type ::= "LLT"
941             | "EN"
942             | "LL"
943
944 htype ::= "htype" ":" INTEGER
945
946 identifier ::= "identifier" ":" STRING
947
948 time ::= "time" ":" INTEGER
949
950 enterprise_id ::= "enterprise-id" ":" INTEGER
951
952 dhcp4o6_port ::= "dhcp4o6-port" ":" INTEGER
953
954 control_socket ::= "control-socket" ":" "{" control_socket_params "}"
955
956 control_sockets ::= "control-sockets" ":" "[" control_socket_list "]"
957
958 control_socket_list ::=
959     | not_empty_control_socket_list
960
961 not_empty_control_socket_list ::= control_socket_entry
962                                | not_empty_control_socket_list "," control_socket_entry
963                                | not_empty_control_socket_list ","
964
965 control_socket_entry ::= "{" control_socket_params "}"
966
967 control_socket_params ::= control_socket_param
968                       | control_socket_params "," control_socket_param
969                       | control_socket_params ","
970
971 control_socket_param ::= control_socket_type
972                      | control_socket_name
973                      | control_socket_address
974                      | control_socket_port
975                      | authentication
976                      | trust_anchor
977                      | cert_file
978                      | key_file
979                      | cert_required
980                      | http_headers
981                      | user_context
982                      | comment
983                      | unknown_map_entry
984

```

(continues on next page)

(continued from previous page)

```

985 control_socket_type ::= "socket-type" ":" control_socket_type_value
986
987 control_socket_type_value ::= "unix"
988                             | "http"
989                             | "https"
990
991 control_socket_name ::= "socket-name" ":" STRING
992
993 control_socket_address ::= "socket-address" ":" STRING
994
995 control_socket_port ::= "socket-port" ":" INTEGER
996
997 cert_required ::= "cert-required" ":" BOOLEAN
998
999 http_headers ::= "http-headers" ":" "[" http_header_list "]"
1000
1001 http_header_list ::=
1002     | not_empty_http_header_list
1003
1004 not_empty_http_header_list ::= http_header
1005                             | not_empty_http_header_list "," http_header
1006                             | not_empty_http_header_list ","
1007
1008 http_header ::= "{" http_header_params "}"
1009
1010 http_header_params ::= http_header_param
1011                    | http_header_params "," http_header_param
1012                    | http_header_params ","
1013
1014 http_header_param ::= name
1015                    | header_value
1016                    | user_context
1017                    | comment
1018                    | unknown_map_entry
1019
1020 header_value ::= "value" ":" STRING
1021
1022 authentication ::= "authentication" ":" "{" auth_params "}"
1023
1024 auth_params ::= auth_param
1025              | auth_params "," auth_param
1026              | auth_params ","
1027
1028 auth_param ::= auth_type
1029             | realm
1030             | directory
1031             | clients
1032             | comment
1033             | user_context
1034             | unknown_map_entry
1035
1036 auth_type ::= "type" ":" auth_type_value

```

(continues on next page)

(continued from previous page)

```

1037
1038 auth_type_value ::= "basic"
1039
1040 realm ::= "realm" ":" STRING
1041
1042 directory ::= "directory" ":" STRING
1043
1044 clients ::= "clients" ":" "[" clients_list "]"
1045
1046 clients_list ::=
1047     | not_empty_clients_list
1048
1049 not_empty_clients_list ::= basic_auth
1050     | not_empty_clients_list "," basic_auth
1051     | not_empty_clients_list ","
1052
1053 basic_auth ::= "{" clients_params "}"
1054
1055 clients_params ::= clients_param
1056     | clients_params "," clients_param
1057     | clients_params ","
1058
1059 clients_param ::= user
1060     | user_file
1061     | password
1062     | password_file
1063     | user_context
1064     | comment
1065     | unknown_map_entry
1066
1067 user_file ::= "user-file" ":" STRING
1068
1069 password_file ::= "password-file" ":" STRING
1070
1071 dhcp_queue_control ::= "dhcp-queue-control" ":" "{" queue_control_params "}"
1072
1073 queue_control_params ::= queue_control_param
1074     | queue_control_params "," queue_control_param
1075     | queue_control_params ","
1076
1077 queue_control_param ::= enable_queue
1078     | queue_type
1079     | capacity
1080     | user_context
1081     | comment
1082     | arbitrary_map_entry
1083
1084 enable_queue ::= "enable-queue" ":" BOOLEAN
1085
1086 queue_type ::= "queue-type" ":" STRING
1087
1088 capacity ::= "capacity" ":" INTEGER

```

(continues on next page)

(continued from previous page)

```

1089 arbitrary_map_entry ::= STRING ":" value
1090
1091 dhcp_ddns ::= "dhcp-ddns" ":" "{" dhcp_ddns_params "}"
1092
1093 sub_dhcp_ddns ::= "{" dhcp_ddns_params "}"
1094
1095 dhcp_ddns_params ::= dhcp_ddns_param
1096                      | dhcp_ddns_params "," dhcp_ddns_param
1097                      | dhcp_ddns_params ","
1098
1099 dhcp_ddns_param ::= enable_updates
1100                  | server_ip
1101                  | server_port
1102                  | sender_ip
1103                  | sender_port
1104                  | max_queue_size
1105                  | ncr_protocol
1106                  | ncr_format
1107                  | user_context
1108                  | comment
1109                  | unknown_map_entry
1110
1111 enable_updates ::= "enable-updates" ":" BOOLEAN
1112
1113 server_ip ::= "server-ip" ":" STRING
1114
1115 server_port ::= "server-port" ":" INTEGER
1116
1117 sender_ip ::= "sender-ip" ":" STRING
1118
1119 sender_port ::= "sender-port" ":" INTEGER
1120
1121 max_queue_size ::= "max-queue-size" ":" INTEGER
1122
1123 ncr_protocol ::= "ncr-protocol" ":" ncr_protocol_value
1124
1125 ncr_protocol_value ::= "UDP"
1126                      | "TCP"
1127
1128 ncr_format ::= "ncr-format" ":" "JSON"
1129
1130 config_control ::= "config-control" ":" "{" config_control_params "}"
1131
1132 sub_config_control ::= "{" config_control_params "}"
1133
1134 config_control_params ::= config_control_param
1135                      | config_control_params "," config_control_param
1136                      | config_control_params ","
1137
1138 config_control_param ::= config_databases
1139                      | config_fetch_wait_time
1140

```

(continues on next page)

(continued from previous page)

```

1141 config_databases ::= "config-databases" ":" "[" database_list "]"
1142
1143 config_fetch_wait_time ::= "config-fetch-wait-time" ":" INTEGER
1144
1145 loggers ::= "loggers" ":" "[" loggers_entries "]"
1146
1147 loggers_entries ::= logger_entry
1148                     | loggers_entries "," logger_entry
1149                     | loggers_entries ","
1150
1151 logger_entry ::= "{" logger_params "}"
1152
1153 logger_params ::= logger_param
1154                 | logger_params "," logger_param
1155                 | logger_params ","
1156
1157 logger_param ::= name
1158                | output_options_list
1159                | debuglevel
1160                | severity
1161                | user_context
1162                | comment
1163                | unknown_map_entry
1164
1165 debuglevel ::= "debuglevel" ":" INTEGER
1166
1167 severity ::= "severity" ":" STRING
1168
1169 output_options_list ::= "output-options" ":" "[" output_options_list_content "]"
1170
1171 output_options_list_content ::= output_entry
1172                               | output_options_list_content "," output_entry
1173                               | output_options_list_content ","
1174
1175 output_entry ::= "{" output_params_list "}"
1176
1177 output_params_list ::= output_params
1178                     | output_params_list "," output_params
1179                     | output_params_list ","
1180
1181 output_params ::= output
1182                | flush
1183                | maxsize
1184                | maxver
1185                | pattern
1186
1187 output ::= "output" ":" STRING
1188
1189 flush ::= "flush" ":" BOOLEAN
1190
1191 maxsize ::= "maxsize" ":" INTEGER
1192

```

(continues on next page)

(continued from previous page)

```

1193 maxver ::= "maxver" ":" INTEGER
1194
1195 pattern ::= "pattern" ":" STRING
1196
1197 compatibility ::= "compatibility" ":" "{" compatibility_params "}"
1198
1199 compatibility_params ::= compatibility_param
1200                        | compatibility_params "," compatibility_param
1201                        | compatibility_params ","
1202
1203 compatibility_param ::= lenient_option_parsing
1204                      | unknown_map_entry
1205
1206 lenient_option_parsing ::= "lenient-option-parsing" ":" BOOLEAN
1207

```

29.3 BNF Grammar for Control Agent

This grammar is generated from `agent_parser.yy`. See *The Kea Control Agent* for more details.

```

1  Grammar
2
3  $accept ::= start EOF
4
5  start ::= START_JSON json
6
7  start ::= START_AGENT agent_syntax_map
8
9  start ::= START_SUB_AGENT sub_agent
10
11 sub_agent ::= "{" global_params "}"
12
13 json ::= value
14
15 value ::= INTEGER
16         | FLOAT
17         | BOOLEAN
18         | STRING
19         | NULL
20         | map
21         | list_generic
22
23 map ::= "{" map_content "}"
24
25 map_value ::= map
26
27 map_content ::=
28             | not_empty_map
29
30 not_empty_map ::= STRING ":" value
31               | not_empty_map "," STRING ":" value

```

(continues on next page)

(continued from previous page)

```

32         | not_empty_map ","
33
34 list_generic ::= "[" list_content "]"
35
36 list_content ::=
37     | not_empty_list
38
39 not_empty_list ::= value
40     | not_empty_list "," value
41     | not_empty_list ","
42
43 unknown_map_entry ::= STRING ":"
44
45 agent_syntax_map ::= "{" global_object "}"
46
47 global_object ::= "Control-agent" ":" "{" global_params "}"
48     | global_object_comma
49
50 global_object_comma ::= global_object ","
51
52 global_params ::= global_param
53     | global_params "," global_param
54     | global_params ","
55
56 global_param ::= http_host
57     | http_port
58     | http_headers
59     | trust_anchor
60     | cert_file
61     | key_file
62     | cert_required
63     | authentication
64     | control_sockets
65     | hooks_libraries
66     | loggers
67     | user_context
68     | comment
69     | unknown_map_entry
70
71 http_host ::= "http-host" ":" STRING
72
73 http_port ::= "http-port" ":" INTEGER
74
75 trust_anchor ::= "trust-anchor" ":" STRING
76
77 cert_file ::= "cert-file" ":" STRING
78
79 key_file ::= "key-file" ":" STRING
80
81 cert_required ::= "cert-required" ":" BOOLEAN
82
83 user_context ::= "user-context" ":" map_value

```

(continues on next page)

(continued from previous page)

```

84
85 comment ::= "comment" ":" STRING
86
87 http_headers ::= "http-headers" ":" "[" http_header_list "]"
88
89 http_header_list ::=
90     | not_empty_http_header_list
91
92 not_empty_http_header_list ::= http_header
93     | not_empty_http_header_list "," http_header
94     | not_empty_http_header_list ","
95
96 http_header ::= "{" http_header_params "}"
97
98 http_header_params ::= http_header_param
99     | http_header_params "," http_header_param
100     | http_header_params ","
101
102 http_header_param ::= name
103     | header_value
104     | user_context
105     | comment
106     | unknown_map_entry
107
108 name ::= "name" ":" STRING
109
110 header_value ::= "value" ":" STRING
111
112 hooks_libraries ::= "hooks-libraries" ":" "[" hooks_libraries_list "]"
113
114 hooks_libraries_list ::=
115     | not_empty_hooks_libraries_list
116
117 not_empty_hooks_libraries_list ::= hooks_library
118     | not_empty_hooks_libraries_list "," hooks_library
119     | not_empty_hooks_libraries_list ","
120
121 hooks_library ::= "{" hooks_params "}"
122
123 hooks_params ::= hooks_param
124     | hooks_params "," hooks_param
125     | hooks_params ","
126     | unknown_map_entry
127
128 hooks_param ::= library
129     | parameters
130
131 library ::= "library" ":" STRING
132
133 parameters ::= "parameters" ":" map_value
134
135 control_sockets ::= "control-sockets" ":" "{" control_sockets_params "}"

```

(continues on next page)

(continued from previous page)

```

136 control_sockets_params ::= control_socket
137                          | control_sockets_params "," control_socket
138                          | control_sockets_params ","
139
140
141 control_socket ::= dhcp4_server_socket
142                | dhcp6_server_socket
143                | d2_server_socket
144                | unknown_map_entry
145
146 dhcp4_server_socket ::= "dhcp4" ":" "{" control_socket_params "}"
147
148 dhcp6_server_socket ::= "dhcp6" ":" "{" control_socket_params "}"
149
150 d2_server_socket ::= "d2" ":" "{" control_socket_params "}"
151
152 control_socket_params ::= control_socket_param
153                        | control_socket_params "," control_socket_param
154                        | control_socket_params ","
155
156 control_socket_param ::= socket_name
157                       | socket_type
158                       | user_context
159                       | comment
160                       | unknown_map_entry
161
162 socket_name ::= "socket-name" ":" STRING
163
164 socket_type ::= "socket-type" ":" socket_type_value
165
166 socket_type_value ::= "unix"
167
168 authentication ::= "authentication" ":" "{" auth_params "}"
169
170 auth_params ::= auth_param
171              | auth_params "," auth_param
172              | auth_params ","
173
174 auth_param ::= auth_type
175             | realm
176             | directory
177             | clients
178             | comment
179             | user_context
180             | unknown_map_entry
181
182 auth_type ::= "type" ":" auth_type_value
183
184 auth_type_value ::= "basic"
185
186 realm ::= "realm" ":" STRING
187

```

(continues on next page)

(continued from previous page)

```

188 directory ::= "directory" ":" STRING
189
190 clients ::= "clients" ":" "[" clients_list "]"
191
192 clients_list ::=
193     | not_empty_clients_list
194
195 not_empty_clients_list ::= basic_auth
196     | not_empty_clients_list "," basic_auth
197     | not_empty_clients_list ","
198
199 basic_auth ::= "{" clients_params "}"
200
201 clients_params ::= clients_param
202     | clients_params "," clients_param
203     | clients_params ","
204
205 clients_param ::= user
206     | user_file
207     | password
208     | password_file
209     | user_context
210     | comment
211     | unknown_map_entry
212
213 user ::= "user" ":" STRING
214
215 user_file ::= "user-file" ":" STRING
216
217 password ::= "password" ":" STRING
218
219 password_file ::= "password-file" ":" STRING
220
221 loggers ::= "loggers" ":" "[" loggers_entries "]"
222
223 loggers_entries ::= logger_entry
224     | loggers_entries "," logger_entry
225     | loggers_entries ","
226
227 logger_entry ::= "{" logger_params "}"
228
229 logger_params ::= logger_param
230     | logger_params "," logger_param
231     | logger_params ","
232
233 logger_param ::= name
234     | output_options_list
235     | debuglevel
236     | severity
237     | user_context
238     | comment
239     | unknown_map_entry

```

(continues on next page)

(continued from previous page)

```

240
241 debuglevel ::= "debuglevel" ":" INTEGER
242
243 severity ::= "severity" ":" STRING
244
245 output_options_list ::= "output-options" ":" "[" output_options_list_content "]"
246
247 output_options_list_content ::= output_entry
248                               | output_options_list_content "," output_entry
249                               | output_options_list_content ","
250
251 output_entry ::= "{" output_params_list "}"
252
253 output_params_list ::= output_params
254                     | output_params_list "," output_params
255                     | output_params_list ","
256
257 output_params ::= output
258                | flush
259                | maxsize
260                | maxver
261                | pattern
262
263 output ::= "output" ":" STRING
264
265 flush ::= "flush" ":" BOOLEAN
266
267 maxsize ::= "maxsize" ":" INTEGER
268
269 maxver ::= "maxver" ":" INTEGER
270
271 pattern ::= "pattern" ":" STRING

```

29.4 BNF Grammar for DHCP-DDNS

This grammar is generated from `d2_parser.yy`. See *The DHCP-DDNS Server* for more details.

```

1  Grammar
2
3  $accept ::= start EOF
4
5  start ::= TOPLEVEL_JSON sub_json
6
7  start ::= TOPLEVEL_DHCPDDNS syntax_map
8
9  start ::= SUB_DHCPDDNS sub_dhcpddns
10
11 start ::= SUB_TSIG_KEY sub_tsig_key
12
13 start ::= SUB_TSIG_KEYS sub_tsig_keys
14

```

(continues on next page)

(continued from previous page)

```

15  start ::= SUB_DDNS_DOMAIN sub_ddns_domain
16
17  start ::= SUB_DDNS_DOMAINS sub_ddns_domains
18
19  start ::= SUB_DNS_SERVER sub_dns_server
20
21  start ::= SUB_DNS_SERVERS sub_dns_servers
22
23  start ::= SUB_HOOKS_LIBRARY sub_hooks_library
24
25  value ::= INTEGER
26          | FLOAT
27          | BOOLEAN
28          | STRING
29          | NULL
30          | map2
31          | list_generic
32
33  sub_json ::= value
34
35  map2 ::= "{" map_content "}"
36
37  map_value ::= map2
38
39  map_content ::=
40              | not_empty_map
41
42  not_empty_map ::= STRING ":" value
43                  | not_empty_map "," STRING ":" value
44                  | not_empty_map ","
45
46  list_generic ::= "[" list_content "]"
47
48  list_content ::=
49                  | not_empty_list
50
51  not_empty_list ::= value
52                  | not_empty_list "," value
53                  | not_empty_list ","
54
55  unknown_map_entry ::= STRING ":"
56
57  syntax_map ::= "{" global_object "}"
58
59  global_object ::= "DhcpDdns" ":" "{" dhcpddns_params "}"
60                  | global_object_comma
61
62  global_object_comma ::= global_object ","
63
64  sub_dhcpddns ::= "{" dhcpddns_params "}"
65
66  dhcpddns_params ::= dhcpddns_param

```

(continues on next page)

(continued from previous page)

```

67         | dhcpddns_params "," dhcpddns_param
68         | dhcpddns_params ","
69
70 dhcpddns_param ::= ip_address
71                 | port
72                 | dns_server_timeout
73                 | ncr_protocol
74                 | ncr_format
75                 | forward_ddns
76                 | reverse_ddns
77                 | tsig_keys
78                 | control_socket
79                 | control_sockets
80                 | hooks_libraries
81                 | loggers
82                 | user_context
83                 | comment
84                 | unknown_map_entry
85
86 ip_address ::= "ip-address" ":" STRING
87
88 port ::= "port" ":" INTEGER
89
90 dns_server_timeout ::= "dns-server-timeout" ":" INTEGER
91
92 ncr_protocol ::= "ncr-protocol" ":" ncr_protocol_value
93
94 ncr_protocol_value ::= "UDP"
95                     | "TCP"
96
97 ncr_format ::= "ncr-format" ":" "JSON"
98
99 user_context ::= "user-context" ":" map_value
100
101 comment ::= "comment" ":" STRING
102
103 forward_ddns ::= "forward-ddns" ":" "{" ddns_mgr_params "}"
104
105 reverse_ddns ::= "reverse-ddns" ":" "{" ddns_mgr_params "}"
106
107 ddns_mgr_params ::=
108                 | not_empty_ddns_mgr_params
109
110 not_empty_ddns_mgr_params ::= ddns_mgr_param
111                             | ddns_mgr_params "," ddns_mgr_param
112                             | ddns_mgr_params ","
113
114 ddns_mgr_param ::= ddns_domains
115                 | unknown_map_entry
116
117 ddns_domains ::= "ddns-domains" ":" "[" ddns_domain_list "]"
118

```

(continues on next page)

(continued from previous page)

```

119 sub_ddns_domains ::= "[" ddns_domain_list "]"
120
121 ddns_domain_list ::=
122     | not_empty_ddns_domain_list
123
124 not_empty_ddns_domain_list ::= ddns_domain
125     | not_empty_ddns_domain_list "," ddns_domain
126     | not_empty_ddns_domain_list ","
127
128 ddns_domain ::= "{" ddns_domain_params "}"
129
130 sub_ddns_domain ::= "{" ddns_domain_params "}"
131
132 ddns_domain_params ::= ddns_domain_param
133     | ddns_domain_params "," ddns_domain_param
134     | ddns_domain_params ","
135
136 ddns_domain_param ::= ddns_domain_name
137     | ddns_key_name
138     | dns_servers
139     | user_context
140     | comment
141     | unknown_map_entry
142
143 ddns_domain_name ::= "name" ":" STRING
144
145 ddns_key_name ::= "key-name" ":" STRING
146
147 dns_servers ::= "dns-servers" ":" "[" dns_server_list "]"
148
149 sub_dns_servers ::= "[" dns_server_list "]"
150
151 dns_server_list ::= dns_server
152     | dns_server_list "," dns_server
153     | dns_server_list ","
154
155 dns_server ::= "{" dns_server_params "}"
156
157 sub_dns_server ::= "{" dns_server_params "}"
158
159 dns_server_params ::= dns_server_param
160     | dns_server_params "," dns_server_param
161     | dns_server_params ","
162
163 dns_server_param ::= dns_server_hostname
164     | dns_server_ip_address
165     | dns_server_port
166     | ddns_key_name
167     | user_context
168     | comment
169     | unknown_map_entry
170

```

(continues on next page)

(continued from previous page)

```

171 dns_server_hostname ::= "hostname" ":" STRING
172
173 dns_server_ip_address ::= "ip-address" ":" STRING
174
175 dns_server_port ::= "port" ":" INTEGER
176
177 tsig_keys ::= "tsig-keys" ":" "[" tsig_keys_list "]"
178
179 sub_tsig_keys ::= "[" tsig_keys_list "]"
180
181 tsig_keys_list ::=
182     | not_empty_tsig_keys_list
183
184 not_empty_tsig_keys_list ::= tsig_key
185     | not_empty_tsig_keys_list "," tsig_key
186     | not_empty_tsig_keys_list ","
187
188 tsig_key ::= "{" tsig_key_params "}"
189
190 sub_tsig_key ::= "{" tsig_key_params "}"
191
192 tsig_key_params ::= tsig_key_param
193     | tsig_key_params "," tsig_key_param
194     | tsig_key_params ","
195
196 tsig_key_param ::= tsig_key_name
197     | tsig_key_algorithm
198     | tsig_key_digest_bits
199     | tsig_key_secret
200     | tsig_key_secret_file
201     | user_context
202     | comment
203     | unknown_map_entry
204
205 tsig_key_name ::= "name" ":" STRING
206
207 tsig_key_algorithm ::= "algorithm" ":" STRING
208
209 tsig_key_digest_bits ::= "digest-bits" ":" INTEGER
210
211 tsig_key_secret ::= "secret" ":" STRING
212
213 tsig_key_secret_file ::= "secret-file" ":" STRING
214
215 control_socket ::= "control-socket" ":" "{" control_socket_params "}"
216
217 control_sockets ::= "control-sockets" ":" "[" control_socket_list "]"
218
219 control_socket_list ::=
220     | not_empty_control_socket_list
221
222 not_empty_control_socket_list ::= control_socket_entry

```

(continues on next page)

(continued from previous page)

```

223         | not_empty_control_socket_list "," control_socket_entry
224         | not_empty_control_socket_list ","
225
226 control_socket_entry ::= "{" control_socket_params "}"
227
228 control_socket_params ::= control_socket_param
229         | control_socket_params "," control_socket_param
230         | control_socket_params ","
231
232 control_socket_param ::= control_socket_type
233         | control_socket_name
234         | control_socket_address
235         | control_socket_port
236         | authentication
237         | trust_anchor
238         | cert_file
239         | key_file
240         | cert_required
241         | user_context
242         | comment
243         | unknown_map_entry
244
245 control_socket_type ::= "socket-type" ":" control_socket_type_value
246
247 control_socket_type_value ::= "unix"
248         | "http"
249         | "https"
250
251 control_socket_name ::= "socket-name" ":" STRING
252
253 control_socket_address ::= "socket-address" ":" STRING
254
255 control_socket_port ::= "socket-port" ":" INTEGER
256
257 trust_anchor ::= "trust-anchor" ":" STRING
258
259 cert_file ::= "cert-file" ":" STRING
260
261 key_file ::= "key-file" ":" STRING
262
263 cert_required ::= "cert-required" ":" BOOLEAN
264
265 authentication ::= "authentication" ":" "{" auth_params "}"
266
267 auth_params ::= auth_param
268         | auth_params "," auth_param
269         | auth_params ","
270
271 auth_param ::= auth_type
272         | realm
273         | directory
274         | clients

```

(continues on next page)

(continued from previous page)

```

275         | comment
276         | user_context
277         | unknown_map_entry
278
279 auth_type ::= "type" ":" auth_type_value
280
281 auth_type_value ::= "basic"
282
283 realm ::= "realm" ":" STRING
284
285 directory ::= "directory" ":" STRING
286
287 clients ::= "clients" ":" "[" clients_list "]"
288
289 clients_list ::=
290     | not_empty_clients_list
291
292 not_empty_clients_list ::= basic_auth
293     | not_empty_clients_list "," basic_auth
294     | not_empty_clients_list ","
295
296 basic_auth ::= "{" clients_params "}"
297
298 clients_params ::= clients_param
299     | clients_params "," clients_param
300     | clients_params ","
301
302 clients_param ::= user
303     | user_file
304     | password
305     | password_file
306     | user_context
307     | comment
308     | unknown_map_entry
309
310 user ::= "user" ":" STRING
311
312 user_file ::= "user-file" ":" STRING
313
314 password ::= "password" ":" STRING
315
316 password_file ::= "password-file" ":" STRING
317
318 hooks_libraries ::= "hooks-libraries" ":" "[" hooks_libraries_list "]"
319
320 hooks_libraries_list ::=
321     | not_empty_hooks_libraries_list
322
323 not_empty_hooks_libraries_list ::= hooks_library
324     | not_empty_hooks_libraries_list "," hooks_library
325     | not_empty_hooks_libraries_list ","
326

```

(continues on next page)

(continued from previous page)

```

327 hooks_library ::= "{" hooks_params "}"
328
329 sub_hooks_library ::= "{" hooks_params "}"
330
331 hooks_params ::= hooks_param
332                 | hooks_params "," hooks_param
333                 | hooks_params ","
334                 | unknown_map_entry
335
336 hooks_param ::= library
337              | parameters
338
339 library ::= "library" ":" STRING
340
341 parameters ::= "parameters" ":" map_value
342
343 loggers ::= "loggers" ":" "[" loggers_entries "]"
344
345 loggers_entries ::= logger_entry
346                  | loggers_entries "," logger_entry
347                  | loggers_entries ","
348
349 logger_entry ::= "{" logger_params "}"
350
351 logger_params ::= logger_param
352                | logger_params "," logger_param
353                | logger_params ","
354
355 logger_param ::= name
356               | output_options_list
357               | debuglevel
358               | severity
359               | user_context
360               | comment
361               | unknown_map_entry
362
363 name ::= "name" ":" STRING
364
365 debuglevel ::= "debuglevel" ":" INTEGER
366
367 severity ::= "severity" ":" STRING
368
369 output_options_list ::= "output-options" ":" "[" output_options_list_content "]"
370
371 output_options_list_content ::= output_entry
372                               | output_options_list_content "," output_entry
373                               | output_options_list_content ","
374
375 output_entry ::= "{" output_params_list "}"
376
377 output_params_list ::= output_params
378                    | output_params_list "," output_params

```

(continues on next page)

(continued from previous page)

```

379         | output_params_list ","
380
381 output_params ::= output
382               | flush
383               | maxsize
384               | maxver
385               | pattern
386
387 output ::= "output" ":" STRING
388
389 flush ::= "flush" ":" BOOLEAN
390
391 maxsize ::= "maxsize" ":" INTEGER
392
393 maxver ::= "maxver" ":" INTEGER
394
395 pattern ::= "pattern" ":" STRING

```

29.5 BNF Grammar for the Kea NETCONF Agent

This grammar is generated from `netconf_parser.yy`. See [YANG/NETCONF](#) for more details.

```

1  Grammar
2
3  $accept ::= start EOF
4
5  start ::= START_JSON json
6
7  start ::= START_NETCONF netconf_syntax_map
8
9  start ::= START_SUB_NETCONF sub_netconf
10
11 sub_netconf ::= "{" global_params "}"
12
13 json ::= value
14
15 value ::= INTEGER
16         | FLOAT
17         | BOOLEAN
18         | STRING
19         | NULL
20         | map
21         | list_generic
22
23 map ::= "{" map_content "}"
24
25 map_value ::= map
26
27 map_content ::=
28             | not_empty_map
29

```

(continues on next page)

(continued from previous page)

```

30 not_empty_map ::= STRING ":" value
31               | not_empty_map "," STRING ":" value
32               | not_empty_map ","
33
34 list_generic ::= "[" list_content "]"
35
36 list_content ::=
37               | not_empty_list
38
39 not_empty_list ::= value
40                 | not_empty_list "," value
41                 | not_empty_list ","
42
43 unknown_map_entry ::= STRING ":"
44
45 netconf_syntax_map ::= "{" global_object "}"
46
47 global_object ::= "Netconf" ":" "{" global_params "}"
48                | global_object_comma
49
50 global_object_comma ::= global_object ","
51
52 global_params ::=
53                | not_empty_global_params
54
55 not_empty_global_params ::= global_param
56                           | not_empty_global_params "," global_param
57                           | not_empty_global_params ","
58
59 global_param ::= boot_update
60                | subscribe_changes
61                | validate_changes
62                | managed_servers
63                | hooks_libraries
64                | loggers
65                | user_context
66                | comment
67                | unknown_map_entry
68
69 boot_update ::= "boot-update" ":" BOOLEAN
70
71 subscribe_changes ::= "subscribe-changes" ":" BOOLEAN
72
73 validate_changes ::= "validate-changes" ":" BOOLEAN
74
75 user_context ::= "user-context" ":" map_value
76
77 comment ::= "comment" ":" STRING
78
79 hooks_libraries ::= "hooks-libraries" ":" "[" hooks_libraries_list "]"
80
81 hooks_libraries_list ::=

```

(continues on next page)

(continued from previous page)

```

82         | not_empty_hooks_libraries_list
83
84 not_empty_hooks_libraries_list ::= hooks_library
85                                | not_empty_hooks_libraries_list "," hooks_library
86                                | not_empty_hooks_libraries_list ","
87
88 hooks_library ::= "{" hooks_params "}"
89
90 hooks_params ::= hooks_param
91               | hooks_params "," hooks_param
92               | hooks_params ","
93               | unknown_map_entry
94
95 hooks_param ::= library
96              | parameters
97
98 library ::= "library" ":" STRING
99
100 parameters ::= "parameters" ":" map_value
101
102 managed_servers ::= "managed-servers" ":" "{" servers_entries "}"
103
104 servers_entries ::=
105                 | not_empty_servers_entries
106
107 not_empty_servers_entries ::= server_entry
108                             | not_empty_servers_entries "," server_entry
109                             | not_empty_servers_entries ","
110
111 server_entry ::= dhcp4_server
112               | dhcp6_server
113               | d2_server
114               | ca_server
115               | unknown_map_entry
116
117 dhcp4_server ::= "dhcp4" ":" "{" managed_server_params "}"
118
119 dhcp6_server ::= "dhcp6" ":" "{" managed_server_params "}"
120
121 d2_server ::= "d2" ":" "{" managed_server_params "}"
122
123 ca_server ::= "ca" ":" "{" managed_server_params "}"
124
125 managed_server_params ::= managed_server_param
126                       | managed_server_params "," managed_server_param
127                       | managed_server_params ","
128
129 managed_server_param ::= model
130                       | boot_update
131                       | subscribe_changes
132                       | validate_changes
133                       | control_socket

```

(continues on next page)

(continued from previous page)

```

134         | user_context
135         | comment
136         | unknown_map_entry
137
138 model ::= "model" ":" STRING
139
140 control_socket ::= "control-socket" ":" "{" control_socket_params "}"
141
142 control_socket_params ::= control_socket_param
143         | control_socket_params "," control_socket_param
144         | control_socket_params ","
145
146 control_socket_param ::= socket_type
147         | socket_name
148         | socket_url
149         | user_context
150         | comment
151         | unknown_map_entry
152
153 socket_type ::= "socket-type" ":" socket_type_value
154
155 socket_type_value ::= "unix"
156         | "http"
157         | "stdout"
158
159 socket_name ::= "socket-name" ":" STRING
160
161 socket_url ::= "socket-url" ":" STRING
162
163 loggers ::= "loggers" ":" "[" loggers_entries "]"
164
165 loggers_entries ::= logger_entry
166         | loggers_entries "," logger_entry
167         | loggers_entries ","
168
169 logger_entry ::= "{" logger_params "}"
170
171 logger_params ::= logger_param
172         | logger_params "," logger_param
173         | logger_params ","
174
175 logger_param ::= name
176         | output_options_list
177         | debuglevel
178         | severity
179         | user_context
180         | comment
181         | unknown_map_entry
182
183 name ::= "name" ":" STRING
184
185 debuglevel ::= "debuglevel" ":" INTEGER

```

(continues on next page)

(continued from previous page)

```
186 severity ::= "severity" ":" STRING
187
188
189 output_options_list ::= "output-options" ":" "[" output_options_list_content "]"
190
191 output_options_list_content ::= output_entry
192                               | output_options_list_content "," output_entry
193                               | output_options_list_content ","
194
195 output_entry ::= "{" output_params_list "}"
196
197 output_params_list ::= output_params
198                     | output_params_list "," output_params
199                     | output_params_list ","
200
201 output_params ::= output
202                | flush
203                | maxsize
204                | maxver
205                | pattern
206
207 output ::= "output" ":" STRING
208
209 flush ::= "flush" ":" BOOLEAN
210
211 maxsize ::= "maxsize" ":" INTEGER
212
213 maxver ::= "maxver" ":" INTEGER
214
215 pattern ::= "pattern" ":" STRING
```

ACKNOWLEDGMENTS

Kea is an open source project designed, developed, and maintained by Internet Systems Consortium, Inc, a 501(c)3 non-profit organization. ISC is primarily funded by revenues from support subscriptions for our open source, and we encourage all professional users to consider this option. To learn more, see <https://www.isc.org/support/>.

We thank all the organizations and individuals who have helped to make Kea possible. Comcast and the Comcast Innovation Fund provided major support for the development of Kea's DHCPv4, DHCPv6, and DDNS modules. Mozilla funded initial work on the RESTful API via a MOSS award.

Kea was initially implemented as a collection of applications within the BIND 10 framework. We thank the founding sponsors of the BIND 10 project: Afilias, IIS.SE, Nominet, SIDN, JPRS, and CIRA; and additional sponsors AFNIC, CNNIC, CZ.NIC, DENIC eG, Google, RIPE NCC, Registro.br, .nz Registry Services, and Technical Center of Internet.