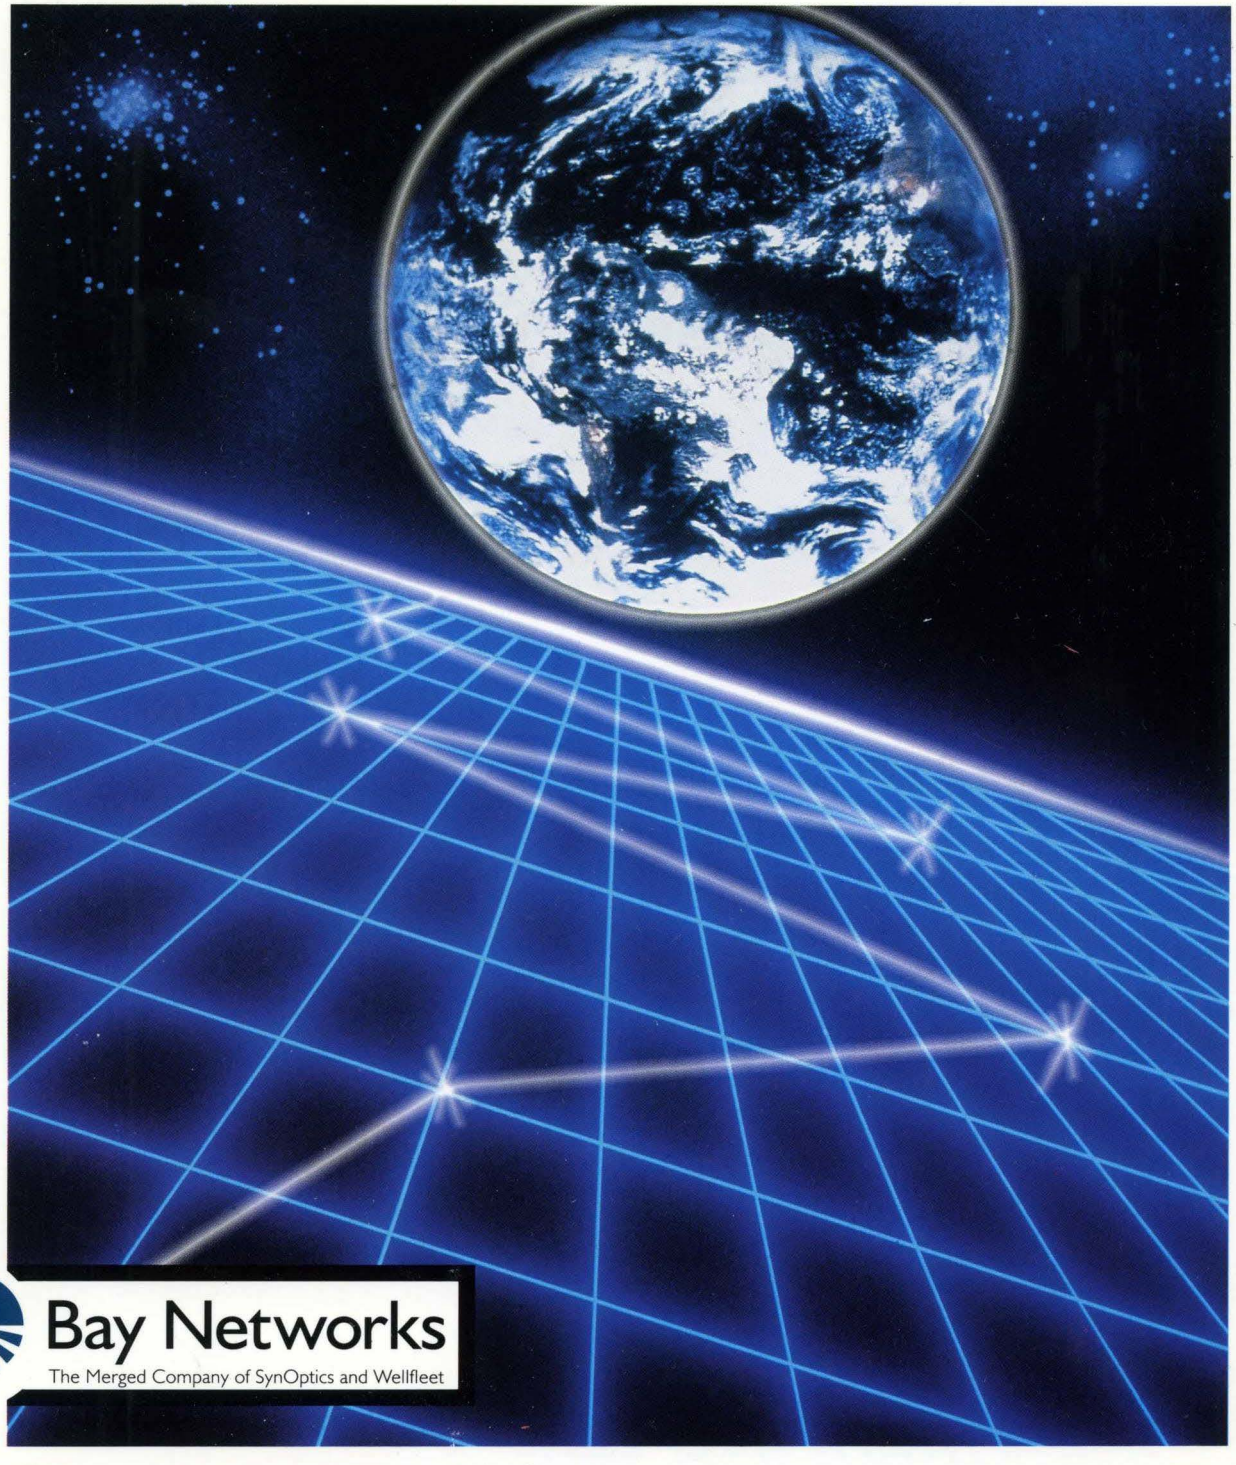

Bay Networks

The Merged Company of SynOptics and Wellfleet

Customizing OSI Services

Part No. 110047 A

Customizing OSI Services

Router Software Version 8.10
Site Manager Software Version 2.10

Part No. 110047 Rev. A
February 1995

Bay Networks

The Merged Company of SynOptics and Wellfleet

Bay Networks, Inc., 8 Federal Street, Billerica, MA 01821

Bay Networks Software License

This Software License shall govern the licensing of all software provided to licensee by Bay Networks ("Software"). Bay Networks will provide licensee with Software in machine-readable form and related documentation ("Documentation"). The Software provided under this license is proprietary to Bay Networks and to third parties from whom Bay Networks has acquired license rights. Bay Networks will not grant any Software license whatsoever, either explicitly or implicitly, except by acceptance of an order for either Software or for a Bay Networks product ("Equipment") that is packaged with Software. Each such license is subject to the following restrictions:

1. Upon delivery of the Software, Bay Networks grants to licensee a personal, nontransferable, nonexclusive license to use the Software with the Equipment with which or for which it was originally acquired, including use at any of licensee's facilities to which the Equipment may be transferred, for the useful life of the Equipment unless earlier terminated by default or cancellation. Use of the Software shall be limited to such Equipment and to such facility. Software which is licensed for use on hardware not offered by Bay Networks is not subject to restricted use on any Equipment, however, unless otherwise specified on the Documentation, each licensed copy of such Software may only be installed on one hardware item at any time.

2. Licensee may use the Software with backup Equipment only if the Equipment with which or for which it was acquired is inoperative.

3. Licensee may make a single copy of the Software (but not firmware) for safekeeping (archives) or backup purposes.

4. Licensee may modify Software (but not firmware), or combine it with other software, subject to the provision that those portions of the resulting software which incorporate Software are subject to the restrictions of this license. Licensee shall not make the resulting software available for use by any third party.

5. Neither title nor ownership to Software passes to licensee.

6. Licensee shall not provide, or otherwise make available, any Software, in whole or in part, in any form, to any third party. Third parties do not include consultants, subcontractors, or agents of licensee who have licensee's permission to use the Software at licensee's facility, and who have agreed in writing to use the Software only in accordance with the restrictions of this license.

7. Third-party owners from whom Bay Networks has acquired license rights to software that is incorporated into Bay Networks products shall have the right to enforce the provisions of this license against licensee.

8. Licensee shall not remove or obscure any copyright, patent, trademark, trade secret, or similar intellectual property or restricted rights notice within or affixed to any Software and shall reproduce and affix such notice on any backup copy of Software or copies of software resulting from modification or combination performed by licensee as permitted by this license.

9. Licensee shall not reverse assemble, reverse compile, or in any way reverse engineer the Software. [Note: For licensees in the European Community, the Software Directive dated 14 May 1991 (as may be amended from time to time) shall apply for interoperability purposes. Licensee must notify Bay Networks in writing of any such intended examination of the Software and Bay Networks may provide review and assistance.]

10. Notwithstanding any foregoing terms to the contrary, if licensee licenses the Bay Networks product "Site Manager," licensee may duplicate and install the Site Manager product as specified in the Documentation. This right is granted solely as necessary for use of Site Manager on hardware installed with licensee's network.

11. This license will automatically terminate upon improper handling of Software, such as by disclosure, or Bay Networks may terminate this license by written notice to licensee if licensee fails to comply with any of the material provisions of this license and fails to cure such failure within thirty (30) days after the receipt of written notice from Bay Networks. Upon termination of this license, licensee shall discontinue all use of the Software and return the Software and Documentation, including all copies, to Bay Networks.

12. Licensee's obligations under this license shall survive expiration or termination of this license.

Contents

Chapter 1
OSI Overview

Chapter 2
OSI Implementation Notes

Chapter 3
Editing OSI Parameters

Appendix A
IP-to-X.121 Address Mapping for DDN

Figures

Tables

About This Guide

If you are responsible for configuring and managing Wellfleet® routers, you need to read this guide.

This guide describes how to customize Wellfleet router software for Open Systems Interconnection (OSI) services.

Refer to this guide for

- An overview of the OSI routing protocol and a description of how Wellfleet routing services work (see the "OSI Overview" chapter)
- Implementation notes that may affect how you configure OSI routing services (see the "OSI Implementation Notes" chapter), including
  - — DECnet™ IV to V Transition
  - — OSI over frame relay
- Instructions on editing OSI global and interface parameters and configuring OSI services (see the "Editing OSI Parameters" chapter), including
  - — Static end system adjacencies
  - — Static external address adjacencies
  - — Static routes

Before You Begin

Before using this guide, you must complete the following procedures:

- Create and save a configuration file that contains at least one OSI interface.
- Retrieve the configuration file in local, remote, or dynamic mode.

Refer to *Configuring Wellfleet Routers* for instructions.

How to Get Help

For additional information or advice, contact the Bay Networks Help Desk in your area:

United States	1-800-2LAN-WAN
Valbonne, France	(33) 92-966-968
Sydney, Australia	(61) 2-903-5800
Tokyo, Japan	(81) 3-328-0052

Conventions

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: if command syntax is **ping** *<ip_address>*, you enter **ping 192.32.10.12**
arrow character (➔)	Separates menu and option names in instructions. Example: Protocols➔AppleTalk identifies the AppleTalk option in the Protocols menu.
brackets ([ ])	Indicate optional elements. You can choose none, one, or all of the options.
user entry text	Denotes text that you need to enter. Example: Start up the Windows environment by entering the following after the prompt: **win**
command text	Denotes command names in text. Example: Use the **xmodem** command.

italic text	Indicates variable values in command syntax descriptions, new terms, file and directory names, and book titles.
`screen text`	Indicates data that appears on the screen. Example: `Set Trap Monitor Filters`
ellipsis points	Horizontal (. . .) and vertical (⋮) ellipsis points indicate omitted information.
quotation marks (" ")	Indicate the title of a chapter or section within a book.
vertical line (|)	Indicates that you enter only one of the parts of the command. The vertical line separates choices. Do not type the vertical line when entering the command. Example: If the command syntax is **show at routes | nets**, you enter either **show at routes** or **show at nets**, but not both.

Acronyms

AAI	Administrative Authority Identifier
ACSE	Association Control Service Element
AFI	Authority and Format Identifier
ANSI	American National Standards Institute
ARP	Address Resolution Protocol
ASN.1	Abstract Syntax Notation One
CCITT	Consultative Committee on International Telegraphy and Telephony
CLNP	Connectionless Network Protocol
CLNS	Connectionless-mode Network Service Protocol
CSNP	Complete Sequence Number Packets
DCA	Defense Communication Agency
DCC	Data Country Code
DCE	Data-Circuit Terminating Equipment
DDN	Defense Data Network

DFI	Domain Format Identifier
DLCI	Data Link Connection Identifier
DSP	Domain Specific Part
DTE	Data-Circuit Terminating Equipment
ES-IS	End System to Intermediate System
FDDI	Fiber Distributed Data Interface
FTAM	File Transfer Access Management
GOSIP	Government OSI Profile
GSA	General Services Administration
HDLC	High Level Data Link Control
ICD	International Code Designator
IDI	Initial Domain Identifier
IDP	Initial Domain Part
IEEE	Institute of Electrical and Electronic Engineers
ILI	Intelligent Link Interface
IP	Internet Protocol
IS-IS	Intermediate System to Intermediate System
ISO	International Organization for Standardization
L1	Level 1
L2	Level 2
LAN	local area network
LSP	Link State Packet
MAC	Media Access Control
MIB	Management Information Base
MOM	Maintenance Operations Module
MOP	Maintenance Operations Protocol
OSI	Open Systems Interconnection
NSAP	Network Service Access Point
PDN	Public Data Network
PPP	Point-to-Point Protocol

PSNP	Partial Sequence Number Packet
PVC	Permanent Virtual Circuit
RFC	Request for Comment
RIP	Routing Information Protocol
SNAP	Subnetwork Access Protocol
SNMP	Simple Network Management Protocol
SNPA	Subnetwork Point of Attachment
TCP	Transmission Control Protocol
VT	Virtual Terminal

Chapter 1
OSI Overview

This chapter provides a general OSI networking overview and describes how OSI routing services for Wellfleet routers work. It includes information on OSI

- Network organization
- Level 1 and Level 2 routing
- Network addressing
- Link-state routing algorithm
- Routing protocols used

Note: This document uses the terms "intermediate system" and "router" interchangeably.

OSI Basic Reference Model

OSI is a nonproprietary distributed processing architecture. The International Organization for Standardization (ISO) developed OSI to provide communication standards. These standards allow computer systems from different vendors to communicate.

The OSI Basic Reference Model combines a structured computer system architecture with a set of common communication protocols. It comprises seven layers. Each layer provides specific functions or services and follows the corresponding OSI communication protocols to perform those services.

OSI is an "open system" architecture. Peer-to-peer common layers between systems abolish the vendor-specific restrictions imposed by other architectures. The principles of the OSI layering scheme include

- Similar services are on the same layer.
- Services provided by lower layers are transparent to the layers above it.
- The lower the layer, the more basic the services it provides.
- The top layer, which interfaces with the computer user, provides the full range of services offered by the layers below it.

OSI services for Wellfleet Version 7.60 software and later are United States Government OSI Profile (GOSIP) Version 2.0 compliant. In addition, Wellfleet router software provides support for the first three layers of the ISO/Consultative Committee on International Telegraphy and Telephony (CCITT) recommended set of standards for international open systems support and vendor interoperability. These layers are physical, data link, and network.

Table 1-1 lists some of the most common ISO standards implemented by OSI.

Table 1-1. OSI Reference Model and Common ISO Standards

Application Layer	8571 File Transfer and Access Management (FTAM) 8649 OSI Association Control Service Element (ACSE) 9040 Virtual Terminal Protocol (VT)
Presentation Layer	8822 OSI connection-oriented and connectionless presentation services 8824 Abstract Syntax Notation One (ASN.1) 9576 OSI connectionless protocol to provide connectionless service
Session Layer	8326 Session service definitions 8327 Session layer protocols
Transport Layer	8072 Transport service definition, both connection and connectionless 8073 Transport connection-oriented protocol definition 8602 Transport definition for connectionless-mode protocol
Network Layer	8473 Connectionless-mode network service 9542 End System to Intermediate System routing exchange protocol 10589 Intermediate System to Intermediate System routing exchange protocol
Data Link Layer	8802 Local area network standards (mostly derived from IEEE standards) 8471 HDLC balanced, link address information 8886 Data link service definition for OSI
Physical Layer	9314 Fiber Distributed Data Interface (FDDI) 9543 Synchronous transmission quality at DTE/DCE interface 9578 Communications connectors used in LANs

OSI Network Organization

The OSI network is made up of end systems and intermediate systems (routers) that are organized hierarchically.

- End systems originate and receive data. They do not perform any routing services.
- Intermediate systems originate and receive data, as well as forward (route) data. The Wellfleet OSI router is an intermediate system.

End systems and intermediate systems are divided administratively into separate routing *areas.* A collection of areas that are under the control of a single administration and operate common routing protocols is a *routing domain.*

A network manager defines the boundaries of routing domains. An entire group of routing domains that are under one administrative authority (for example, a company or a university) is an *administrative domain* (Figure 1-1).

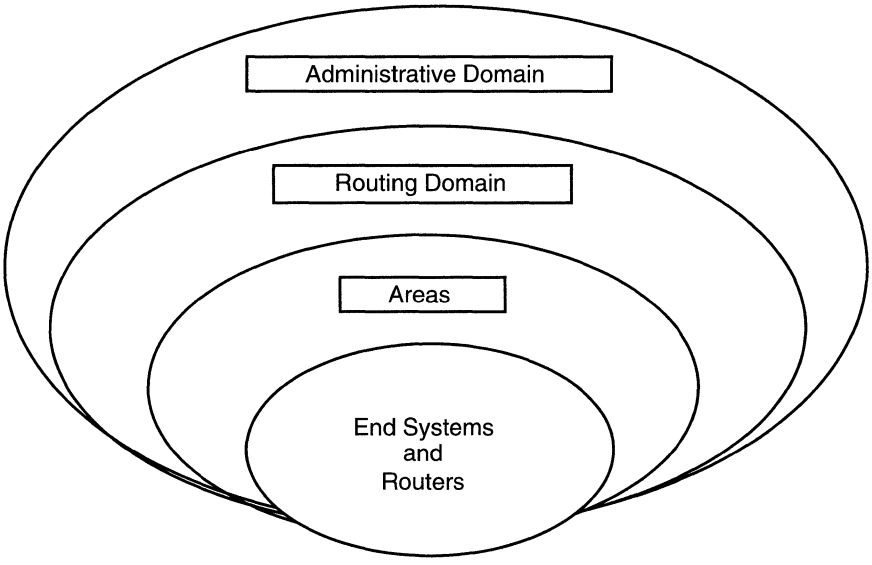

Figure 1-1. OSI Network Organization

Level 1 and Level 2 Routing

In an OSI network, the Wellfleet router running OSI transfers data in a connectionless (packet) format using the Connectionless Network Protocol (CLNP). The router routes data through the network using

- *Level 1 (L1) routing,* for routing data within an area
- *Level 2 (L2) routing*, for routing data between areas

You can configure a Wellfleet router running OSI to function as either an L1 router or L1/L2 router. An L1 router performs strictly L1 routing services. It can only route data to systems located within its local area. The L1 router forwards packets destined for a different area or domain to the nearest L1/L2 router for processing.

An L1/L2 router can perform both types of routing services. It can route data to systems located within its local area (using L1 routing) or to systems located in a different area (using L2 routing), depending on

its configuration. In addition, an L1/L2 router can route data between routing domains (called external routing), as long as you statically define the external link.

When you configure a Wellfleet L1/L2 router, you can select the routing level supported on each of the router's OSI interfaces. For example, if you want the router to perform both L1 and L2 routing over all of its interfaces, then you configure all of its interfaces to support both types of routing services. On the other hand, if you want the router to perform only L2 routing over a certain interface, then you configure that interface to support only L2 routing services.

Note: To support routing between areas, every area must contain at least one L1/L2 router configured to support L2 routing services (refer to Figure 1-2).

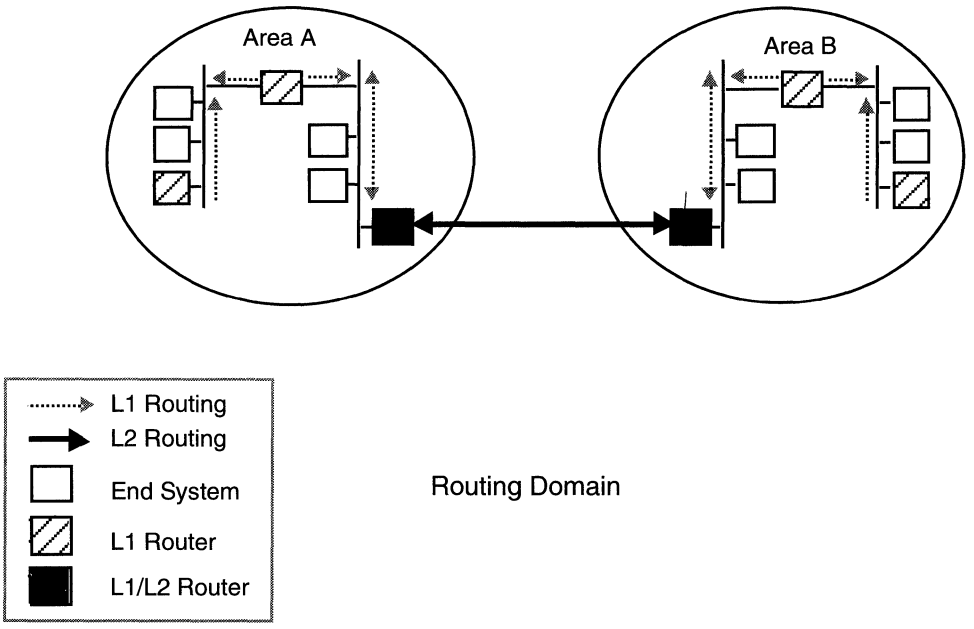

Figure 1-2. L1 and L2 Routing

OSI Network Addressing

The OSI addressing scheme is based on the hierarchical structure of the OSI global network. A unique *Network Service Access Point (NSAP)* address identifies each system within an OSI network. The NSAP address specifies the point at which the end system or intermediate system performs OSI network layer services.

The complete set of NSAP addresses contained within the OSI network is the global network addressing domain. This domain is divided into subsets called *network addressing domains* (which can be further divided into various subdomains). A network addressing domain is a set of NSAP addresses regulated by the same *addressing authority.* The addressing authority is the administration responsible for allocating unique NSAP addresses to OSI networks.

Each addressing authority operates independently of other authorities at the same level. However, if the network addressing domain administers several subdomains, then the addressing authority for the higher domain can authorize the addressing authorities for the subdomains to assign NSAP addresses (Figure 1-3). The subdomain specifies the format of the NSAP addresses allocated to the network.

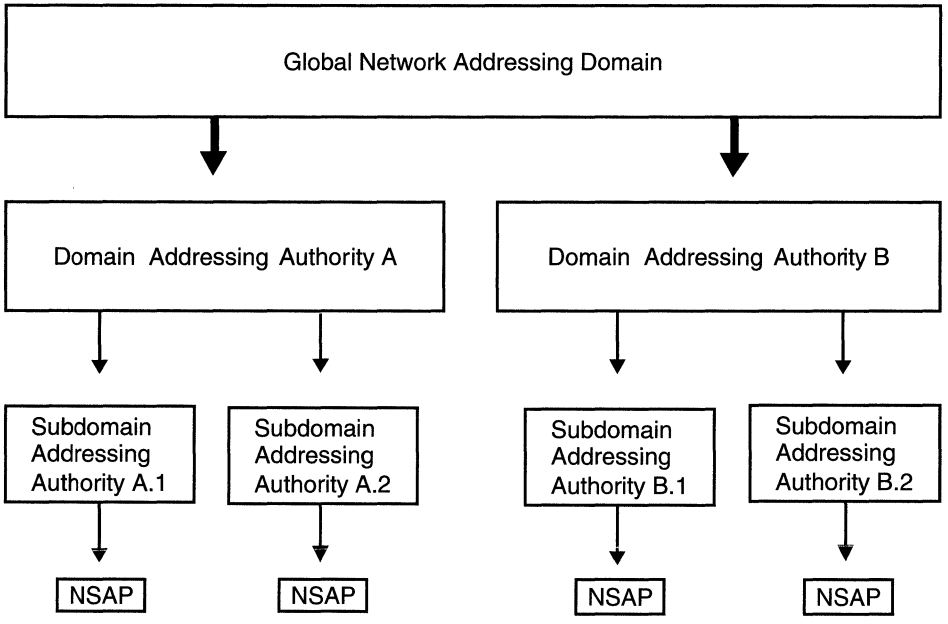

Figure 1-3. Hierarchical Addressing Authority Structure

Two of the addressing authorities that administer NSAP addresses for OSI networks in the United States include the United States General Services Administration (GSA, which allocates NSAPs that are intended primarily for government use) and the American National Standards Institute (ANSI).

NSAP Structure

The basic NSAP address structure reflects the hierarchal assignment of NSAPs throughout the global network addressing domain. NSAP addresses must be globally unique. They can be up to 20 bytes in length and contain two basic parts: the Initial Domain Part (IDP) and the Domain Specific Part (DSP) (Figure 1-4).

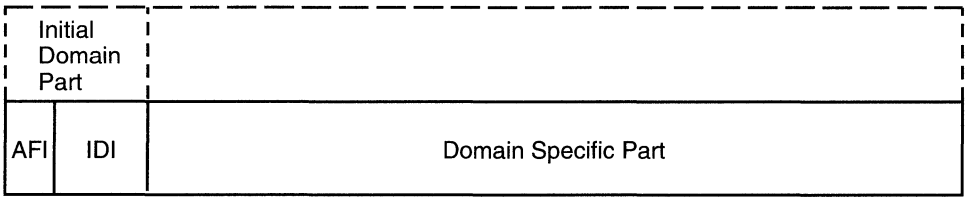

Figure 1-4. Basic NSAP Address Structure

The IDP consists of an Authority and Format Identifier (AFI) and Initial Domain Identifier (IDI). The AFI is one octet in length and specifies the format of the IDI, the network addressing authority responsible for allocating values to the IDI, and the abstract syntax of the DSP.

The IDI is variable in length. It specifies the addressing subdomain from which values of the DSP are allocated and the network addressing authority responsible for allocating values of the DSP from that subdomain. The authority identified by the IDI determines the structure and semantics of the DSP.

For example, if you register your OSI network with the United States General Services Administration, the GSA will probably assign your network to the ISO International Code Designator (ICD) 0005 subdomain. The DSP portion of the NSAP addresses allocated by this subdomain follows the GOSIP Version 2 structure described in Figure 1-5.

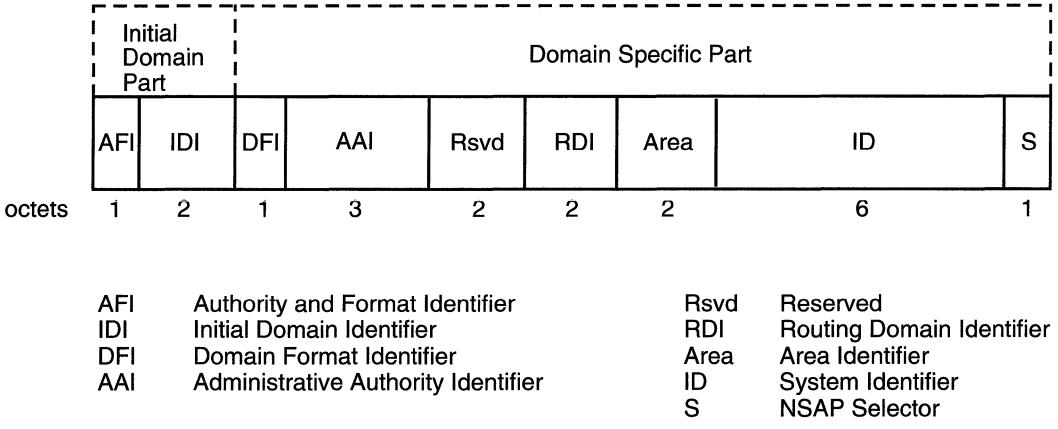

Figure 1-5. GOSIP NSAP Address Format

The AFI for these NSAP addresses is 47, which shows that the network belongs to an ICD subdomain. The IDI is 0005, specifying the ICD 0005 subdomain, which is reserved for use by the U.S. Government. The Domain Format Identifier (DFI) is 80, specifying that the DSP portion of NSAP is in GOSIP format. (Currently, the only DSP format defined by the ICD 0005 subdomain is that defined by GOSIP.)

The Administrative Authority Identifier (AAI) portion of these NSAP addresses is a globally unique number assigned by the ICD 0005 subdomain. It identifies the network within the ICD 0005 subdomain, where the NSAP resides, and the authority responsible for organizing the network into routing domains and areas. Note that the authority specified by the AAI assigns values to the Routing Domain ID, Area ID, System ID, and NSAP Selector portions of the NSAP address. Table 1-2 describes the contents of each field for this type of NSAP address.

Table 1-2. NSAP Address Structure (Assigned by the ICD 0005 Subdomain)

Field	Value	Meaning
AFI	47	Identifies the subdomain as ICD. Specifies the syntax of the DSP as binary octets.
IDI	0005	Indicates that the subdomain is ICD 0005.
DFI	80	Specifies that the format of the DSP is GOSIP.
AAI	variable	Identifies the network within the ICD 0005 subdomain where the NSAP resides, and the authority responsible for organizing the network into routing domains and areas.
RSVD	0000	Indicates that this field is reserved.
RDI	variable	Specifies the routing domain where the NSAP resides (assigned by the authority identified in the AAI field).
Area	variable	Identifies the local area where the NSAP resides (assigned by either the authority identified in the AAI field or the local administrative authority that the AAI authority has delegated to this routing domain).
ID	variable	Specifies the system where the NSAP resides (assigned by the local area administrator that a higher authority has delegated to this area).
S	0 or 1	Selects the transport layer entity the system uses. This entity is specified in the ID field.

Similarly, if you register your OSI network with the ANSI, your network is assigned to the ISO Data Country Code (DCC) 840 subdomain. Currently, the structure of the DSP portion of NSAP addresses allocated by the DCC 840 subdomain is not standardized. However, the most recent proposal suggests a structure identical to that specified by GOSIP, with the Administrative Authority Identifier field replaced by an Organization Identifier field (refer to Figure 1-6).

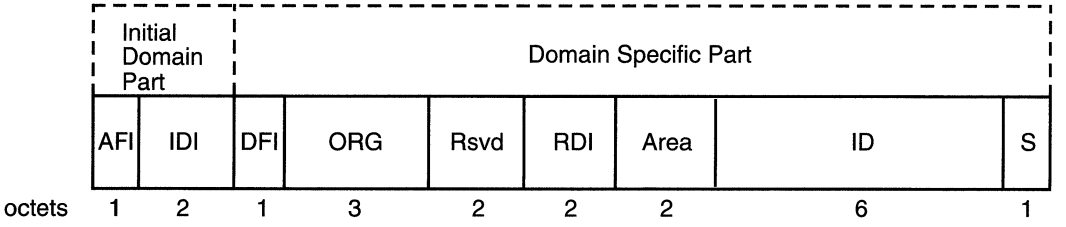

AFI	Authority and Format Identifier	Rsvd	Reserved
IDI	Initial Domain Identifier	RDI	Routing Domain Identifier
DFI	Domain Format Identifier	Area	Area Identifier
ORG	Organization Identifier	ID	System Identifier
		S	NSAP Selector

Figure 1-6. ANSI NSAP Address Format

The AFI for these NSAP addresses is 39, which shows that the network is registered with ANSI and belongs to a DCC subdomain. The IDI is 840, specifying the DCC 840 subdomain, which is reserved for use by networks located in the United States. The DFI is not standardized and is assigned by the DCC 840 subdomain.

The Organization (ORG) Identifier portion of the NSAP address is a globally unique number that is assigned by the DCC 840 subdomain. It identifies the network within the DCC 840 subdomain where the NSAP resides and the authority responsible for organizing the network into routing domains and areas. (The Organization Identifier serves the same purpose as the Administrative Authority portion of a NSAP assigned by the ICD 0005 subdomain; refer to Table 1-2.) Table 1-3 describes the contents of each field for this type of NSAP address.)

Table 1-3. NSAP Address Structure (Assigned by the DCC 840 Subdomain)

Field Name	Value	Meaning
AFI	39	Identifies the subdomain as DCC 840. Specifies the syntax of the DSP as binary octets.
IDI	840	Indicates that the subdomain is DCC 840.
DFI	variable	Identifies the format of the DSP. The subdomain identified in the IDI specifies this value.
ORG	variable	Specifies the network within the DCC 840 subdomain, where the NSAP resides, and the authority responsible for organizing the network into routing domains and areas.
Rsvd	0000	Indicates that this field is reserved.
RDI	variable	Identifies the routing domain where the NSAP resides (assigned by the authority identified in the ORG field).
Area	variable	Specifies the local area where the NSAP resides (assigned by either the authority identified in the ORG field or the local administrative authority that the ORG authority has delegated to this routing domain).
ID	variable	Identifies the system where the NSAP resides (assigned by the local area administrator that a higher authority has delegated to this area).
S	0 or 1	Selects the transport layer entity the system uses. This entity is specified in the ID field.

The IDP and the first part of the DSP (called the High Order Part of the DSP) is the NSAP's *area address*. The area address identifies the area in an OSI network where an NSAP resides (refer to Figure 1-7).

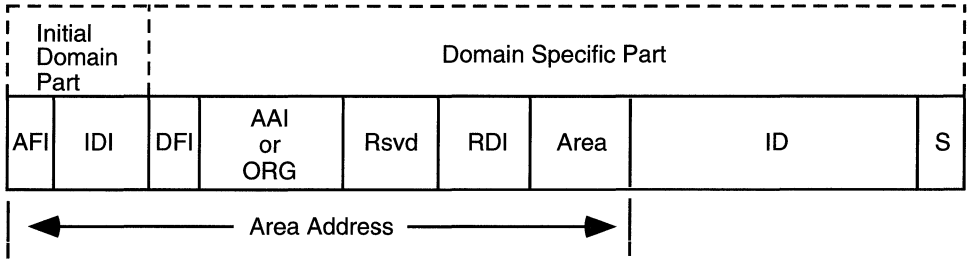

Figure 1-7. NSAP Area Address

When a router receives a packet, it examines the contents of the packet's NSAP destination area address fields. The router compares its own NSAP area address(es) with the NSAP destination address contained in the packet's header. If they match, then the destination system is in that router's area. If the addresses do not match, then the destination system is located in a different area and the router must route the packet outside of the local area using L2 routing services.

Allocating NSAP Addresses

To demonstrate how NSAP addresses are allocated, Figure 1-8 shows a sample OSI network set up on a college campus in the United States. To obtain and allocate NSAP addresses for the OSI network, the network administrator did the following:

1. Divided the campus OSI network into areas.

   The administrator divided the campus OSI network into Areas A, B, and C. These three areas make up the campus routing domain.

2. Assigned identifiers to the campus routing domain and local areas as follows:

   Campus Routing Domain Identifier = 0001
   Area A Identifier = 0001
   Area B Identifier = 0002
   Area C Identifier = 0003

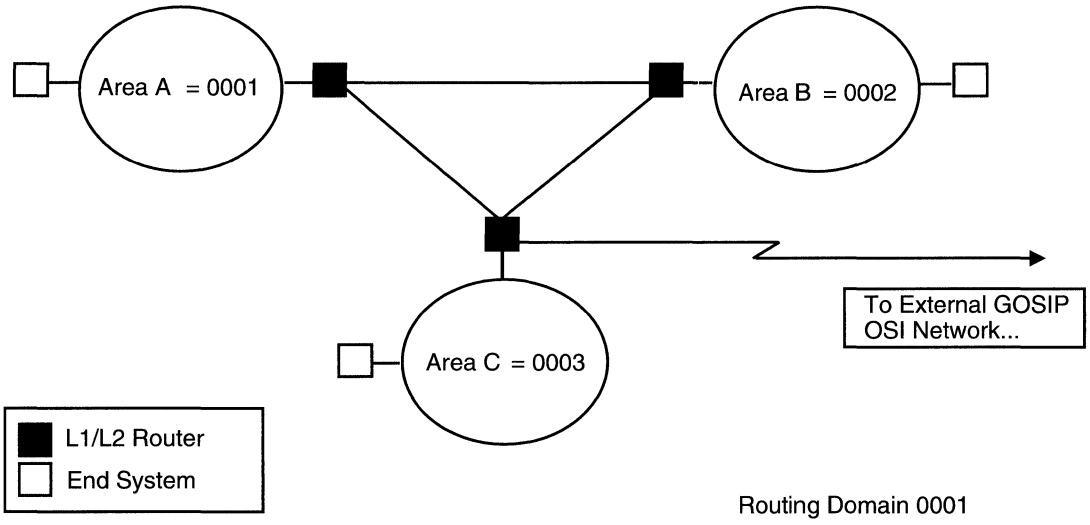

Figure 1-8. Campus Routing Domain

3. Registered the campus network with the Addressing Authorities.

   Because Area A and Area B are not linked to any areas outside of the campus routing domain, the administrator obtained NSAP addresses for Area A and Area B by simply registering the campus network with ANSI. ANSI assigned the network to the DCC 840 subdomain, which in turn assigned an Organization Identifier of 113527 to the network.

   Area C, however, is linked to an external domain that is operated by the federal government. So besides registering the network with ANSI, the administrator also registered the network with the GSA (in order to receive NSAP addresses in GOSIP format for those systems residing in Area C). The GSA assigned the network to the ICD 0005 subdomain, which in turn assigned an Administrative Authority Identifier of 00004e to the network.

4. Assigned full NSAP addresses to the routers and end systems in Area A, Area B, and Area C.

After receiving the Organization ID for the campus network from the DCC 840 subdomain, the administrator assigned full NSAP addresses to the routers and end systems in Area A and Area B (Figure 1-9). Note that the DSP portion is structured according to DCC 840 subdomain standard format.

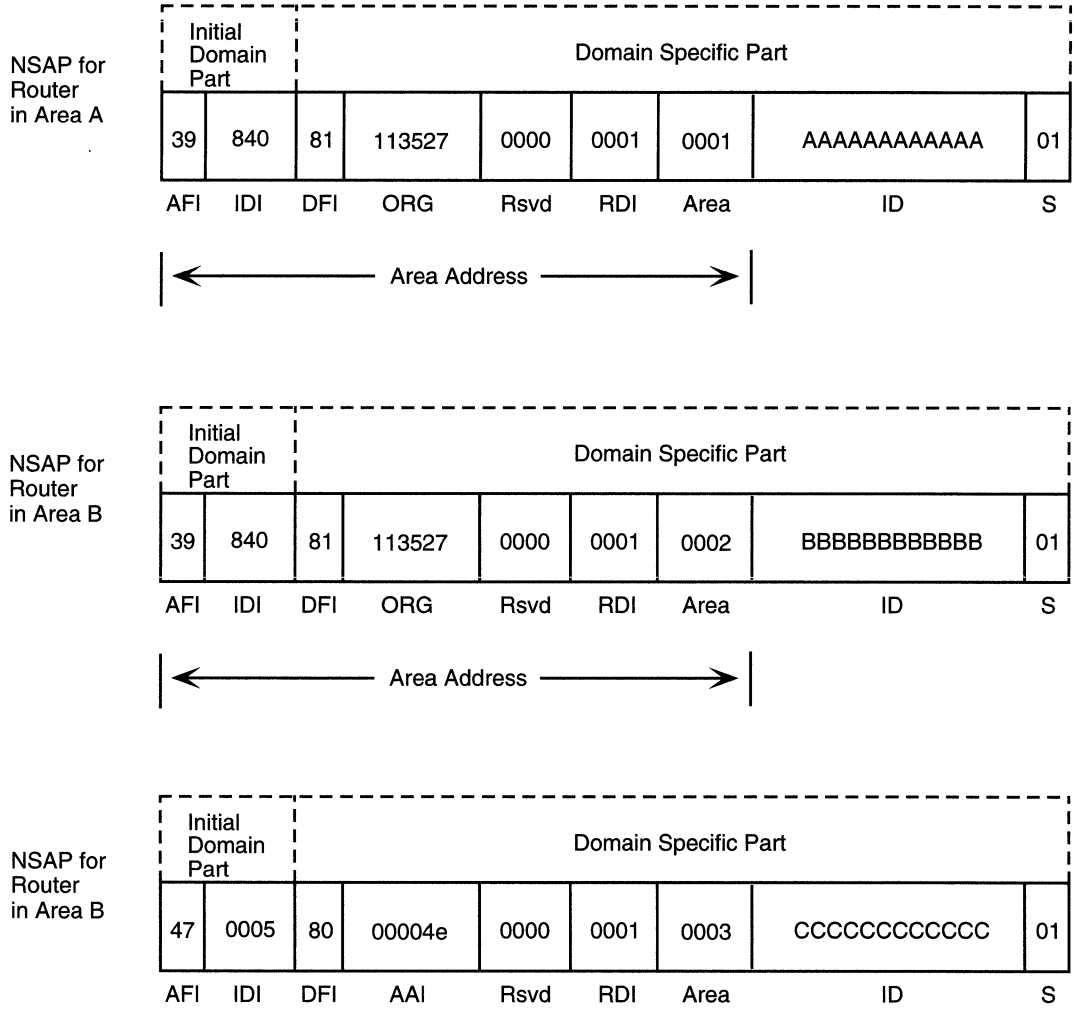

Figure 1-9. Assigning NSAP Addresses

Similarly, after receiving the AAI for the campus network from the ICD 0005 subdomain, the administrator assigned a full NSAP address to the router and end systems in Area C. The DSP portion is structured according to ICD 0005 subdomain standard format.

OSI Basic Routing Algorithm

The OSI routing algorithm is based on link state information. That is, each OSI router periodically generates link state protocol data units, which describe the status of all of the router's immediate or adjacent data links. These link state advertisements are propagated throughout the network. The router compiles a database of the link state information and uses it to calculate the paths to all reachable destinations in the domain.

The OSI routing algorithm uses these three processes:

- Update Process

  In response to changes in network topology, routers transmit and receive LSPs. Each time a router receives a LSP, the router uses it to update its link state database with the new link state information.

- Decision Process

  Each router calculates the shortest paths from itself to all other systems that it can reach, using information it retrieves from its link state database. It then stores the paths in a forwarding database.

- Forwarding Process

  When the router receives a CLNP packet, it forwards the packet to the next hop specified in its forwarding database.

Each router in the OSI network performs routing services. The following sections describe each process in more detail.

Update Process

In an OSI network, every router is responsible for finding out the identity and reachability of its immediate or adjacent neighbors. In other words, a router must decide which network addresses it can directly reach. It uses this information, together with the assigned link cost of reaching each neighbor, to construct a *link state packet* (LSP).

LSPs describe what the router knows about the network topology. Depending on its configuration, the router generates different types of LSPs (see Table 1-4). L1 routers generate only L1 LSPs while L1/L2 routers generate both L1 and L2 LSPs.

In addition on broadcast subnetworks, the subnetwork itself is conceptually viewed as a node (called a *pseudonode*) in the OSI network. One router on the subnetwork is elected as the *designated router* for the pseudonode. The designated router is responsible for creating and transmitting a LSP on behalf of the pseudonode. Thus, the designated router generates a *pseudonode LSP*. By generating a single LSP that represents the pseudonode, the router reduces the amount of link state information that traverses the subnetwork.

The L1 designated router and the L2 designated router for a subnetwork are elected independently. If there is only a single L1 or L1/L2 router on a LAN segment, it becomes the designated L1 or L2 router by default.

Note: Because you can configure multiple OSI interfaces on the same Wellfleet router, the router can act as the designated router for some subnetworks, while not for others.

Table 1-4. Link State Packet Types

Router type	Generates LSP type	Describing	Sent to
L1 designated router	L1 pseudonode	The links to all dynamically learned L1 routers and end systems in the local area that are reachable over the broadcast subnetwork.	All L1 routers within the area
L1 router	L1 non-pseudonode	The links to the L1 designated router and static links.	All L1 routers within the area
L2 designated router	L2 pseudonode	The links to all L1 and L1/L2 routers in the domain that are reachable over the broadcast subnetwork, as any routes to external domains that exist.	All L1/L2 routers within the domain
L2 router	L2 non-pseudonode	The links to the L1/L2 designated router and static external links.	All L1/L2 routers within the domain

OSI routers generate LSPs periodically and when there is a change in the network topology. For example, if a new end system is added to Area A, Router 1 generates an L1 LSP and floods it to all other L1 routers in the area. Each router that receives the LSP uses it to update its link state database, then floods it out all interfaces except for the one that it was received upon (refer to Figure 1-10).

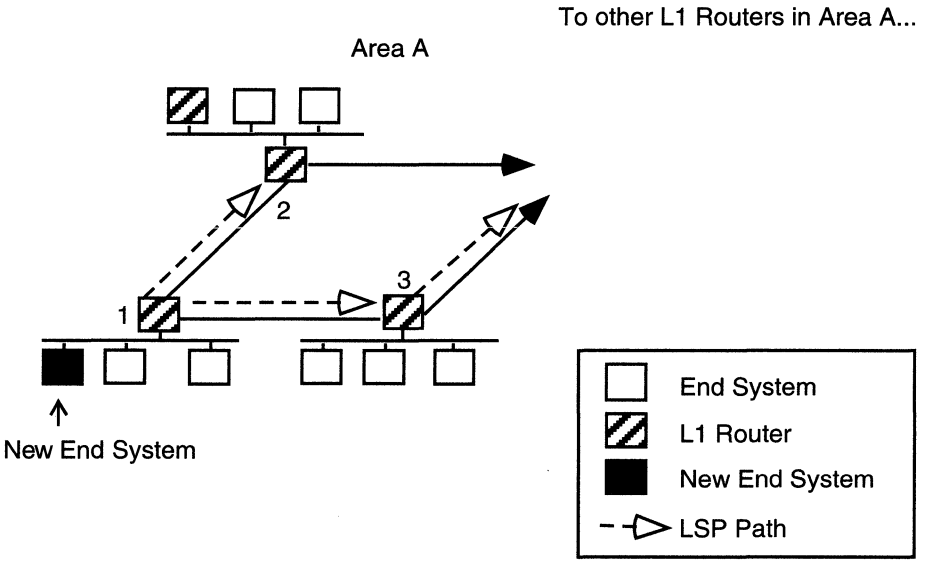

Figure 1-10. Flooding L1 LSPs within an Area

Similarly, if a new L1/L2 router is added to the network, L1/L2 routers flood both L1 and L2 LSPs throughout the domain. When an L1/L2 router receives a new LSP, it updates its corresponding L1 or L2 link state database with the new information. The router then forwards the LSP on all other links except the one that it was received upon. Note that the L1/L2 routers that support both types of traffic maintain separate L1 and L2 link state databases.

The router refers to its link state database(s) when deciding the shortest path between itself and all other routers it can reach.

Decision Process

During the decision process, the OSI router uses the link state database information that it has accumulated during the update process to

- Define a set of paths from the router to every reachable destination in the domain
- Calculate the shortest path to each destination
- Record the identity of the first hop on the shortest path to each destination into a forwarding database

The router uses a shortest path first (SPF) algorithm to define the set of paths to a destination. The shortest path, however, may not necessarily be that which is nearest to the destination. The OSI router considers the relative cost (metric) of routing a packet along each path, defining the shortest path as the lowest cost path.

Every circuit on the OSI network is assigned a relative cost by the network manager. During the decision process, the OSI router calculates the total path cost of forwarding a packet along each possible path toward the destination. The total path cost is the sum of the costs of the circuits that make up the path. The least cost path is the one preferred by the OSI router.

When deciding between multiple paths to a destination, the router will choose the path that is assigned a lower path cost over one assigned a higher cost, even if the lower cost path is longer (refer to Figure 1-11).

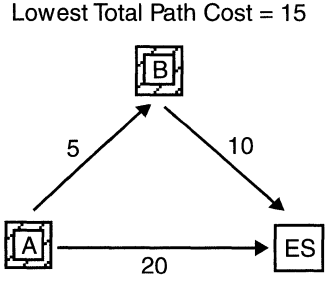

Figure 1-11. Example of Lowest Cost Path to a Destination

Once the router determines the lowest cost path to a destination, it stores the identity of the corresponding adjacent router into its forwarding database. The adjacent router is the next hop on the path toward the destination.

The router executes the decision process separately for each routing level. The router keeps separate forwarding databases for L1 and L2 routing. The L1 link state database is used to calculate the L1 forwarding database, which describes the shortest paths to destination systems located in the same area. If a router also routes L2 traffic, it uses its L2 link state database to create an L2 forwarding database, which describes the shortest paths to other destination areas.

Because the OSI router's link state database is updated every time there is a change to the network, the router can base its routing decisions on the most current network topology.

Note: When you configure the Wellfleet OSI router, you assign a cost metric to each OSI interface. So, if you want to limit the use of a certain low-speed interface, you can assign it a high cost. See the section "Editing OSI Interface Parameters" for instructions.

Forwarding Process

The OSI router performs the forwarding process after it receives a packet. First, it examines the destination address contained in the packet to determine if the packet requires L1 routing or L2 routing. It then refers to the corresponding forwarding database for information on where to forward the packet.

If the router is an L1 router and the packet's destination address is within the local area, the router checks its L1 forwarding database and forwards the packet to the next hop along the path to the destination. If the destination address is not local, the router checks its forwarding database for the location of the nearest L1/L2 router in the area. It then forwards the packet to the next hop along that path.

When the L1/L2 router receives the packet, it checks its L2 forwarding database to see which L1/L2 router is the next hop on the path to the destination area. It then forwards the packet to that L1/L2 router. The packet will continue to be forwarded between L1/L2 routers until it arrives at its destination area, at which point it will be routed (using L1 routing) to its destination system.

The Wellfleet OSI router also supports source routing and record route options. That is, if a packet has a statically entered path in the optional field of the packet header, the router forwards the packet toward the next hop. The record route function records the path(s) followed by a packet as it traverses a series of routers.

OSI Routing Protocols

This section summarizes the following OSI routing protocols the Wellfleet OSI router uses at the networking level:

- *ISO 8473 Connectionless-mode Network Service Protocol (CLNS),* which defines the OSI router data packet format procedures for the connectionless transmission of data and control information.
- *ISO 9542 End System to Intermediate System Routing Exchange Protocol,* which defines how end systems and intermediate systems exchange configuration and routing information to facilitate the routing and relaying functions of the network layer.
- *ISO 10589 Intermediate System to Intermediate System Routing Exchange Protocol*, which defines how L1 and L2 routing works.

Connectionless-mode Network Service Protocol

Connectionless-mode Network Service Protocol (ISO 8473) is the network layer protocol that specifies the procedures for the connectionless transmission of data and control information from one network system to a peer network system using CLNP packets.

Each CLNP packet that an OSI router receives is processed independently and does not require any previously established

network connection. A router bases its decision on how to process a CLNP packet solely on the information found in the packet header. That is, the packet header tells the router if the packet has reached its destination or if the packet requires additional processing.

If the size of the initial packet is greater than the maximum size supported by the network, CLNP allows a packet to be partitioned into two or more new packets (segments). The values contained in the header fields of the segmented packets are identical to those contained in the original packet (except for the segment length and checksum fields). Once the packet is partitioned, each packet segment is then sent out onto the network. When all of the packet segments finally arrive at the destination system, the system reconstructs the original packet before sending it up to the next layer for further processing.

In order to control data misdirection and congestion throughout the network, CLNP includes a lifetime control function. The originating system can assign a specific lifetime value (in units of 500 milliseconds) to the lifetime field of the packet header before the packet is sent out onto the network. Every system that receives the packet decrements its lifetime. If the lifetime value reaches zero before the packet reaches its destination system, the packet is dropped.

A packet will also be discarded if its checksum is incorrect, or if the network is congested and it cannot be processed, or if the destination address is unknown. CLNP includes an error reporting option that, when enabled, sends an error report data packet back to the originating system whenever a data packet is lost or discarded.

End System to Intermediate System Routing Exchange Protocol

The End System to Intermediate System Routing Exchange Protocol (ISO 9542) defines the way in which end systems and intermediate systems (routers) that are attached to the same subnetwork exchange configuration and routing information. Note that this protocol does not address communication issues between routers.

Configuration Reporting

The ISO 9542 configuration report function allows end systems and routers that are attached to the same physical network (subnetwork) to dynamically discover each other's identity by periodically generating and exchanging *hello* packets. The hello packet exchange process tells the router which NSAPs it can access.

End systems generate hello packets that contain the end system's subnetwork address, and specify which NSAPs the end system services. When a router receives an end system hello packet, the router extracts the configuration information from the packet (matching the subnetwork address with the corresponding NSAPs) and stores it in its routing information base. Routers generate hello packets that contain the router's own subnetwork address. When an end system receives a router hello packet, the end system extracts the router's subnetwork address and stores it in its own routing information base.

Two types of timers control how often hello packets are exchanged: a configuration timer and a holding timer. The configuration timer, which is maintained by each individual system, determines how often a system reports its availability or any change in its configuration to the other systems attached to the same subnetwork. The holding timer, which is a value set by the originating system, is contained in the holding time field of a hello packet. It specifies how long a receiving system should retain the configuration information before it is flushed from the routing information base.

Route Redirecting

The ISO 9542 route redirection function allows routers to inform end systems of the most desirable route to a particular destination—either through a different router or directly to the destination end system, if it is attached to the same subnetwork. After the router forwards a data packet to the next hop toward the destination end system, the router checks if a more direct route exists.

That is, the router determines if the next hop is

- The destination system, and if it is attached to the same subnetwork as the originating system (Figure 1-12, Example 1)
- Another router that is connected to the same subnetwork as the originating end system (Figure 1-12, Example 2)

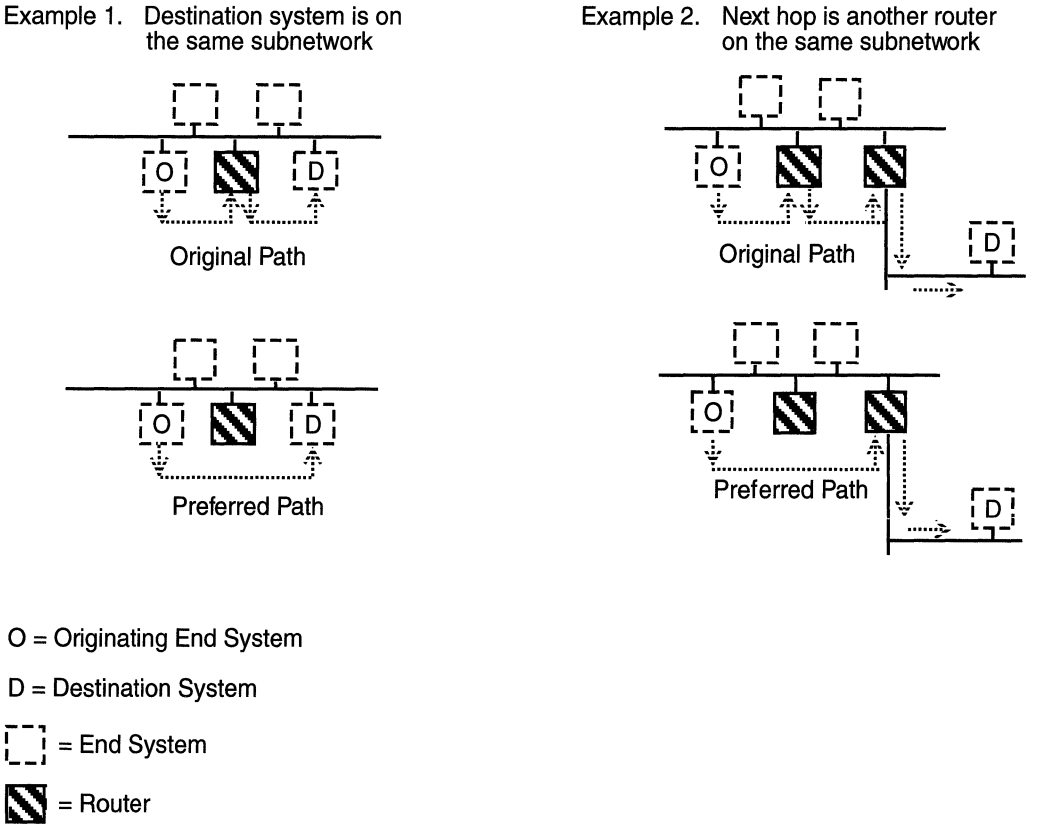

Figure 1-12. Route Redirecting

If the next hop is either a destination system or another router on the same subnetwork, then there is a better path (one that does not

traverse the router) to the destination. The router then constructs a redirect packet (RD packet), which contains the following information:

- Destination address of the original packet
- Subnetwork address of the preferred next hop
- Network entity title of the next hop, unless it is the destination end system
- Holding Timer and Maintenance, Security, and Priority options

The router sends the RD packet back to the originating end system, which has the option to use the RD packet to update its routing information base with the more direct route.

Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol

The Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol (ISO 10589) defines the way in which intermediate systems (routers) within a routing domain exchange configuration and routing information. It works in conjunction with ISO 8473 and ISO 9542 to define how routers can communicate and route packets within and between areas.

Intra-Domain Routing

This protocol is designed to function within a single routing domain. The domain may consist of various types of subnetworks that have been administratively divided into separate routing areas.

Under this protocol, L1 routers keep track of the routing that occurs within their own areas. Thus, each L1 router must know the topology of its local area, including the location of all other routers and end systems (from LSP and hello packets that are exchanged throughout the network). Note that an L1 router does not need to know the identity of those systems residing outside of its local area, because it forwards all packets destined for other areas to the nearest L1/L2 router.

Similarly, each L1/L2 router must know the topology of the other L1/L2 routers located in the domain and the addresses that are reachable through each L1/L2 router (again, through LSPs and hello packets). The set of all L1/L2 routers is a type of "backbone" network for interconnecting all areas in the domain. Note that an L1/L2 router that supports L1 routing also needs to know the topology within its local area.

For example, when an L1 router receives a data packet, it compares the destination area address in the packet with its own area address. If the destination area address is different, then the packet is destined for another area and needs to be routed using L2 routing. The router forwards the packet to the nearest L1/L2 router in its own area, regardless of what the destination area is. The L1/L2 router then forwards the packet to a peer L1/L2 router that is the next hop on the path to the destination system. The packet will continue to be routed between L1/L2 routers until it reaches its destination area, where it will be forwarded (using L1 routing) to the destination end system.

Inter-Domain Routing

Inter-domain routing is possible when paths to other domains are statically defined. In order for inter-domain routing to occur, you must manually enter the set of reachable address prefixes into each L1/L2 router (called a bordering router) that is linked to an external domain. The address prefixes describe which NSAP addresses are reachable over that L1/L2 router's external link. The next time the L1/L2 routers in the domain exchange LSPs, they become aware of the existence of the reachable external addresses and update their link state databases with this information.

As traffic is routed throughout the network, any packets with destination addresses that match the statistically defined reachable address prefixes are directed to the bordering router. The bordering router then transmits the packet out of the domain. The other domain must assume responsibility for routing the packet to its final destination (refer to Figure 1-13).

Note: Inter-domain routing is strictly between L1/L2 routers.

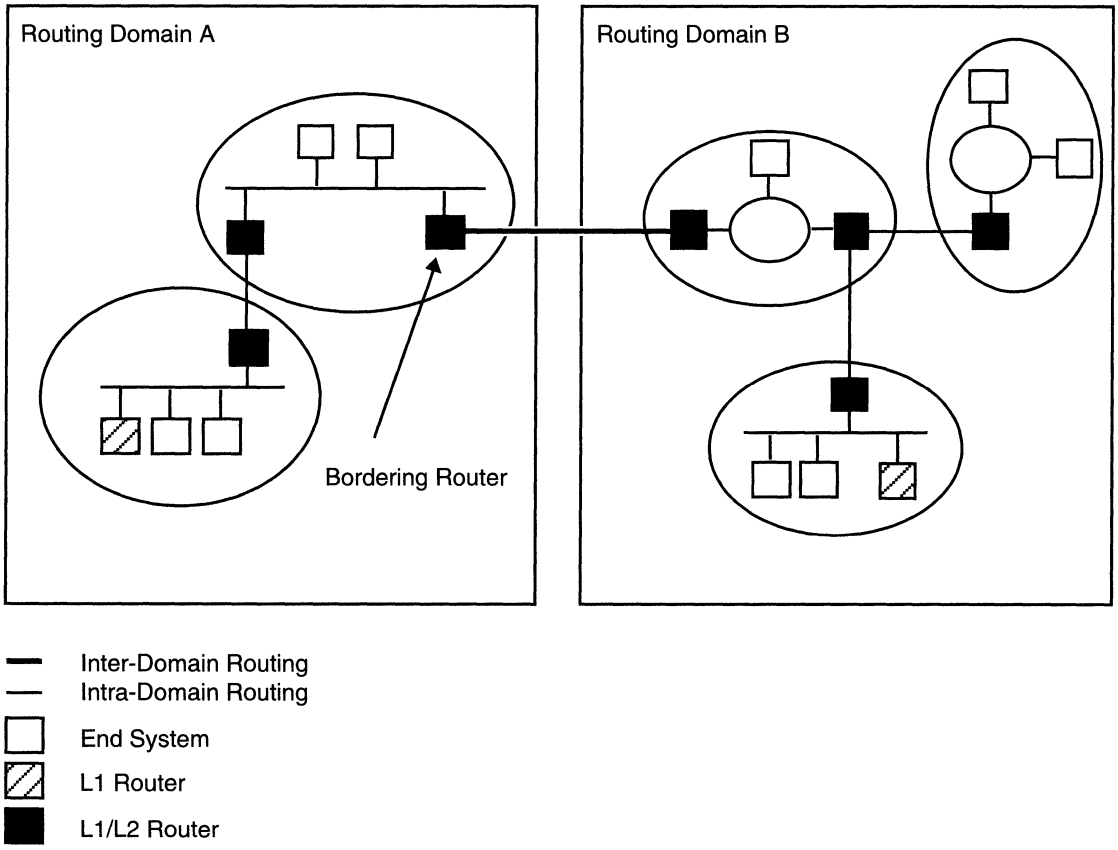

Figure 1-13. Static Inter-Domain Routing

Chapter 2
OSI Implementation Notes

This chapter contains some basic guidelines on adding Wellfleet OSI routers to your network. It also addresses special configuration features that may match your network requirements.

Configuring Area Address Aliases

You configure area address aliases if you plan on dividing a large area into two or more smaller areas. An area address alias is a second (or third) area address configured for systems residing in a single area. When used appropriately, the area address alias feature can make network management easier.

For example, consider the OSI network shown in Figure 2-1. All routers and end systems belong to the area XY. This area had originally been assigned the area address 123. Sometime in the near future, the network administrator plans to divide the area into two smaller, more manageable areas: Area X and Area Y.

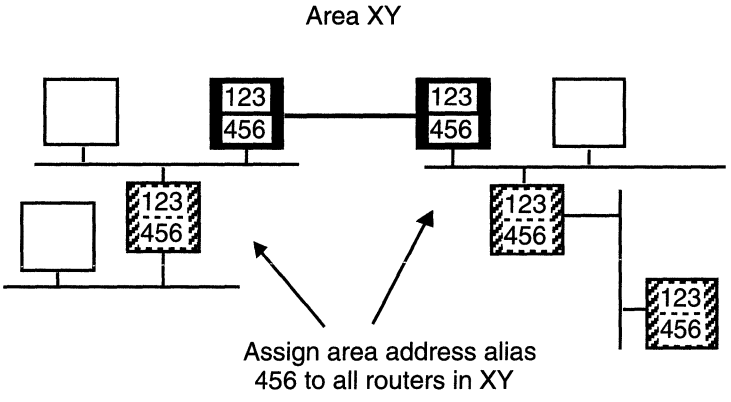

Figure 2-1. Original Area Addresses for Area XY

Taking advantage of the area address alias feature, the administrator

1. Assigns the area address alias 456 to all routers within area XY (Figure 2-2).

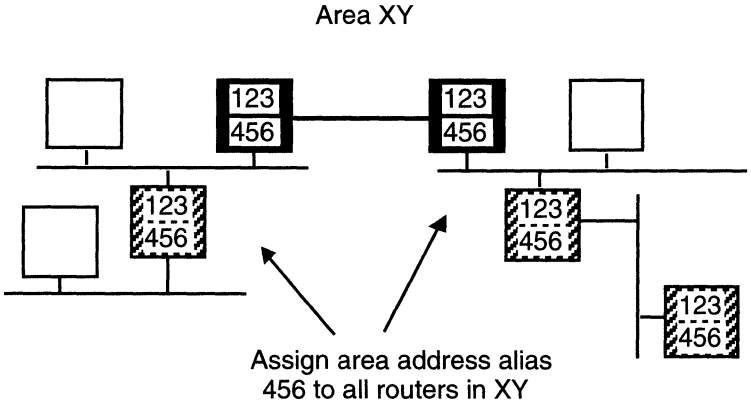

Figure 2-2. Assign Area Address Alias 456 to all Routers in Area XY

2. Assigns the area address alias 456 to those end systems that will eventually belong to area Y when area XY is divided (Figure 2-3).

   All end systems are still able to communicate using the originally assigned area address 123, so this can be done gradually.

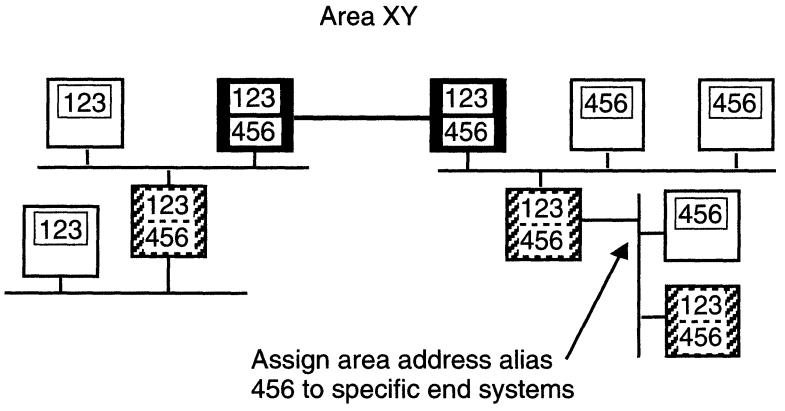

Figure 2-3. Assign Area Address 456 to Specific End Systems

3. Finally, in order to divide Area AB completely, the manager simply deletes area address alias 456 from those routers that will remain

in area X, and deletes area address 123 from those routers and end systems that will be part of the new area Y.

Because the end systems in both area X and area Y have already been assigned corresponding area addresses, they do not have to be reconfigured, and the division is complete (Figure 2-4).

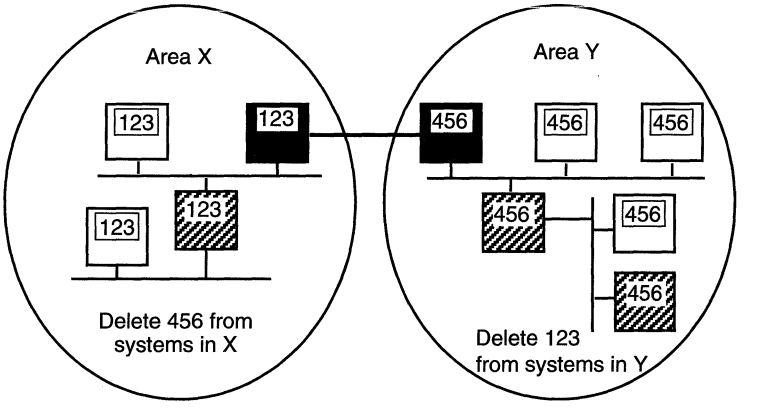

Figure 2-4. Divide Area AB into Area X and Area Y

See "Editing OSI Interface Parameters" for instructions on how to configure Area Address Alias parameters.

Area Partitions

An area is partitioned when all nodes in the area cannot communicate with each other either directly or indirectly at Level 1. Partitions happen through improper network design or when one or more links fail in an area. (Area partition repair as specified in *ISO 10589 Intermediate System to Intermediate System Routing Exchange Protocol* is currently not supported.) See Chapter 1 for a discussion of the role of areas and Level 1 and 2 routing in OSI network organization.

Figure 2-5 demonstrates an improper network design.

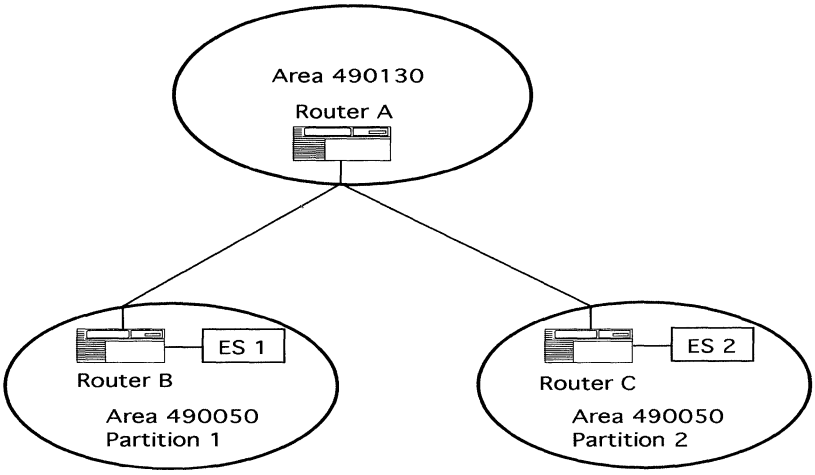

Figure 2-5. An Area Partition Due to Improper Network Design

In this hub and spoke topology, Router A in Area 490130 recognizes two separate routes to Area 490050. Routers B and C do not have a Level 1 link between them; therefore, each is in a different partition of the area. They cannot exchange Level 1 information and neither one knows about end systems in the other partition. If Router A sends a packet to an end system in Area 490050, it may choose Router B in Partition 1 as the lowest cost route. If the packet is intended for an end system attached to Router C, Router B will reject the packet, because it does not know about the end system in Partition 2.

One solution is to modify the topology by creating a link between Routers B and C. Another solution is to create another area for Router C or B. The routers could then use Level 2 routing to communicate.

Configuring Static External Adjacencies

You configure static external adjacencies if you want to route traffic between domains. To do so, you must

- Configure external routing support on each interface that connects the L1/L2 router to an external domain.

You do this by setting the Routing Level parameter to a selection that includes external (for example, Level 2 & External). See the section "Editing OSI Interface Parameters" in Chapter 3 for details.

- Manually enter the set of reachable address prefixes into each L1/L2 bordering router that is linked to an external domain.

  The address prefixes describe which NSAP addresses are reachable over that L1/L2 router's external link. See the section "Configuring Static External Address Adjacencies" in Chapter 3 for details.

Configuring OSI over DDN X.25

The X.25 Defense Data Network (DDN) provides end-to-end connectivity between a router and remote Data-Circuit Terminating Equipment (DTE) that supports X.25 DDN Standard Service. Internet Protocol (IP) uses DDN service to transmit IP datagrams over the X.25 network.

Each network interface that connects to the X.25 network uses an X.121 address. (For additional information about the X.25 network and X.121 addresses, see the *Customizing X.25 Services* guide.)

If you want to run OSI over DDN X.25, you must

- Configure IP over an X.25 DDN circuit. See the *Configuring Wellfleet Routers* guide for details.
- Convert the remote IP address to an X.121 address. You use the converted address as the Subnetwork Point of Attachment (SNPA) for a static end system adjacency or a static external address adjacency. (See Chapter 3 for details on the SNPA parameter and Appendix A for details on address conversion.)

Configuring OSI over Frame Relay

Frame relay is a high-speed, shared-bandwidth, wide-area networking protocol. Frame relay performs only basic processing on each packet, allowing frame relay networks to operate at high speeds with few delays but with little error detection. See *Customizing Frame Relay Services* for general information about the protocol.

Running OSI over frame relay requires you to consider the following issues:

- Configuration Overview
- Frame Relay Circuit Modes
  - Direct Access
  - Group Access
  - Hybrid
  - Mixed Access
- Topology
  - Full Mesh
  - Partial Mesh
  - Area Partitions
- Route Redirecting
- Designated Router Selection
- IS Neighbor Detection
- Circuits Per Slot

These issues are discussed in the following sections.

Configuration Overview

If you want to run OSI over frame relay, you must:

1. Configure a frame relay circuit using Site Manager.

   See *Configuring Wellfleet Routers* for initial frame relay configuration information.

2. Configure OSI to operate over frame relay.

   See *Configuring Wellfleet Routers* for initial OSI configuration information.

3. Customize frame relay and OSI depending on your network's circuit mode and topology.

   See the following sections for information on OSI over frame relay based on the circuit mode and topology of your network.

4. In direct access mode, repeat Steps 1 to 3 for each permanent virtual circuit (PVC). See the "Direct Access" section.

Frame Relay Circuit Modes

Our implementation of OSI over Frame Relay operates as subnetworks in either of these two types of Intermediate System to Intermediate System (IS-IS) operation modes:

- Point-to-Point
- Broadcast

The OSI router implements these IS-IS operation modes over frame relay circuits. Table 2-1 lists the frame relay modes used for IS-IS operations.

Table 2-1. Frame Relay Modes Used for OSI IS-IS Operations

Frame Relay Mode	IS-IS Operation Mode
Direct access	Point-to-Point
Group access	Broadcast
Hybrid	Broadcast

Direct Access

In direct access mode, OSI treats a PVC as a point-to-point connection. OSI views each PVC as an individual network interface Figure 2-6.

OSI point-to-point operation over frame relay utilizes circuit bandwidth more efficiently than OSI broadcast operation. It also complies with the ISO standards for point-to-point operation. However, point-to-point operation uses proportionally more memory resources on the router per PVC than broadcast operation.

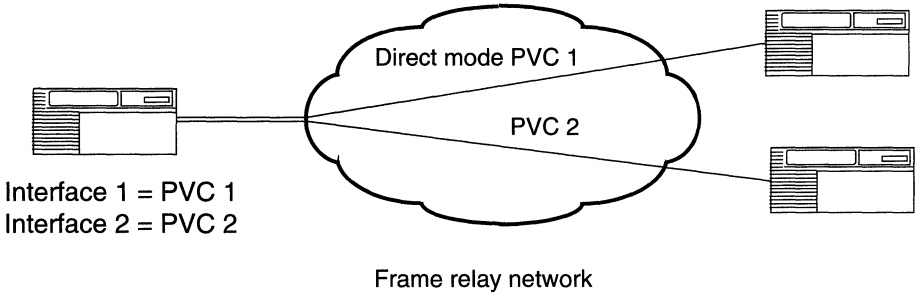

Figure 2-6. Frame Relay Direct Access Mode

If you use direct access mode, you must configure each frame relay PVC manually and configure the OSI protocol to run over it. The OSI router will treat each PVC as a separate OSI interface. See *Customizing Frame Relay Services* for information about configuring PVCs.

Group Access

In group access mode, OSI treats each frame relay network interface as a single access point to the subnetwork (Figure 2-7). DLCIs on the subnetwork are treated like MAC addresses on actual broadcast media. When an OSI packet needs to be "broadcast" on a particular frame relay circuit, it is sent over all known PVCs on that circuit. OSI assumes that all systems on the subnetwork will receive a broadcast packet.

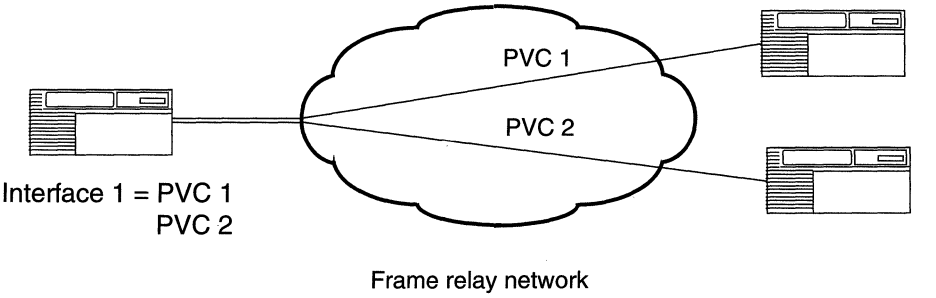

Figure 2-7. Frame Relay Group Access Mode

Group access works best in either full mesh environments or partial mesh environments set up in a hub and spoke topology where communication between systems that are not directly connected to one another go through the hub.

OSI over frame relay in group mode has several unique features and limitations to consider in planning your network as noted in the following sections.

Hybrid

For OSI, hybrid frame relay circuit mode is the same as group access.

Mixed Access

It is possible to mix both group and direct access mode in a configuration as long as the group access restrictions are not violated.

Figure 2-8 shows a designated router with direct access on PVC1 and group access on PVC2.

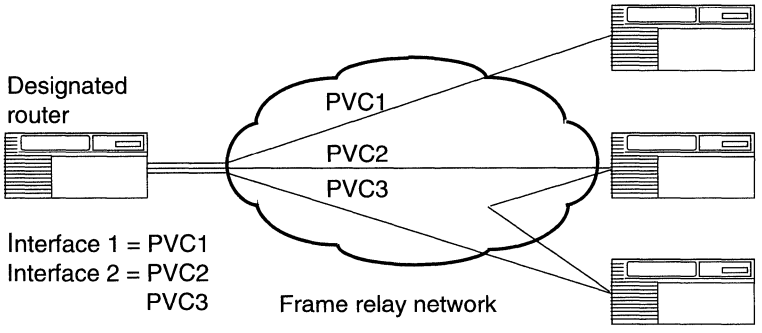

Figure 2-8. Frame Relay Mixed Access Modes (Direct and Group)

Topology

Consider the following issues in implementing OSI over group access mode frame relay circuits in a full or partial mesh topology.

Full Mesh Topology

Using group access mode in a full mesh topology models the frame relay network as a LAN (Figure 2-9).

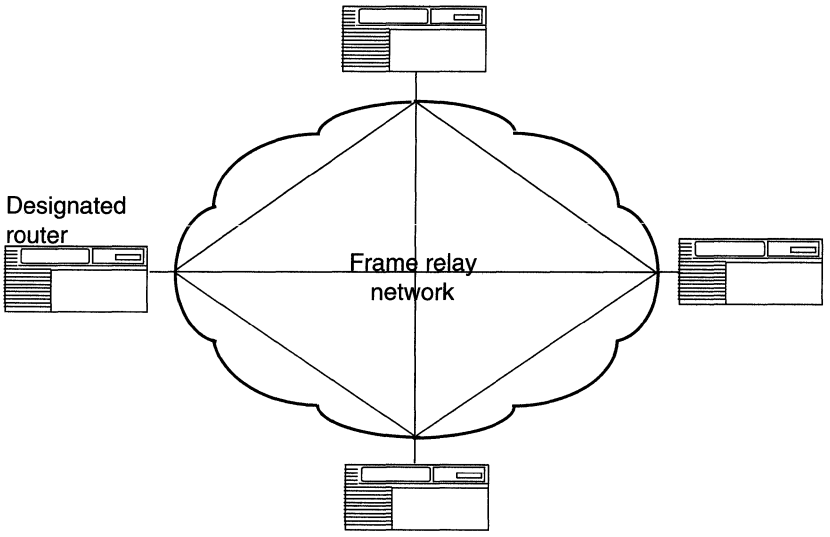

Figure 2-9. Full Mesh Topology

If a router fails or the link to the frame relay network fails, the topology remains full mesh. If a PVC fails, however, the network changes from a full mesh to a partial mesh topology. This can introduce connectivity problems in the resulting network. For example, if a non-designated router loses a PVC to the designated router, it will attempt to elect another designated router. Since the other systems are still in contact with the active designated router, the link state databases of the routers will not be synchronized, which could result in connectivity problems between systems.

Partial Mesh Topology

If you use a partial mesh topology with group access mode you need to arrange the network in a hub and spoke topology with the designated router as the hub (Figure 2-10).

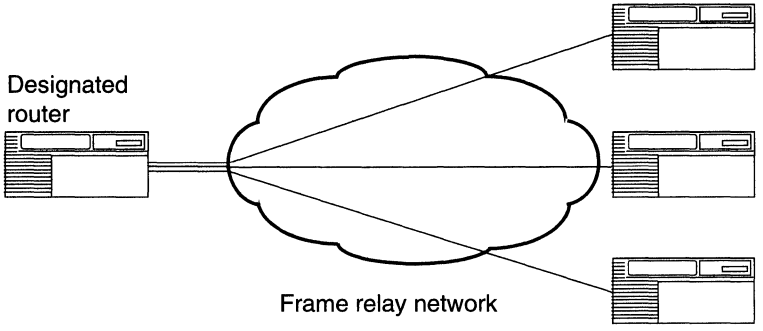

Figure 2-10. Partial Mesh in Hub and Spoke Topology

A PVC that goes down will only cause communication failure between the hub (designated router) and a spoke. However, in a partial mesh topology, losing the hub router causes all communication links on the subnetwork to fail.

Route Redirecting

When you configure OSI over frame relay, the Redirect Enable/Disable parameter appears on the OSI Interface Lists window. (See Chapter 3 for the Redirect parameter description.) Redirects specify whether an OSI interface sends a redirect packet (ES-IS message) back to the originating system informing it of a more direct path to a destination system. This function is valid in full mesh topology because all systems can communicate directly.

Redirects are invalid when running OSI over frame relay in group access mode in a hub and spoke topology, because the spoke systems cannot communicate directly with each other.

Set the Redirect Enable/Disable parameter to Disabled when operating OSI over frame relay in group mode in a hub and spoke topology. Accept the default value of Enabled in full mesh topologies.

Designated Router Selection

OSI over group access frame relay uses the highest System ID for designated router selection. This feature is needed to break a tie when the designated router priority is the same for two or more routers on a subnetwork. Normally, the IS-IS specification in OSI calls for the comparison of local SNPA addresses in breaking ties in designated router elections. Frame relay interfaces do not have a local SNPA address. See the "Update Process" section of Chapter 1 for more information about designated routers.

IS Neighbor Detection

Two-way connectivity checking in adjacency establishment does not operate in OSI over group mode frame relay. Normally, two Intermediate Systems on an OSI broadcast subnetwork report each other in their LAN hello packets. An IS must see its own subnet address in a LAN hello packet from a neighbor to form an active adjacency. A local subnet address does not exist on a frame relay interface, so this function is not used.

Circuits Per Slot

A maximum of 48 OSI interfaces per slot are supported in this release.

Chapter 3
Editing OSI Parameters

Once you enable an OSI interface, you can use the Site Manager to edit OSI parameters and customize OSI services.

This chapter describes how to

- Edit OSI parameters.
- Add, edit, or delete a static route, static adjacency, or the DECnet IV to V Transition feature.
- Delete OSI globally from the Wellfleet router.

Note: The instructions in this chapter assume that you have already configured at least one OSI interface on the router. If you have not yet configured an OSI interface, or want to add additional OSI interfaces, see the *Configuring Wellfleet Routers* guide for instructions.

Accessing OSI Parameters

You access all OSI parameters from the Wellfleet Configuration Manager window shown in Figure 3-1. Refer to the *Configuring Wellfleet Routers* guide for details on accessing this window.

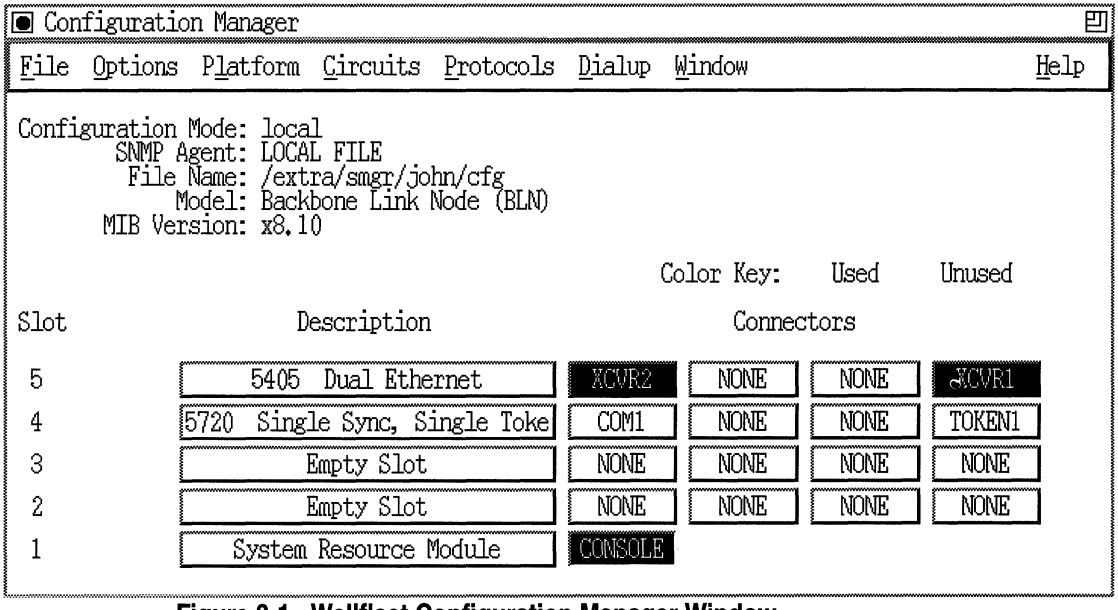

Figure 3-1. Wellfleet Configuration Manager Window

To customize the router software for OSI services, you can edit any of these OSI parameters:

- Global
- Interface
- Static adjacency
- Static route
- DECnet IV to V Transition

Note: You can only access the DECnet IV to V Transition parameters using OSI. To enable the DECnet IV to V Transition feature, you must configure at least one DECnet interface on the router. (See *Configuring Wellfleet Routers* for details.)

For each OSI parameter, this chapter describes the Wellfleet default setting, all valid setting options, the parameter function, instructions for setting the parameter, and the MIB object ID. See the section that applies to the type of parameter you want to edit.

Editing OSI Global Parameters

To edit the OSI global parameters:

1. Select the Protocols→OSI→Global option from the Wellfleet Configuration Manager window (refer to Figure 3-1). The Edit OSI Global Parameters window appears (Figure 3-2).

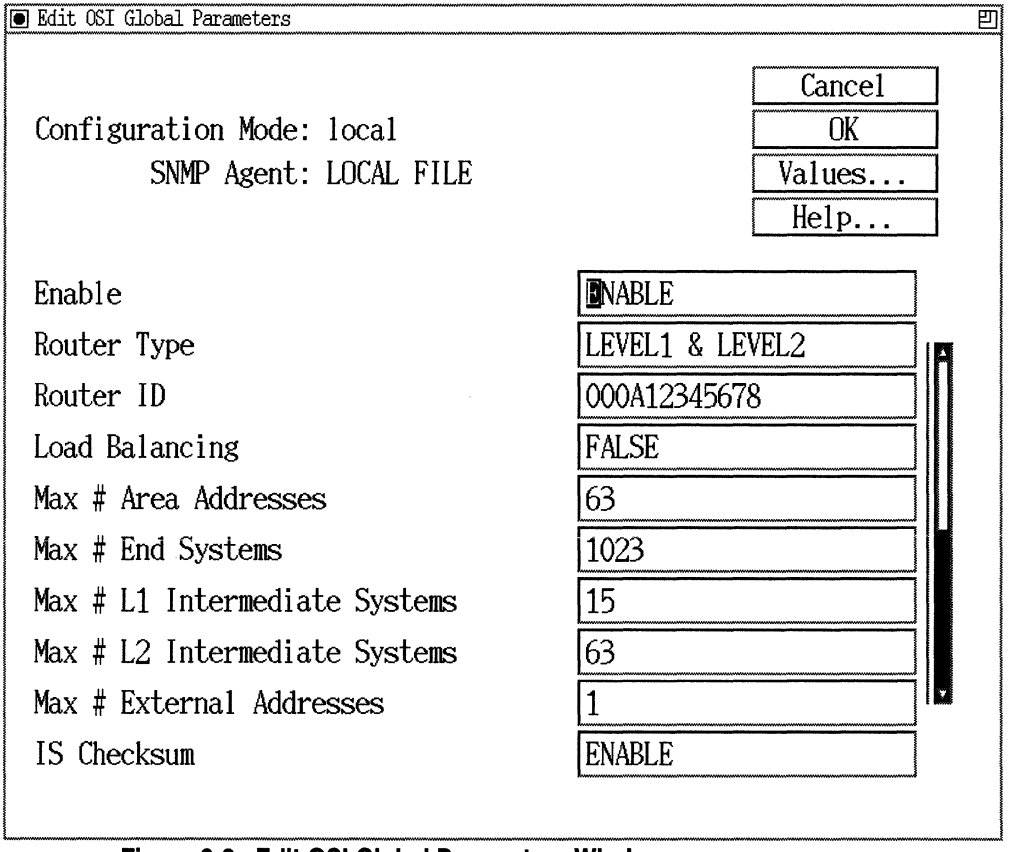

Figure 3-2. Edit OSI Global Parameters Window

2. Edit the parameters, using the descriptions in the next section as a guide.

3. Click on OK to save your changes and exit the window. Site Manager returns you to the Wellfleet Configuration Manager window.

OSI Global Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the Edit OSI Global Parameters window (refer to Figure 3-2):

Parameter:	**Enable**
Default:	Enable
Options:	Enable | Disable
Function:	Enables or Disables OSI routing on the router.
Instructions:	Disable only if you want to globally disable OSI routing on all interfaces on which it is configured.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.2

Parameter:	**Router Type**
Default:	Level 1 and Level 2
Options:	Level 1 | Level 1 and Level 2
Function:	Specifies whether the router functions as an L1 router (Level 1) or an L1/L2 router (Level 1 and Level 2). An L1 router can support only Level 1 routing within its own area. An L1/L2 router can support Level 1 routing, Level 2 routing between areas, and external routing between domains. You can further define the type of traffic that router supports by editing interface parameters. For example, if you want a certain interface to route only Level 2 traffic, then you can designate the individual interface as an L2 interface (see "Editing OSI Interface Parameters" for instructions).
Instructions:	Select the appropriate Router Type.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.5

Note: To support routing between areas, you must specify at least one L1/L2 router per area. However, each L1/L2 router can service only a single area.

Parameter:	**Router ID**
Default:	Variable
Options:	Any valid 6-byte System ID
Function:	Identifies the router within its local area. The System ID is the ID portion of the router's NSAP address. (See the section "OSI Network Addressing" for more information.)
Instructions:	Enter a 6-byte System ID in hexadecimal format. If the System ID is not 6 bytes, add leading zeroes. Since every router in a domain must have a unique System ID, using a router's MAC address for its System ID ensures this requirement.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.6

Parameter:	**Load Balancing**
Default:	False
Options:	True | False
Function:	Specifies if the router should balance the data traffic flow over two equal cost paths to the same destination. Load balancing keeps one path from becoming overloaded, while taking advantage of the bandwidth available on an additional path. The paths must be of equal cost.
Instructions:	To enable load balancing, reset this parameter to True.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.7

Parameter:	**Max # Area Addresses**
Default:	63
Range:	1 to 1000
Function:	Specifies the maximum number of local areas in the domain.
Instructions:	Unless there are more than 63 areas in the router's domain, accept the default value 63.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.8

Parameter:	**Max # End Systems**
Default:	512
Range:	1 to 4000
Function:	Specifies the maximum number of end systems contained within this local area.
Instructions:	Unless there are more than 1023 end systems in the local area, accept the default value 1023.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.9

Parameter:	**Max # L1 Intermediate Systems**
Default:	15
Range:	1 to 1000
Function:	Specifies the maximum number of Level 1 OSI routers contained within this local area.
Instructions:	Unless there are more than 15 Level 1 OSI routers residing within this local area, accept the default value 15.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.10

Parameter:	**Max # L2 Intermediate Systems**
Default:	63
Range:	1 to 1000
Function:	Specifies the maximum number of L1/ L2 OSI routers contained within this local area.
Instructions:	Unless there are more than 63 L1/L2 OSI routers residing within this local area, accept the default value 63.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.11

Parameter:	**Max # External Addresses**
Default:	1
Range:	1 to 500
Function:	Specifies the number of external domain addresses imported into the local domain.
Instructions:	If you do not have any links to external domains, then accept the default value 1. Otherwise, enter the maximum number of external domains linked to the local domain.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.12

Parameter:	**IS Checksum**
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables the generation of a non-zero checksum for IS packets.
Instructions:	To allow checksum processing, accept the default value Enable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.13

Parameter:	**L1 LSP Password**
Default:	None
Options:	Any text string, 8 characters maximum
Function:	Assigns a password to the Level 1 link state packets (LSP), partial sequence number packets (PSNP), and complete sequence number packets (CSNP) that the router (L1 or L1/L2) generates and accepts. The router uses LSP information to make routing decisions, and PSNP and CSNP information to make sure that its LSP database is up-to-date. The router will route only data though those routers with which it has exchanged LSPs. Therefore, the L1 LSP password is a security device. You assign identical L1 LSP passwords to all routers located in the same area through which you wish to route data. When the OSI router floods Level 1 LSPs through the area, only those routers that have been assigned the same password accept the LSPs.
Instructions:	If you do not want to assign a L1 LSP password to this router, then leave this field blank. If you assign a L1 LSP password to this router, then you must assign the same L1 LSP password to every router in the area with which this router communicates.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.14

Parameter: **L2 LSP Password**

Default: None

Options: Any text string, 8 characters maximum

Function: Assigns a password to the Level 2 link state packets (LSP), partial sequence number packets (PSNP), and complete sequence number packets (CSNP) that the router (L1/L2) generates and accepts.

The router uses LSP information to make routing decisions, and PSNP and CSNP information to make sure that its LSP database is up-to-date. The router will route data only though those routers with which it has exchanged LSPs. Therefore, the L2 LSP password is a security device. You assign identical L2 LSP passwords to all routers located in the domain through which you wish to route data. When the OSI router floods Level 2 LSPs through the area, only those routers that have been assigned the same password accept the LSPs.

Instructions: If you do not want to assign a L2 LSP password to this router, then leave this field blank. If you assign a L2 LSP password to this router, then you must assign the same L2 LSP password to every router in the domain with which this router communicates.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.1.15

Note: If you set the Router Type for this router to Level 1 only, then this parameter is ignored.

Parameter: **Area Address**

Default: 490040

Options: Any area address entered in hexadecimal format with a 3-byte minimum and 13-byte maximum length.

Function: Identifies the local area in the routing domain where the router resides.

Instructions: If you have registered your OSI network with an addressing authority, then the area address will also reflect the location of the router in the global addressing domain. Enter the entire area address portion of the NSAP address allocated to your network as follows:

— Check with your administrative authority to determine the NSAP addresses that have been allocated to your OSI network.

— Enter the entire area address portion of the NSAP address that reflects the location of the router—including the routing domain and area portions that identify where in the local network the router resides. Either you or your administrative authority should provide the identifiers for the local routing domain and area portions of the address.

— If you have *not* registered your OSI network with an addressing authority, then, you can accept the default area address of 490040.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.1.16

Note: You must assign the same Area Address to all routers residing in the same local area. You must assign different area addresses to routers that reside in different areas.

Parameter:	**Area Address Alias 1 (hex)**
Default:	None
Options:	Any valid area address.
Function:	Assigns the first area address alias to the router. An area address alias is a different area address that is assigned to the same router. For the DECnet IV to V Transition feature, the area address alias defines the Phase IV prefix and Phase IV area fields of the Phase IV–compatible address.
Instructions:	Enter the area address alias in hexadecimal format. For the DECnet IV to V Transition feature, enter the Phase IV prefix (from 1 to 9 bytes) followed by 2 bytes of the Phase IV area address. Otherwise, leave this field blank.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.17

Parameter:	**Area Address Alias 2**
Default:	None
Options:	Any valid area address
Function:	Assigns the second area address alias to the router.
Instructions:	Enter the area address alias in hexadecimal format. Otherwise, leave this field blank.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.18

Parameter:	**Max # Learned End Systems**
Default:	1024
Range:	1 to 4000
Function:	Specifies the maximum number of end systems per slot that the router can learn about dynamically through the exchange of hello packets.
Instructions:	Unless the area in which this router resides contains more than 1024 end systems, accept the default value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.28

Parameter:	**Max # Learned L1 Intermediate Systems**
Default:	64
Range:	1 to 4000
Function:	Specifies the maximum number of L1 routers per slot that this router can learn about dynamically through the exchange of hello packets.
Instructions:	Unless the area in which this router resides contains more than 64 Level 1 intermediate systems, accept the default value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.29

Parameter:	**Max # Learned L2 Intermediate Systems**
Default:	64
Range:	1 to 4000
Function:	Specifies the maximum number of L2 routers per slot that the router can learn about dynamically through the exchange of hello packets.
Instructions:	Unless the domain in which this router resides contains more than 64 L2 routers, accept the default value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.30

Parameter:	**CLNP Source Route Support**
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables the processing of source routing options in CLNP packets.
Instructions:	Set to Disable if this router requires GOSIP v2 support.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.38

Editing OSI Interface Parameters

To edit an OSI interface:

1. Select the Protocols→OSI→Interfaces option from the Wellfleet Configuration Manager window (refer to Figure 3-1). The OSI Interface Lists window appears (Figure 3-3). It displays all interfaces on which OSI is enabled.

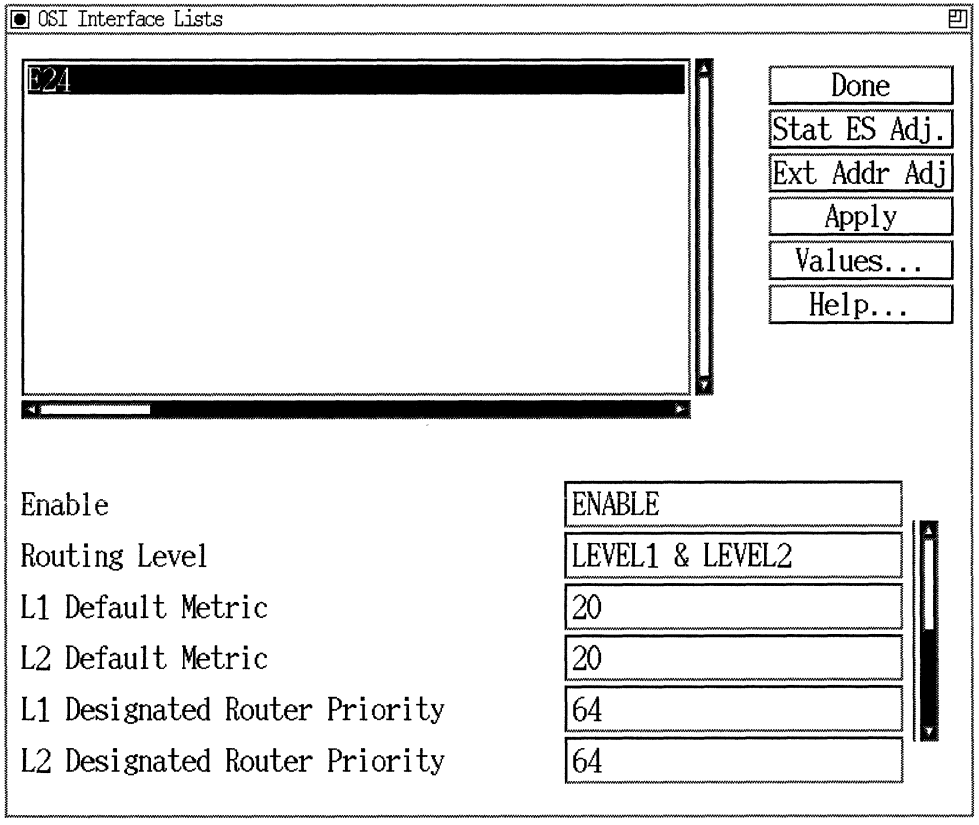

Figure 3-3. OSI Interface Lists Window

2. Click on an interface to select it.

3. Edit the parameters, using the descriptions in the next section as a guide. Use the scroll bar to scroll through the list of parameters for the interface.
4. Implement your changes by clicking on Apply.
5. Exit the window by clicking on Done. Site Manager returns you to the Wellfleet Configuration Manager window.

Note: When you reconfigure an interface in dynamic configuration mode, OSI restarts on that interface.

OSI Interface Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the OSI Interface Lists window (refer to Figure 3-3):

Parameter:	**Enable**
Default:	Enable
Options:	Enable | Disable
Function:	Enables OSI routing on this interface.
Instructions:	Disable only if you want to disable OSI routing on this interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.2

Parameter:	**Routing Level**
Default:	Level 1 and Level 2
Options:	Level 1 | Level 2 | Level 1 and Level 2 | External | L2 External | L1 and L2 External | ES-IS-only
Function:	Specifies the type of traffic that is routed over this interface.
Instructions:	Select the Routing Level that matches the level of traffic you want to route on this interface. Note that if you set the global Router Type parameter to Level 1, then you can only route Level 1 traffic on this interface. See "Editing OSI Global Parameters" for instructions on setting the global Router Type parameter. If this interface will be used to route traffic between domains, then select an option that includes External. In addition, you must statically define the external adjacencies with which this router communicates. See "Configuring Static External Address Adjacencies" for instructions.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.5

Parameter:	**L1 Default Metric**
Default:	20
Range:	1 to 63
Function:	Specifies the default metric (relative cost) of routing Level 1 traffic over this interface. OSI determines path costs on the basis of the *sum* of the individual *circuit costs*. The cost that you assign to a particular circuit typically reflects the speed of the transmission medium. Low costs reflect high-speed media, while high costs reflect slower media. Refer to Table 3-1 for a list of suggested OSI circuit costs. The OSI router always selects the interfaces with the lowest cost when defining a path, so assigning each interface a cost is, in effect, a way of assigning it a priority.
Instructions:	If you do not want to use this interface to route Level 1 traffic on a regular basis, assign it a high cost. Otherwise, accept the default 20.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.6

Table 3-1. Suggested OSI Circuit Cost Values

Speed	Cost	Speed	Cost
100 Mb/s	1	64 Kb/s	54
16 Mb/s	19	56 Kb/s	55
10 Mb/s	20	38.4 Kb/s	56
4 Mb/s	21	32 Kb/s	57
1.54 Mb/s	45	19.2 Kb/s	58
1.25 Mb/s	48	9.6 Kb/s	59
833 Kb/s	49	7.2 Kb/s	60
625 Kb/s	50	4.8 Kb/s	61
420 Kb/s	51	2.4 Kb/s	62
230.4 Kb/s	52	1.2 Kb/s	63
125 Kb/s	53		

Parameter: **L2 Default Metric**

Default: 20

Range: 1 to 63

Function: Specifies the relative cost of routing Level 2 traffic over this interface.

OSI determines path costs on the basis of the sum of the individual circuit costs. The cost that you assign to a particular circuit typically reflects the speed of the transmission medium. Low costs reflect high-speed media, while high costs reflect slower media. Refer to Table 3-1 for a list of suggested OSI circuit costs.

The OSI router always selects the interfaces with the lowest cost when defining a path, so assigning each interface a cost is, in effect, a way of assigning it a priority.

Instructions: If you do not want this interface to be used to route Level 2 traffic on a regular basis, assign it a high cost. Otherwise, accept the default 20.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.7

Parameter: L1 Designated Router Priority

Default: 64

Range: 1 to 127

Function: Can specify which L1 router becomes the L1 designated router for the LAN segment. (See the section entitled "Update Process" for more information about the designated router.)

You can control which L1 router becomes the L1 designated router for the LAN segment by assigning a priority value to each L1 router. Then, the L1 router assigned the highest priority becomes the L1 designated router for that LAN segment.

If all routers have the same priority, then the L1 router with the highest MAC address becomes the L1 designated router for the LAN segment.

Instructions: If you want this L1 router to become the L1 designated router for the LAN segment, then assign it the highest priority value among L1 routers on the LAN.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.8

Note: If the network is synchronous (for example, point-to-point or X.25), then the routers on each end of the connection must have different L1 DR Priority values or this parameter is ignored. This applies only to Wellfleet standard point-to-point and X.25 point-to-point service. It does not apply to a synchronous circuit running Point-to-Point Protocol (PPP) or X.25 Public Data Network (PDN) (or DDN) service.

Parameter:	**L2 Designated Router Priority**
Default:	64
Range:	1 to 127
Function:	Can specify which L2 router becomes the L2 designated router for the LAN segment. (See the section entitled "Update Process" for information on the pseudonode.) You can control which L2 router becomes the L2 designated router for the LAN segment by assigning a priority value to each L2 router. Then, the L2 router assigned the highest priority becomes the L2 designated router for that LAN segment. If all routers have the same priority, then the L2 router with the highest MAC address becomes the L2 designated router for the LAN segment.
Instructions:	If you want this L2 router to become the L2 designated router for the LAN segment, then assign it the highest priority value among L2 routers on the LAN.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.9

Parameter: **IIH Hello Timer**

Default: 8

Options: 2 | 4 | 8 | 15 | 30 | 60 | 120 | 300 | 600 | 1800 | 2400 | 3600

Function: Specifies in seconds how often other routers need to send Intermediate System Hello messages to this router. This router includes this value in the Intermediate System Hellos it sends to the other routers.

Instructions: Accept the default value, or select any valid option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.10

Parameter: **ISH Hello Timer**

Default: 30

Options: 2 | 4 | 8 | 15 | 30 | 60 | 120 | 300 | 600 | 1800 | 2400 | 3600

Function: Specifies the interval in seconds between LAN hello messages transmitted across the interface between a router (L1 or L1/L2) and an end system in the local area.

Instructions: Accept the default value, or select any valid option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.11

Parameter:	**ESH Configuration Time**
Default:	600
Options:	2 | 4 | 8 | 15 | 30 | 60 | 120 | 300 | 600 | 1800 | 2400 | 3600
Function:	Specifies in seconds how often end systems need to send system hello messages to this router. This value is included in the intermediate system hello messages the router sends to end systems.
Instructions:	Accept the default value, or select any valid option.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.12

Parameter:	**Circuit Password**
Default:	None
Range:	Any text string, 8 characters maximum
Function:	Assigns a password to the interface. A router will only route packets to those routers that have been assigned the same Circuit Password. The Circuit Password is carried to other routers when IS-IS hello packets are exchanged. If a router discovers that another router has a different password, it will not route traffic to that router. Therefore, to communicate, adjacent routers on either end of a point-to-point connection must have the same Circuit Password.
Instructions:	To assign a circuit password, enter a text string.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.3.13

Parameter: **IIH Hold Time Multiplier**

Default: 3

Range: 1 to 5

Function: Specifies the multiplier value used to compute the hold time set in the IIH packets transmitted on this interface. This multiplies the IIH Hello Timer value by this factor.

Instructions: Set to the appropriate value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.64

Parameter: **ISH Hold Time Multiplier**

Default: 3

Range: 1 to 5

Function: Specifies the multiplier value used to compute the hold time set in the ISH packets transmitted on this interface. This multiplies the ISH Hello Timer value by this factor.

Instructions: Set to the appropriate value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.65

Parameter: **Redirect Enable/Disable**

Default: Enabled

Options: Enabled | Disabled

Function: Specifies whether an OSI interface sends a redirect packet back to the originating system informing it of a more direct path to a destination system.

Redirects should be disabled when they are inappropriate for particular media and topology combinations. For example, when operating OSI

over a frame relay circuit configured for group access and the underlying topology is hub and spoke, redirects should be disabled.

Instructions: Set the Redirect Enable/Disable parameter to Disable to prevent redirect packets from being sent over the OSI interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.6.3.66

Configuring Static End System Adjacencies

You must define a static end system adjacency with any end system serviced by the router that 1) resides in the same area as the OSI router, 2) is reachable over a single interface, and 3) does not have ISO ESIS 9542 enabled.

To configure a static end system adjacency:

1. Select the Protocols→OSI→Interfaces option from the Wellfleet Configuration Manager window (refer to Figure 3-1). The OSI Interface Lists window appears (refer to Figure 3-3).
2. Click on Static ES Adjacencies. The OSI Static ES Adjacency List window appears (Figure 3-4). It lists all static end system adjacencies that are defined. If you did not add any end system adjacencies, none will be listed.

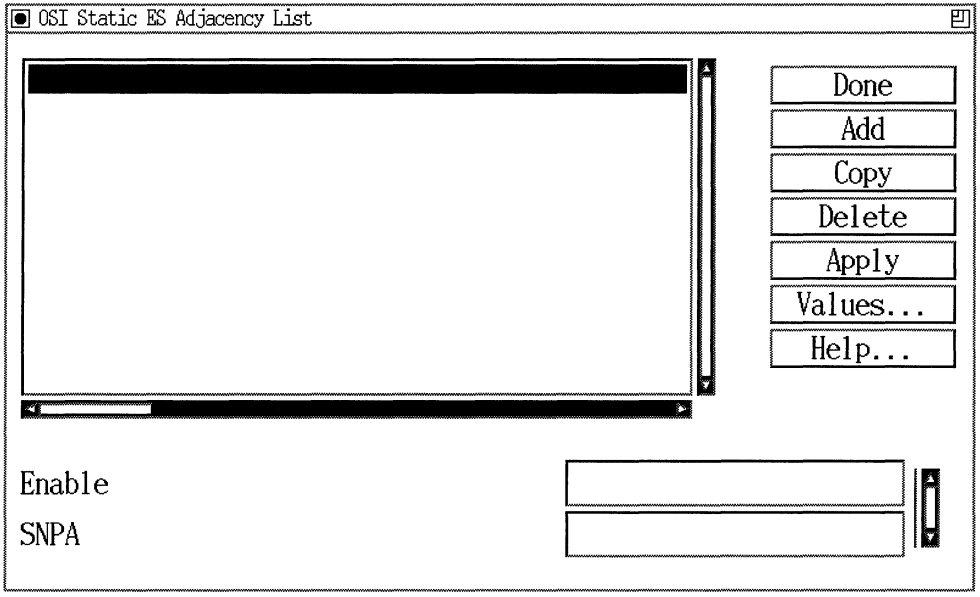

Figure 3-4. OSI Static ES Adjacency List Window

Continue to the following sections to add, copy, edit, or delete static end system adjacencies.

Adding a Static End System Adjacency

To add a static end system adjacency:

1. Click on Add in the OSI Static ES Adjacency List window (refer to Figure 3-4). The OSI Static ES Adjacency Configuration window appears (Figure 3-5).

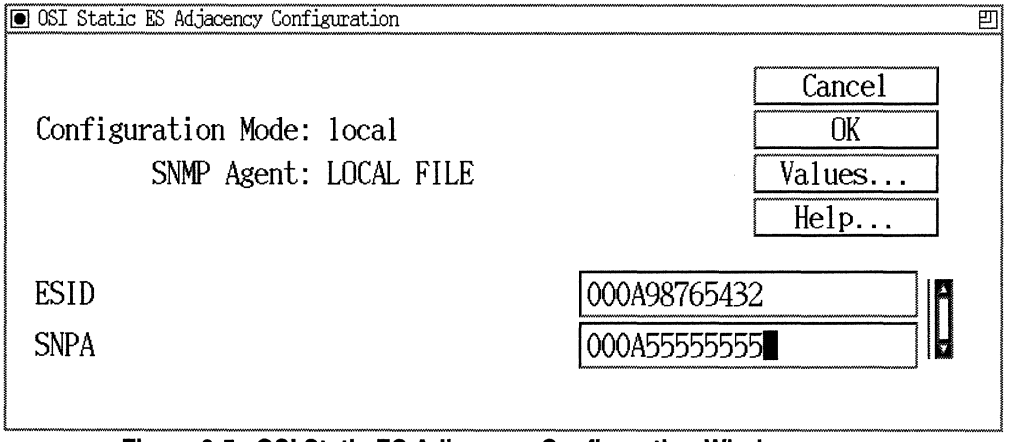

Figure 3-5. OSI Static ES Adjacency Configuration Window

2. Define the static end system parameters, using the descriptions in the next section as a guide.
3. Click on OK. The End System Adjacency List window displays the new adjacency you defined.
4. Repeat Steps 1–3 to add additional static end system adjacencies.

Continue to the following sections to copy, edit, or delete static end system adjacencies.

Static End System Adjacency Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the OSI Static ES Adjacency List and OSI Static ES Adjacency Configuration windows (refer to Figures 3-4 and 3-5):

Parameter:	**Enable**
Default:	Enable
Options:	Enable | Disable
Function:	Enables the end system adjacency as defined by the ESID and SNPA parameters.
Instructions:	Set to Enable to enable the end system adjacency.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.5.1.2

Parameter:	**ESID**
Default:	None
Options:	Any valid 6-byte End System ID
Function:	Specifies the End System ID of the adjacent end system.
Instructions:	Enter the 6-byte End System ID assigned to the adjacent end system in hexadecimal format.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.5.1.3

Parameter:	**SNPA**
Default:	None
Options:	Depends on the circuit type (see Instructions).
Function:	Specifies an SNPA for the adjacent end system.
Instructions:	Enter the SNPA for the adjacent end system: — If this circuit is an X.25 PDN circuit, then enter any valid X.121 address in decimal format. — If this circuit is an X.25 DDN circuit, then enter a valid X.121 address for the remote router in decimal format. — If this circuit uses PPP, then leave this field blank. — If this circuit is of any other type, then enter any valid MAC address.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.5.1.5

Note: To enter a valid X.121 address for an X.25 DDN circuit, you must convert the remote IP address to an X.121 address. (See Appendix A for the conversion algorithm.)

Copying a Static End System Adjacency

To copy a static end system adjacency:

1. Click on the adjacency you want to copy from the list in the OSI Static ES Adjacency List window (refer to Figure 3-4).
2. Define the ESID parameter for that adjacency.
3. Click on OK. The OSI Static ES Adjacency List window displays the new adjacency you copied.
4. Repeat Steps 1–3 to copy additional static end system adjacencies.
5. Click on Done to exit the window.

If you want to edit a static end system adjacency, continue to the next section. For details on deleting an adjacency, go to "Deleting a Static End System Adjacency."

Editing a Static End System Adjacency

To edit a static end system adjacency:

1. Select the adjacency you want to edit from the list in the OSI Static ES Adjacency List window (refer to Figure 3-4).
2. Edit the static adjacency parameters you want to change.
3. Click on Apply to implement your changes.
4. Repeat Steps 1–3 to edit additional static adjacencies.
5. Click on Done to exit the window.

If you want to delete a static end system adjacency, continue to the next section.

Deleting a Static End System Adjacency

To delete a static end system adjacency:

1. Select the adjacency you want to delete from the list in the OSI Static ES Adjacency List window (refer to Figure 3-4).
2. Click on Delete. The static end system adjacency is no longer listed.
3. Repeat Steps 1 and 2 to delete additional adjacencies.
4. Click on Done to exit the window.

Configuring Static External Address Adjacencies

You configure static external adjacencies to enable interdomain routing (routing between domains).

To configure a static external address adjacency:

1. Select the Protocols→OSI→Interfaces option from the Wellfleet Configuration Manager window (refer to Figure 3-1). The OSI Interface Lists window appears (refer to Figure 3-3).
2. Click on External Address Adjacency. The OSI External Address Adjacency List window appears (refer to Figure 3-6). It lists all external address adjacencies that are defined. If you did not add any adjacencies, none will be listed.

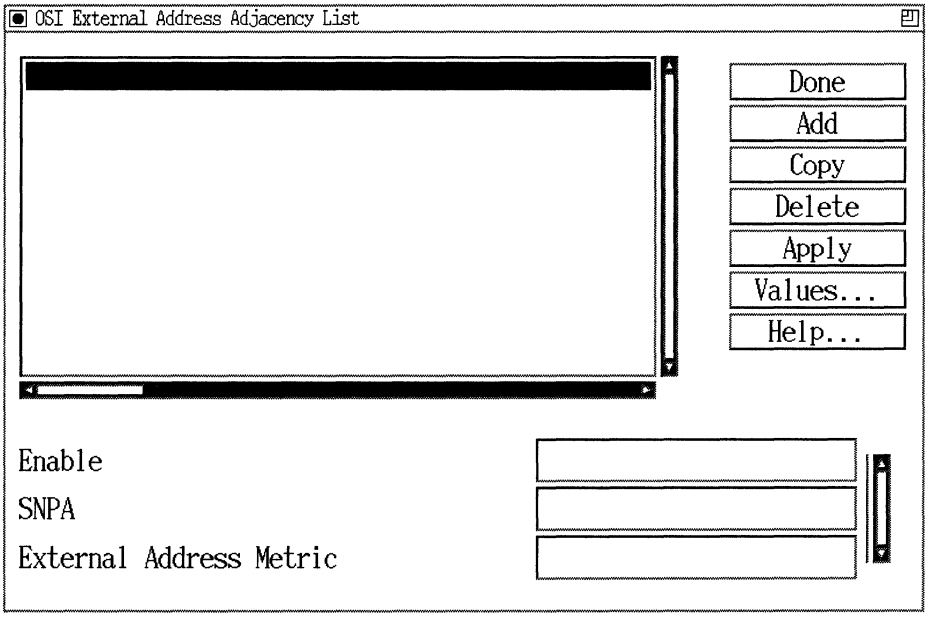

Figure 3-6. OSI External Address Adjacency List Window

Note: To configure static external address adjacencies for the OSI interface, set the Routing Level parameter in the OSI Interface Lists window to an option that includes External (for example, Level 2 and External).

Continue to the following sections to add, remove, copy, or edit external address adjacencies from this window as described in the following sections.

Adding Static External Address Adjacencies

To add a static external address adjacency:

1. Click on Add from the OSI External Address Adjacency List window (refer to Figure 3-6). The OSI External Address Adjacency Configuration window appears (Figure 3-7).

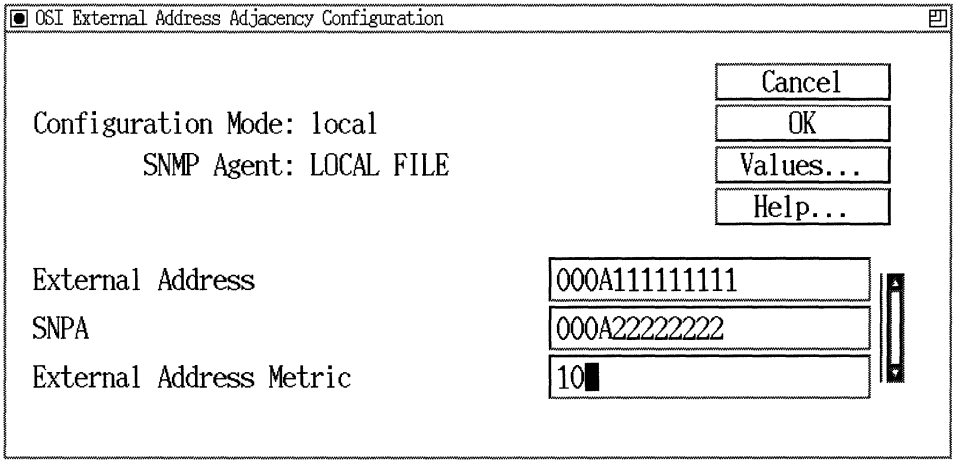

Figure 3-7. OSI External Address Adjacency Configuration Window

2. Define the static external address adjacency parameters, using the descriptions in the next section as a guide.
3. Click on OK to implement your changes and exit the window. The OSI External Address Adjacency List window displays the new adjacency you defined.
4. Repeat Steps 1–3 to add additional adjacencies.

Static External Address Adjacency Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the OSI External Address Adjacency Configuration window (refer to Figure 3-7):

Parameter:	**Enable**
Default:	Enable
Options:	Enable | Disable
Function:	Enables the external adjacency defined by the SNPA parameter.
Instructions:	Set to Enable to enable this external adjacency.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.4.2

Parameter:	**External Address**
Default:	None
Options:	Any valid address
Function:	Specifies the destination address of the external adjacency.
Instructions:	Enter the address assigned to the external adjacency in hexadecimal format.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.4.5

Parameter:	**SNPA**
Default:	None
Options:	Depends on the circuit type (see "Instructions")
Function:	Specifies a SNPA for the adjacent end system.
Instructions:	Enter the SNPA for the adjacent end system as follows: — If this circuit is an X.25 PDN circuit, then enter a valid X.121 address for the remote router in decimal format. — If this circuit is an X.25 DDN circuit, then enter a valid X.121 address for the remote router in decimal format. — If this circuit uses PPP, then leave this field blank. — If this circuit is of any other type, then enter any valid MAC address.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.4.6

Note: To enter a valid X.121 address for an X.25 DDN circuit, you must convert the remote IP address to an X.121 address. (See Appendix A for the conversion algorithm.)

Parameter:	**External Address Metric**
Default:	20
Range:	1 to 63
Function:	Specifies the relative cost of using this interface to reach the external adjacency. If there are multiple interfaces configured to the same external adjacency, the OSI router will route all external domain traffic using the interface that has been assigned the lowest External Address Metric.
Instructions:	If you only have a single link to the external adjacency, or have no preference regarding which interface is used to access the external domain, accept the default value. If there are multiple interfaces configured to the same external adjacency, and you want this interface to be used regularly, then assign it the lowest External Address Metric. Similarly, assign it a high cost if you do not want it to be used regularly.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.4.7

Copying Static External Address Adjacencies

To copy a static external address adjacency:

1. Select the adjacency you want to copy from the list in the OSI External Address Adjacency List window (refer to Figure 3-6).
2. Click on Copy.
3. Define the External Address for the new adjacency.
4. Click on Save.
5. Repeat Steps 1–4 to copy additional adjacencies.

Editing Static External Address Adjacencies

To edit a static external address adjacency:

1. Select the adjacency you want to edit from the list in the OSI External Address Adjacency List window (refer to Figure 3-6).
2. Edit the static external address adjacency parameters.
3. Click on Update to implement your changes.
4. Repeat Steps 1–3 to edit additional adjacencies.

Deleting Static External Address Adjacencies

To delete a static external address adjacency:

1. Select the adjacency you want to delete from the list in the OSI External Address Adjacency List window (refer to Figure 3-6).
2. Click on Delete. The static external address adjacency is no longer listed.
3. Repeat Steps 1 and 2 to delete additional adjacencies.
4. Click on Done to exit the window.

Configuring Static Routes

You configure static routes when you want to control which path the router uses to route OSI traffic.

To configure a static route, select the Protocols→OSI→Static Routes option in the Wellfleet Configuration Manager window (refer to Figure 3-1). The OSI Static Routes window appears (Figure 3-8). It lists all static routes that are defined. If you did not add any static routes, none will be listed.

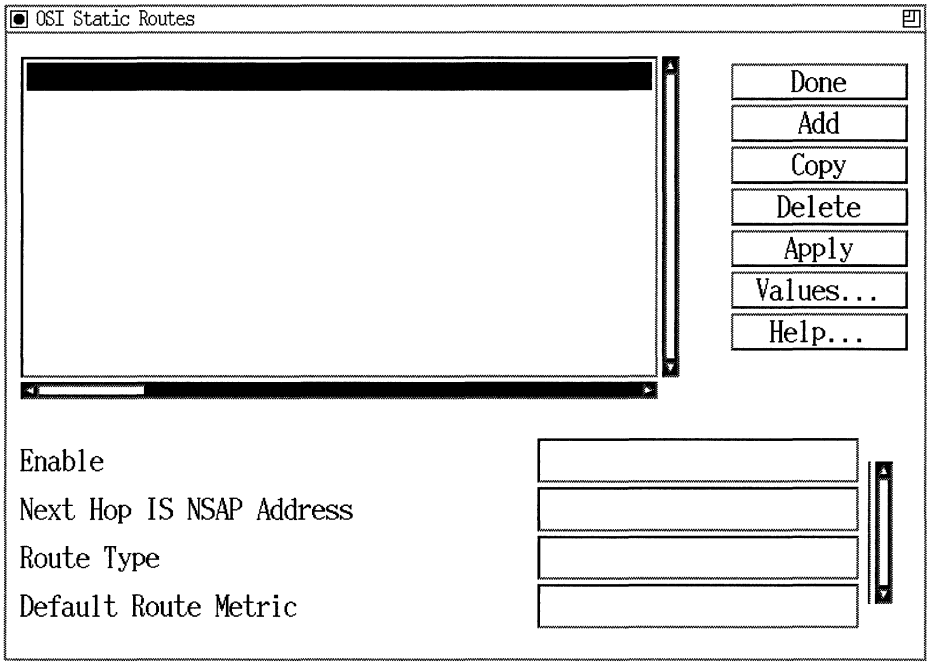

Figure 3-8. OSI Static Routes Window

Continue to the following sections to add, copy, edit, or delete static routes.

Adding Static Routes

To add a static route:

1. Click on Add in the OSI Static Routes window (refer to Figure 3-8). The Static Route Configuration window appears (Figure 3-9).

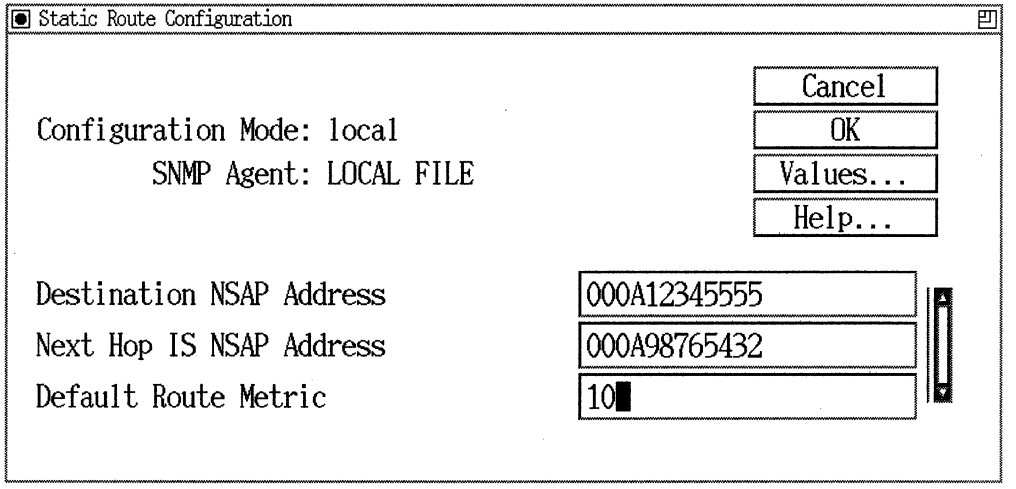

Figure 3-9. Static Route Configuration Window

2. Define the static route parameters, using the descriptions in the next section as a guide.
3. Click on OK to implement your changes. The OSI Static Routes window displays the new static route you defined.
4. Repeat Steps 1–3 to add additional static routes.

Continue to the following sections to copy, edit, or delete static routes.

OSI Static Route Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the OSI Static Routes and OSI Static Route Configuration windows (refer to Figures 3-8 and 3-9):

Parameter:	**Enable**
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables the selected static route.
Instructions:	To disable the static route, set to Disable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.2.1.2

Parameter:	**Destination NSAP Address**
Default:	None
Options:	Any valid NSAP Address
Function:	Specifies the NSAP address of the destination end system.
Instructions:	Enter the address assigned to the destination end system in hexadecimal format.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.2.1.4

Parameter:	**Route Type**
Default:	None
Options:	End System | Area | External Domain
Function:	Specifies the route type.
Instructions:	Select the route type for this static route.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.2.1.6

Parameter:	**Next Hop IS NSAP Address**
Default:	None
Options:	Any valid NSAP address
Function:	Specifies the NSAP address of the intermediate system that is the next hop on the path to the destination end system.
Instructions:	Enter the address assigned to the next-hop intermediate system in hexadecimal format.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.2.1.5

Note: The next hop that you specify for the Next Hop IS NSAP Address parameter must be an intermediate system with which this router has a dynamic or static adjacency.

Parameter:	**Default Route Metric**
Default:	20
Range:	1 to 1023
Function:	Specifies the default metric (relative cost) of routing Level 1 traffic over this interface. The OSI router always selects the circuit(s) with the lowest cost when defining a path, so assigning each circuit a cost is, in effect, a way of assigning it a priority.
Instructions:	If you do not want to use this interface to route Level 1 traffic on a regular basis, assign it a high cost. Otherwise, accept the default (20).
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.2.1.7

Copying Static Routes

To copy a static route:

1. Select the static route you want to copy from the list in the OSI Static Routes window (refer to Figure 3-8).
2. Click on Copy.
3. Define the static route parameters.
4. Click on OK to implement your changes. The OSI Static Routes window displays the new static route you defined.
5. Repeat Steps 1–4 to copy additional static routes.
6. Click on Done to exit the screen.

Editing Static Routes

To edit a static route:

1. Select the static route you want to edit from the list in the OSI Static Routes window (refer to Figure 3-8).
2. Edit the static route parameters.
3. Click on Apply to implement your changes.
4. Repeat Steps 1–3 to edit additional static routes.
5. Click on Done to exit the screen.

Deleting Static Routes

To delete a static route:

1. Select the static route you want to delete from the list in the OSI Static Routes window (refer to Figure 3-8).
2. Click on Delete. The static route is no longer listed.
3. Repeat Steps 1 and 2 to delete additional static routes.
4. Click on Done to exit the screen.

Configuring DECnet IV to V Transition

You create, edit, and delete DECnet IV to V Transition from the Wellfleet Configuration Manager window. Continue to the next section if you want to create DECnet IV to V Transition.

Note: See *Customizing DECnet Services* for more information about the DECnet IV to V Transition feature.

Creating the DECnet IV to V Transition

From the Wellfleet Configuration Manager window, select Protocols→OSI→Create DECnet IV to V Transition (Figure 3-10). This enables the DECnet IV to V Transition feature. If you select Protocols→OSI, you see that the edit and delete options are now available.

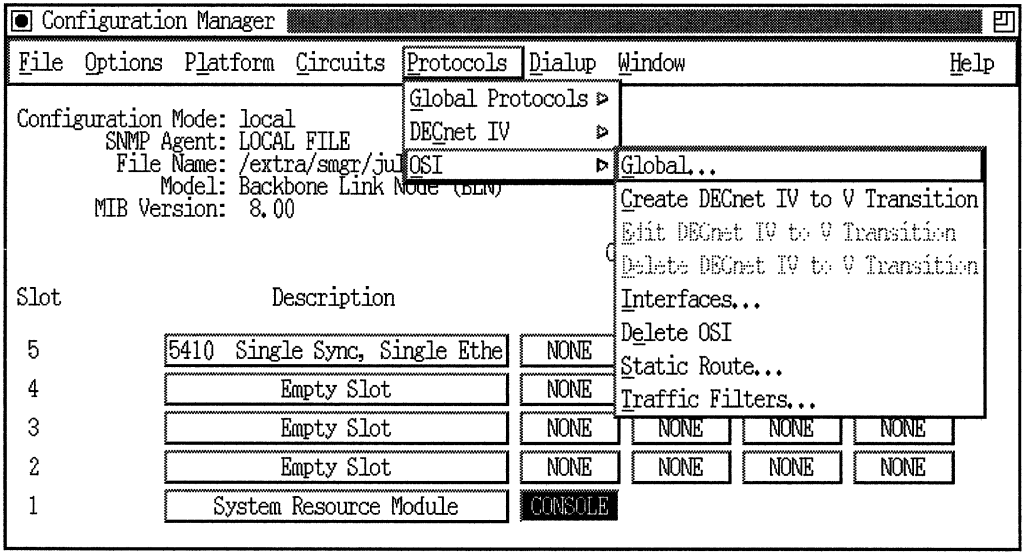

Figure 3-10. Selecting Protocols→OSI→Create DECnet IV to V Transition

If you want to edit the DECnet IV to V Transition parameters, continue to the next section. For details on deleting DECnet IV to V Transition, go to "Deleting DECnet IV to V Transition."

Editing the DECnet IV to V Transition Parameters

To edit the DECnet IV to V Transition parameters:

1. Select Protocols→OSI→Edit DECnet IV to V Transition from the Wellfleet Configuration Manager window (refer to Figure 3-10). The Edit DECnet IV to V Transition Parameters window appears (Figure 3-11).

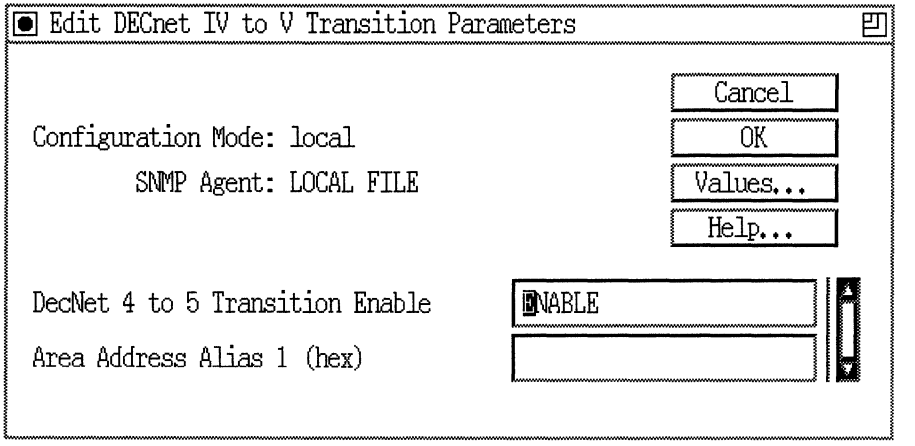

Figure 3-11. Edit DECnet IV to V Transition Parameters Window

2. Edit the parameters, using the descriptions in the next section as a guide.
3. Click on OK to implement your changes and exit the screen.

If you want to delete DECnet IV to V Transition from the router, go to "Deleting DECnet IV to V Transition."

DECnet IV to V Transition Parameter Descriptions

Use the following descriptions as a guide when you configure the parameters in the Edit DECnet IV to V Transition Parameters window (refer to Figure 3-11):

Parameter:	**DECnet 4 to 5 Transition Enable**
Default:	None
Options:	Enable | Disable
Function:	Enables or disables DECnet IV to V Transition.
Instructions:	To enable the transition, set this value to Enable. Otherwise, set this value to Disable to turn it off.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.12.2

Parameter:	**Area Address Alias 1 (hex)**
Default:	None
Options:	Any valid area address.
Function:	Assigns the first area address alias to the router. An area address alias is a different area address that is assigned to the same router. For the DECnet IV to V Transition feature, the area address alias defines the Phase IV prefix and Phase IV area fields of the Phase IV-compatible address.
Instructions:	Enter the area address alias in hexadecimal format. For the DECnet IV to V Transition feature, enter the Phase IV prefix (from 1 to 9 bytes) followed by 2 bytes of the Phase IV area address. Otherwise, leave this field blank.
MIB Object ID:	1.3.6.1.4.1.18.3.5.6.1.17

Deleting DECnet IV to V Transition

To delete DECnet IV to V Transition:

1. Select Protocols→OSI→Delete DECnet IV to V Transition from the Wellfleet Configuration Manager window (refer to Figure 3-10). A window pops up and prompts:

```
Do you REALLY want to delete OSI DECnet IV to V Transition?
```

2. Click on OK. The system returns you to the Wellfleet Configuration Manager window. The DECnet IV to V Transition feature is no longer configured on the router.

Deleting OSI from the Router

To delete the OSI routing protocol from all router circuits on which it is currently enabled:

1. Select the Protocols→OSI→Delete OSI option from the Wellfleet Configuration Manager window (refer to Figure 3-1). A window pops up and prompts:

```
Do you REALLY want to delete OSI?
```

2. Click on OK. The Wellfleet Configuration Manager window appears. OSI is no longer configured on the router.

If you examine the Wellfleet Configuration Manager window, you see that the connectors for circuits on which OSI was the *only* protocol enabled are no longer highlighted. You must reconfigure the circuits for these connectors. See *Configuring Wellfleet Routers* for details on configuring circuits.

Appendix A
IP-to-X.121 Address Mapping for DDN

This appendix describes how to convert an IP address to an X.121 address if you are configuring OSI over DDN X.25. You enter this converted address when you add static end system adjacencies or an external address. (See Chapters 2 and 3 for additional information.)

This appendix includes

- An overview of the IP address classes
- Address conversion methods
- Example address conversions

Note: The information in this appendix was taken from RFC 1236, *IP to X.121 Address Mapping for DDN*.

IP-to-X.121 Address Mapping

This section defines a standard way of converting IP addresses to CCITT X.121 addresses and is the recommended standard for use on the Internet, specifically for the Defense Data Network (DDN). This section provides information for the Internet community. It does not specify an Internet standard.

Overview

The Defense Communication Agency (DCA) has stated that "DDN specifies a standard for mapping Class A addresses to X.121 addresses." Additionally, DCA has stated that Class B and C IP-to-X.121 address mapping standards "are the responsibility of the administration of the Class B or C network in question." Therefore, there is no defined standard way of converting Class B and Class C IP addresses to X.121 addresses.

This is an important issue because currently there is no way for administrators to define IP-to-X.121 address mapping. Without a single standard, in a multi-vendor network environment there is no assurance that devices using IP and DDN X.25 will communicate with each other.

The IP-to-X.121 address mapping of Class B and Class C IP addresses shall be implemented as described below. This translation method is a direct expansion of the algorithm described in the MIL-STD: X.25, DDN X.25 Host Interface Specification[1]. The translation method described below is totally independent of IP subnetting and of any masking that may be used in support of IP subnetting.

1. MIL-STD: X.25 "Defense Data Network X.25 Host Interface Specification," Defense Communications Agency, BBN Communications Corporation, 1983 December, Volume 1 of the *DDN Protocol Handbook* (NIC 50004). Also available on-line at the DDN NIC as NETINFO:X.25.DOC.

Background

All Internet hosts are assigned a four-octet (32-bit) address composed of a network field and a local address field (also known as the REST field[2]); refer to Figures A-1 through A-3. Two basic forms of addresses are provided: (1) physical addresses, which correspond to the node number and DCE port number of the node to which the DTE is connected and (2) logical addresses, which are mapped transparently by DCE software into a corresponding physical network address.

To provide flexibility, Internet addresses are divided into three primary classes: Class A, Class B, and Class C. These classes allow for a large number of small and medium-sized networks. The network addresses used within the Internet in Class A, B, and C networks are divided between Research, Defense, Government (Non-Defense), and Commercial uses.

As described in the MIL-STD: X25, an IP address consists of the ASCII text string representation of four decimal numbers separated by periods, corresponding to the four octets of a thirty-two-bit Internet address. The four decimal numbers are referred to in this appendix as network (**n**), host (**h**), logical address (**l**), and Interface Message Processor (IMP) or Packet Switch Node (PSN) (**i**). Thus, an Internet address may be represented as **n.h.l.i** (Class A), **n.n.h.i** (Class B), or **n.n.n.hi** (Class C), depending on the Internet address class. Each of these four numbers will have one, two, or three decimal digits and will never have a value greater than 255. For example, in the Class A IP address 26.9.0.122, **n** = 26, **h**= 9, **l**= 0, and **i**= 122.

2. MIL-STD: 1777 "Internet Protocol," 1983 August, Volume 1 of the *DDN Protocol Handbook* (NIC 50004).

The different classes of Internet addresses[3] are illustrated:

Class A:

- The highest-order bit is set to 0.
- 7 bits define the network number.
- 24 bits define the local address.
- This allows up to 126 Class A networks.
- Networks 0 and 127 are reserved.

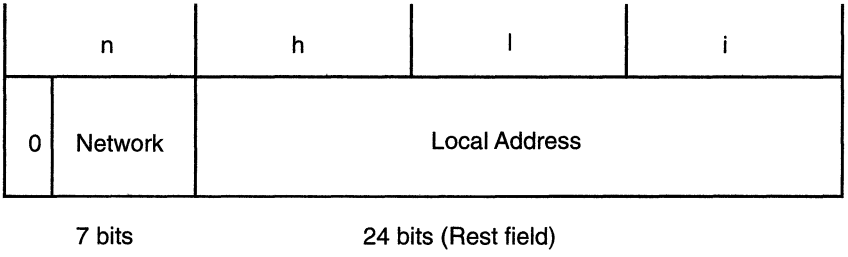

Figure A-1. Class A Internet Address

3. Kirkpatrick, S., M. Stahl, and M. Recker, *Internet Numbers*, RFC 1166, DDN NIC, July 1990.

Class B:

- The two highest-order bits are set to 1-0.
- 14 bits define the network number.
- 16 bits define the local address.
- This allows up to 16,384 Class B networks.

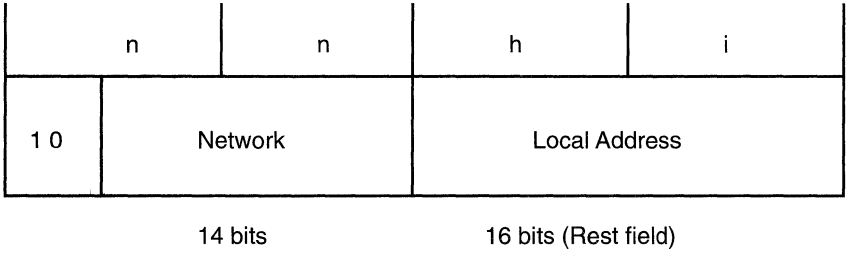

Figure A-2. Class B Internet Address

Class C:

- The three highest-order bits are set to 1-1-0.
- 21 bits define the network number.
- 8 bits define the local address.
- This allows up to 2,097,152 Class C networks.

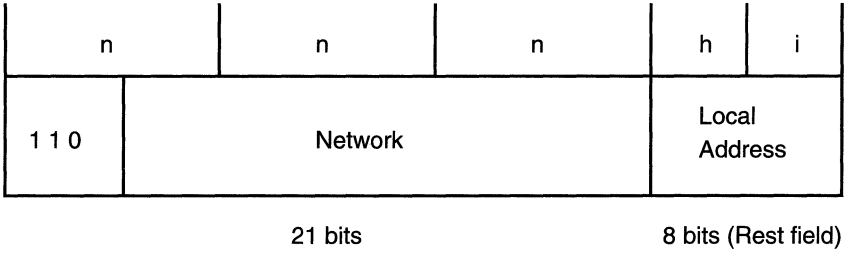

Figure A-3. Class C Internet Address

The fourth type of address, Class D, is used as a multicast address. The four highest-order bits are set to 1-1-1-0.

Note: No addresses are allowed with the four highest-order bits set to 1-1-1-1. These addresses, called *Class E*, are reserved.

The MIL-STD: X.25 states that "All DDN addresses are either twelve or fourteen BCD (binary-coded decimal) digits in length." The last two digits are referred to as the Sub-Address and are not used on the DDN. The Sub-Address is carried across the network without modification. Its presence is optional. Therefore, a DTE may generate either twelve or fourteen BCD X.121 address, but must accept both twelve and fourteen BCD X.121 addresses.

Standard IP to X.121 Address Mapping

This section describes the algorithm that should be used to convert IP addresses to X.121 addresses. You will note that **h** is always listed as greater than or less than the number 64. This number is used to differentiate between PSN physical and logical host port addresses. Note that at the time of this writing, the DDN does not make use of the PSN's logical addressing feature, which allows hosts to be addressed independently of their physical point of attachment to the network.

Derivation of DDN X.25 Addresses

The following describes Class A, B, and C IP address to DDN X.25 address conversion.

Class A IP Address to DDN X.25 Address Conversion

To convert a Class A IP address to a DDN X.25 address:

If the host field (**h**) is less than 64 (**h** < 64), the address corresponds to the following DDN X.25 physical address:

ZZZZ F III HH ZZ (SS)

Where:

- **ZZZZ** = 0000
- **F** = 0 because the address is a physical address
- **III** is a three decimal digit representation of **i**, right-adjusted and padded with leading zeros if required
- **HH** is a two decimal digit representation of **h**, right-adjusted and padded with leading zeros if required
- **ZZ** = 00 is optional
- **(SS)** is an optional Sub-Address field that is ignored in the DDN; this field is either left out or filled with zeros

The address 26.9.0.122 corresponds to the DDN X.25 physical address 000001220900.

If the host field (**h**) is greater than or equal to 64 (**h** ≥ 64), the address corresponds to the following DDN X.25 physical address:

ZZZZ F RRRRR ZZ (SS)

Where:

- **ZZZZ** = 0000
- **F** = 1 because the address is a logical address
- **RRRRR** is a five-decimal-digit representation of the result **r** of the calculation
- **r** = **h*** 256 + **i** (note that the decimal representation of **r** will always require five digits)
- **ZZ** = 00
- **(SS)** is optional

The address 26.83.0.207 corresponds to the DDN X.25 logical address 000012145500.

Class B IP Address-to-DDN X.25 Address Conversion

For Class B IP addresses, the **h** and **i** fields will always consist of 8 bits, each taken from the REST field of the Internet address. The mapping follows the same rules as Class A.

Class C IP Address-to-DDN X.25 Address Conversion

For Class C IP addresses, the **h** and **i** fields will always consist of 4 bits, each taken from the REST field of the Internet address. The mapping follows the same rules as Class A.

Examples

The following are examples of IP-to-X.121 address mappings for Class A, Class B, and Class C IP addresses.

Class A Example

This is an example of the mapping of an X.121 address for Class A networks.

For **h** < 64:

Example: 26.29.0.122, format: **n.h.l.i**

ZZZZ F III HH ZZ (SS)

0000 0 122 29 00 00 = X.21 address

For **h** > or = 64:

Example: 26.80.0.122, format: **n.h.l.i**

ZZZZ F RRRRR ZZ (SS)

0000 1 20602 00 00 = X.21 address

Where **r**= **h*** 256 + **i**

Class B Example

This is an example of the mapping of an X.121 address for Class B networks.

For **h** < 64:

Example: 137.80.1.5, format: **n.n.h.i**

ZZZZ F III HH ZZ (SS)

0000 0 005 01 00 00 = X.121 address

For **h** > or = 64:

Example: 137.80.75.2, format: **n.n.h.i**

ZZZZ 1 RRRRR ZZ (SS)

0000 1 19202 00 00 = X.121 address

Where **r**= **h*** 256 + **i**

Class C Example

This is an example of the mapping of an X.121 address for Class C networks.

For **h** < 64:

Example: 192.33.50.19, format: **n.n.n.hi**

h i

n.n.n.0001 0011

1 3

Subnet 1
Subhost 3

ZZZZ F III HH ZZ (SS)

0000 0 003 01 00 00 = X.121 address

Note: The mapping of X.121 address for Class C networks for **h** > 64 is not applicable since the **h** field can never exceed 15.

Index

M

N

O

P

R

S

T

X

Bay Networks

The Merged Company of SynOptics and Wellfleet

8 Federal Street
Billerica, MA 01821

Printed in U.S.A. on Recycled Paper