



Management Network Connectivity

This chapter provides an overview of Cisco ONS 15600 data communications network (DCN) connectivity. Cisco optical network communication is based on IP, including communication between Cisco Transport Controller (CTC) computers and ONS 15600s, and communication among networked ONS 15600 nodes. The chapter provides scenarios showing ONS 15600s in common IP network configurations as well as information about provisionable patchcords, the IP routing table, external firewalls, and open gateway network element (GNE) networks.



Note

This chapter does not provide a comprehensive explanation of IP networking concepts and procedures, nor does it provide IP addressing examples to meet all networked scenarios. For ONS 15600 networking setup instructions, refer to the “Turn Up a Node” chapter of the *Cisco ONS 15600 Procedure Guide*.

Although ONS 15600 DCN communication is based on IP, ONS 15600s can be networked to equipment that is based on the Open System Interconnection (OSI) protocol suites. This chapter describes the ONS 15600 OSI implementation and provides scenarios that show how ONS 15600 can be networked within a mixed IP and OSI environment.

Chapter topics include:

- [9.1 IP Networking Overview, page 9-1](#)
- [9.2 ONS 15600 IP Addressing Scenarios, page 9-2](#)
- [9.3 Provisionable Patchcords, page 9-19](#)
- [9.4 Routing Table, page 9-20](#)
- [9.5 External Firewalls, page 9-22](#)
- [9.6 Open GNE, page 9-24](#)
- [9.7 TCP/IP and OSI Networking, page 9-26](#)



Note

To set up ONS 15600s within an IP network, you must work with a LAN administrator or other individual at your site who has IP networking training and experience.

9.1 IP Networking Overview

ONS 15600s can be connected in many different ways within an IP environment:

- You can connect ONS 15600 nodes and LANs through direct connections or a router.

- IP subnetting can create ONS 15600 node groups, which allow you to provision nodes in a network that are not connected using the data communications channel (DCC).
- Different IP functions and protocols allow you to achieve specific network goals. For example, Proxy Address Resolution Protocol (ARP) enables one LAN-connected ONS 15600 to serve as a gateway for ONS 15600s that are not connected to the LAN.
- You can create static routes to enable connections among multiple CTC sessions with ONS 15600s that reside on the same subnet but have different destination IP addresses.
- If ONS 15600s are connected to Open Shortest Path First (OSPF) networks, ONS 15600 network information is automatically communicated across multiple LANs and WANs.

9.2 ONS 15600 IP Addressing Scenarios


ONS 15600 IP addressing generally has eight common scenarios or configurations. Use the scenarios as building blocks for more complex network configurations. [Table 9-1](#) provides a general list of items to check when setting up ONS 15600s in IP networks.

Table 9-1 General ONS 15600 IP Troubleshooting Checklist

Item	What to Check
Link integrity	Verify that link integrity exists between: <ul style="list-style-type: none"> • CTC computer and network hub/switch • ONS 15600s (backplane ports or active Timing and Shelf Controller [TSC] card port) and network hub/switch • Router ports and hub/switch ports
ONS 15600 hub/switch ports	If connectivity problems occur, set the hub or switch port that is connected to the ONS 15600 to 10 Mbps half-duplex.
Ping	Ping the node to test connections between computers and ONS 15600s.
IP addresses/subnet masks	Verify that ONS 15600 IP addresses and subnet masks are set up correctly.
Optical connectivity	Verify that ONS 15600 optical trunk ports are in service; DCC is enabled on each trunk port.

9.2.1 Scenario 1: CTC and ONS 15600s on the Same Subnet

Scenario 1 shows a basic ONS 15600 LAN configuration ([Figure 9-1](#)). The ONS 15600s and CTC computer reside on the same subnet. All ONS 15600s connect to LAN A, and all ONS 15600s have DCC connections.

Figure 9-1 Scenario 1: CTC and ONS 15600s on Same Subnet

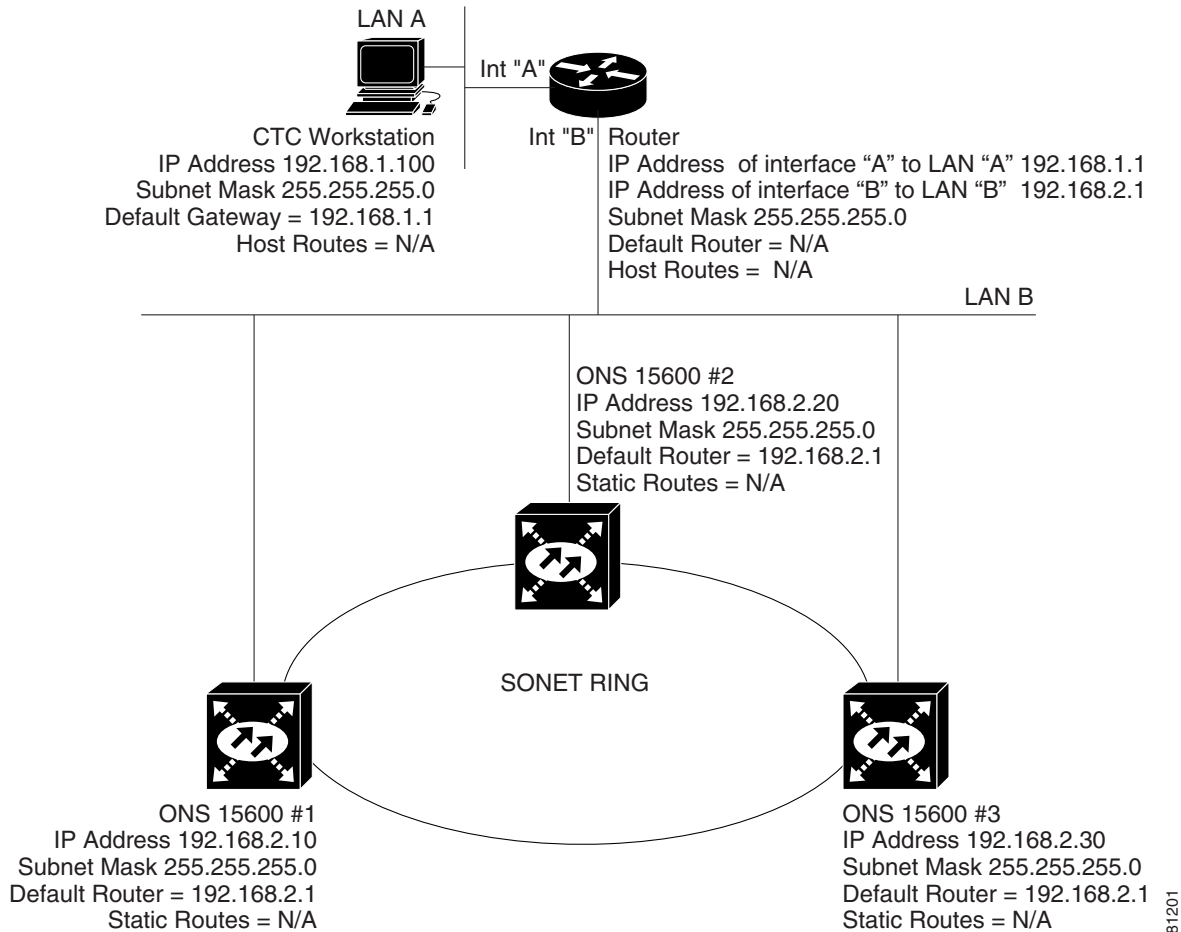
81200

9.2.2 Scenario 2: CTC and ONS 15600s Connected to Router

In Scenario 2, the CTC computer resides on a subnet (192.168.1.0) and attaches to LAN A (Figure 9-2). The ONS 15600s reside on a different subnet (192.168.2.0) and attach to LAN B. A router connects LAN A to LAN B. The IP address of router interface A is set to LAN A (192.168.1.1), and the IP address of router interface B is set to LAN B (192.168.2.1).

On the CTC computer, the default gateway is set to router interface A. If the LAN uses Dynamic Host Configuration Protocol (DHCP), the default gateway and IP address are assigned automatically. In the Figure 9-2 example, a DHCP server is not available.

Figure 9-2 Scenario 2: CTC and ONS 15600s Connected to Router



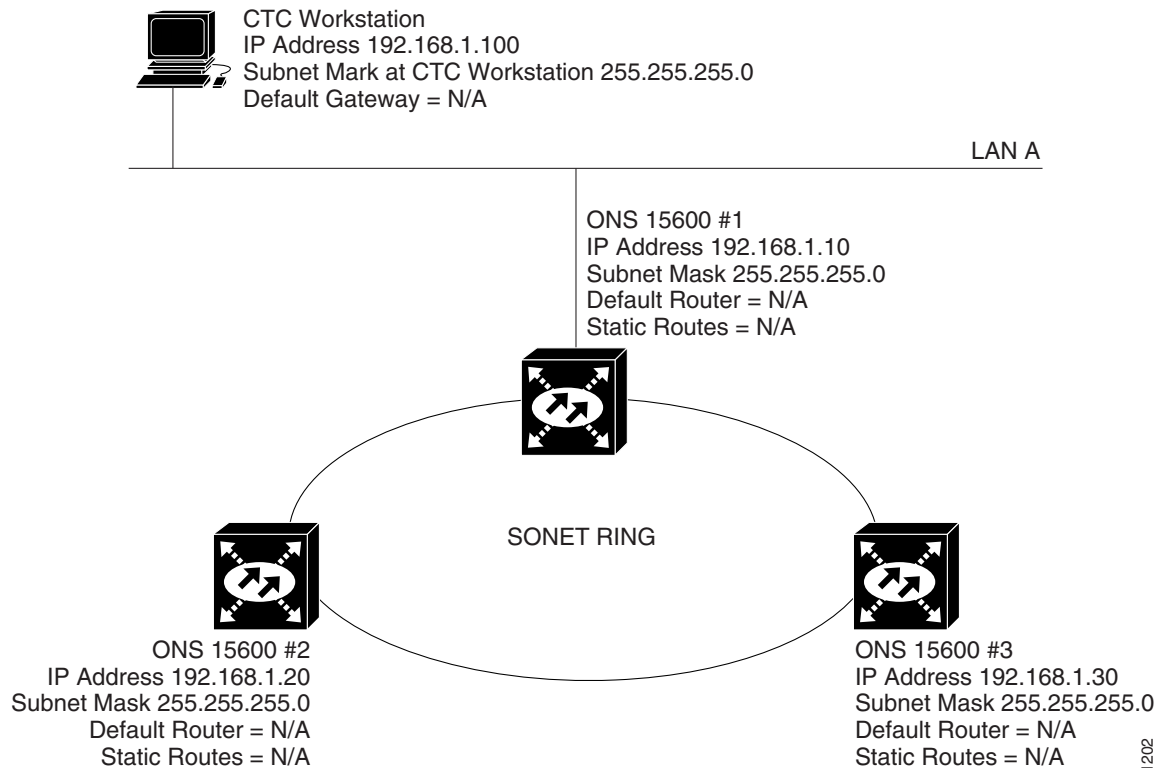
9.2.3 Scenario 3: Using Proxy ARP to Enable an ONS 15600 Gateway

Scenario 3 is similar to Scenario 1, but only one ONS 15600 (Node 1) connects to the LAN (Figure 9-3). Two ONS 15600s (Nodes 2 and 3) connect to Node 1 through the SONET data communications channel (DCC). Because all three ONS 15600s are on the same subnet, Proxy ARP enables Node 1 to serve as a gateway for Nodes 2 and 3.



Note

This scenario assumes that all CTC connections are to Node 1. If you connect a laptop to either Node 2 or Node 3, network partitioning will occur; neither the laptop or the CTC computer will be able to see all nodes. If you want laptops to connect directly to end network elements (ENEs), you will need to create static routes (refer to the “9.2.5 Scenario 5: Using Static Routes to Connect to LANs” section on page 9-6) or enable the ONS 15600 proxy server (see the “9.2.7 Scenario 7: Provisioning the ONS 15600 Proxy Server” section on page 9-11).

Figure 9-3 Scenario 3: Using Proxy ARP

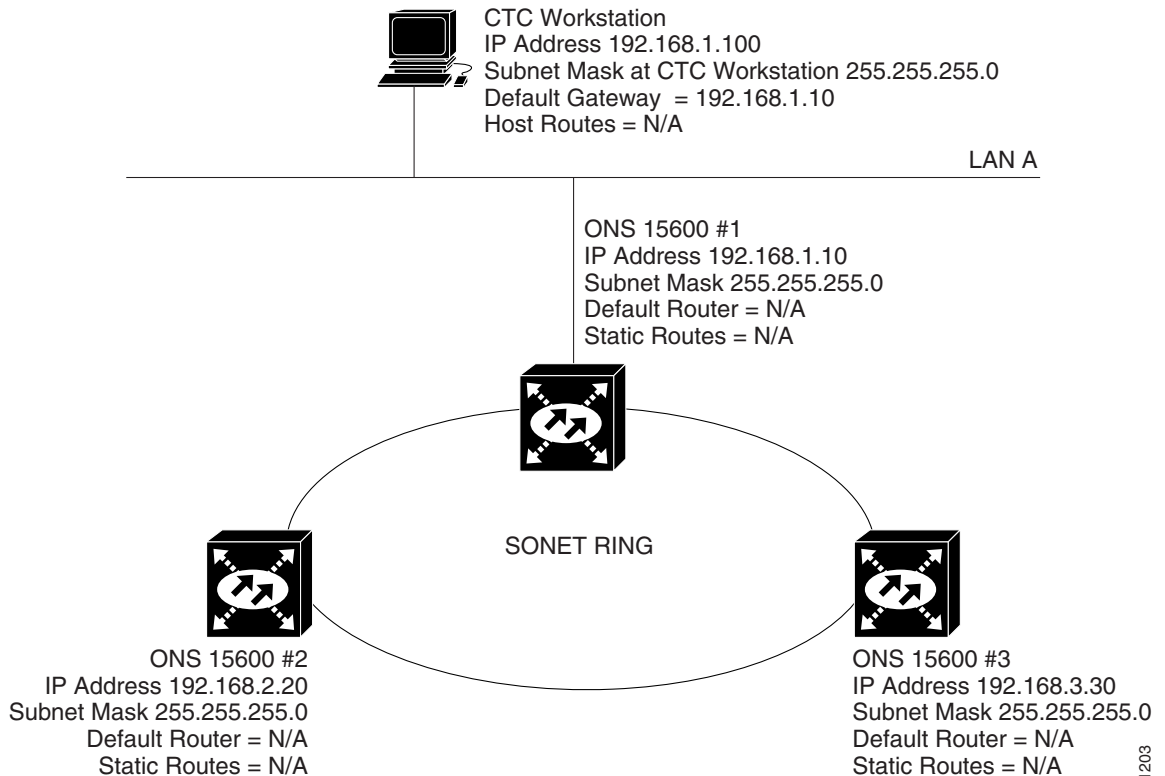
ARP matches higher-level IP addresses to the physical addresses of the destination host. It uses a lookup table (called the ARP cache) to perform the translation. When the address is not found in the ARP cache, a broadcast is sent out on the network with a special format called the ARP request. If one of the machines on the network recognizes its own IP address in the request, it sends an ARP reply back to the requesting host. The reply contains the physical hardware address of the receiving host. The requesting host stores this address in its ARP cache so that all subsequent datagrams (packets) to this destination IP address can be translated to a physical address.

Proxy ARP enables one LAN-connected ONS 15600 to respond to the ARP request for ONS 15600s that are not connected to the LAN. (ONS 15600 Proxy ARP requires no user configuration.) For this response to occur, the DCC-connected ONS 15600s must reside on the same subnet. When a LAN device sends an ARP request to an ONS 15600 that is not connected to the LAN, the gateway ONS 15600 returns its MAC address to the LAN device. The LAN device then sends the datagram for the remote ONS 15600 to the MAC address of the proxy ONS 15600. The proxy ONS 15600 uses its routing table to forward the datagram to the non-LAN ONS 15600.

9.2.4 Scenario 4: Default Gateway on CTC Computer

Scenario 4 is similar to Scenario 3, but Nodes 2 and 3 reside on different subnets, 192.168.2.0 and 192.168.3.0, respectively (Figure 9-4). Node 1 and the CTC computer are on subnet 192.168.1.0. For the CTC computer to communicate with Nodes 2 and 3, you would enter Node 1 as the default gateway on the CTC computer.

Figure 9-4 Scenario 4: Default Gateway on a CTC Computer



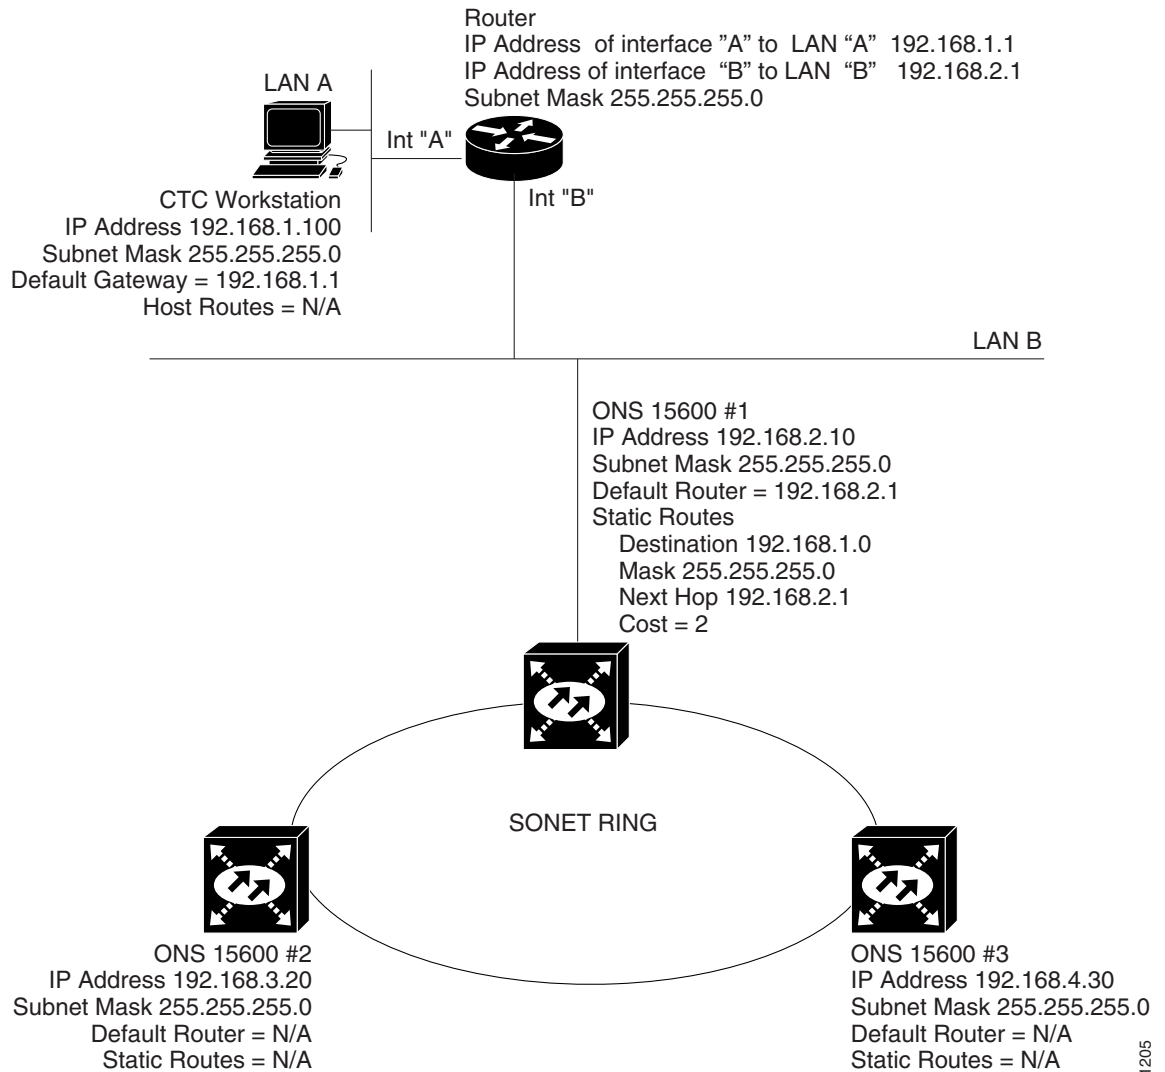
81203

9.2.5 Scenario 5: Using Static Routes to Connect to LANs

Use static routes for the following two reasons:

- To connect ONS 15600s to CTC sessions on one subnet connected by a router to ONS 15600s residing on another subnet. (These static routes are not needed if OSPF is enabled. Scenario 6 shows an OSPF example.)
- To enable multiple CTC sessions among ONS 15600s residing on the same subnet.

In [Figure 9-5](#), one CTC residing on subnet 192.168.1.0 connects to a router through interface A. (The router is not set up with OSPF.) ONS 15600s residing on subnet 192.168.2.0 are connected through Node 1 to the router through interface B. To connect to CTC computers on LAN A, you would create a static route on Node 1.

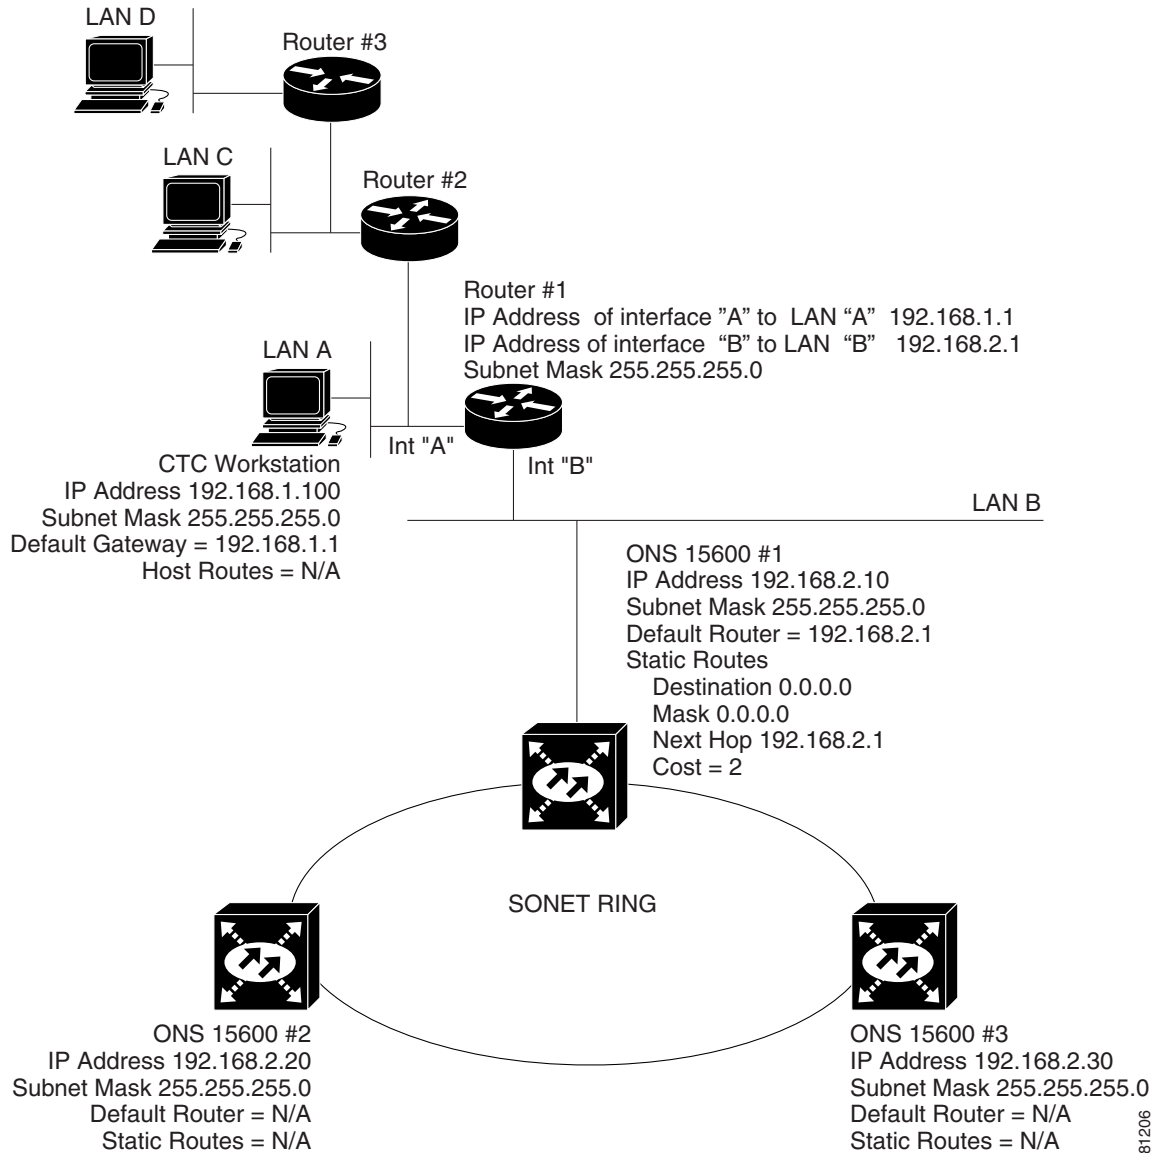
Figure 9-5 Scenario 5: Static Route with One CTC Computer Used as a Destination

The destination and subnet mask entries control access to the ONS 15600s:

- If a single CTC computer will be connected to a router, enter the complete CTC "host route" IP address as the destination with a subnet mask of 255.255.255.255.
- If CTC computers on a subnet are connected to a router, enter the destination subnet (in this example, 192.168.1.0) and a subnet mask of 255.255.255.0.
- If all CTC computers are connected to router, enter a destination of 0.0.0.0 and a subnet mask of 0.0.0.0. [Figure 9-6](#) shows an example.

In this figure, the IP address of router interface B is entered as the next hop (the next router that a packet traverses to reach its destination), and the cost (number of hops from source to destination) is 2.

Figure 9-6 Scenario 5: Static Route with Multiple LAN Destinations



81206

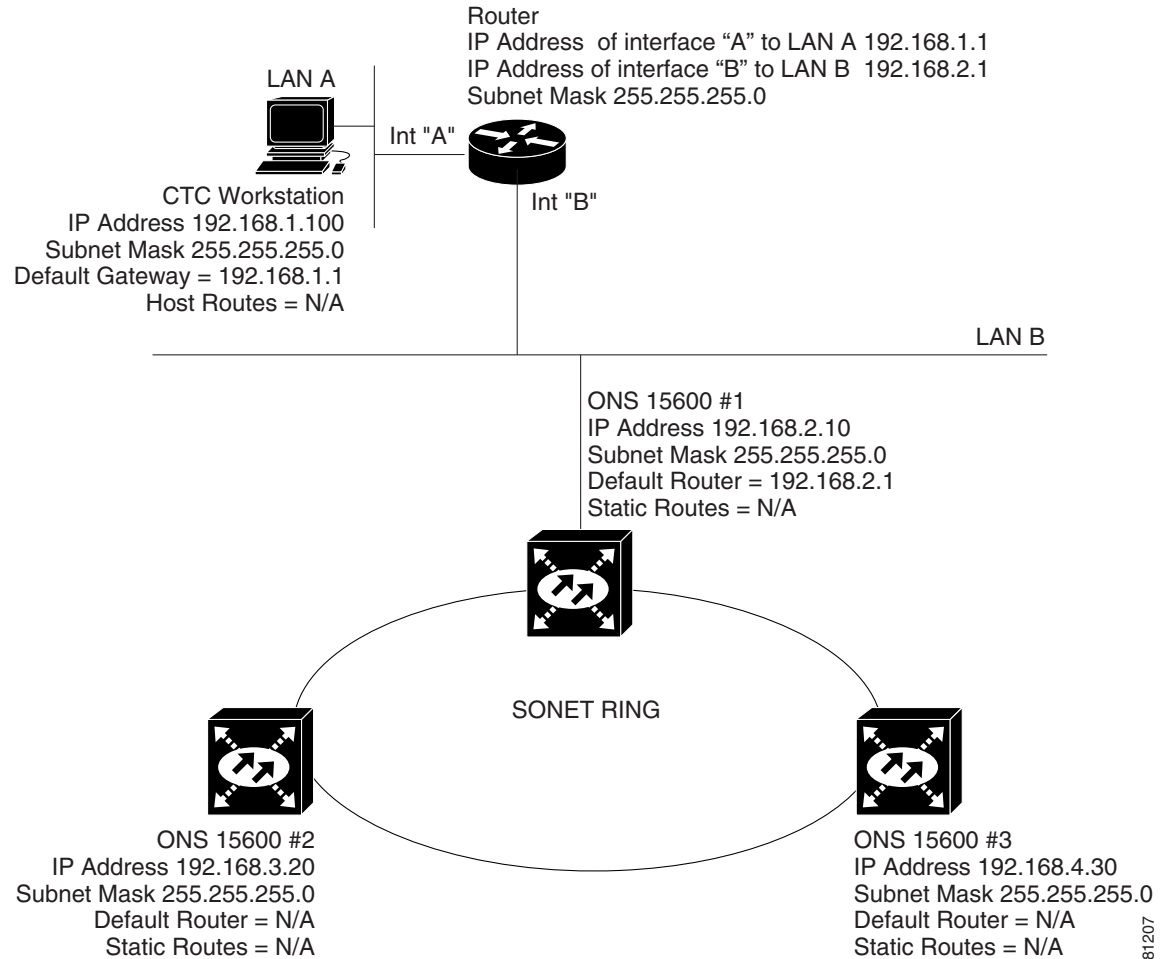
9.2.6 Scenario 6: Using OSPF

OSPF is a link state Internet routing protocol. Link state protocols use a "hello protocol" to monitor their links with adjacent routers and to test their links with their neighbors. Link state protocols advertise their directly connected networks and their active links. Each link state router captures the link state advertisements (LSAs) and puts them together to create a topology of the entire network or area. From this database, the router calculates a routing table by constructing a shortest path tree. The router continuously recalculates to capture ongoing topology changes.

ONS 15600s use the OSPF protocol in internal ONS 15600 networks for node discovery, circuit routing, and node management. You can enable OSPF on the ONS 15600s so that the ONS 15600 topology is sent to OSPF routers on a LAN. Advertising the ONS 15600 network topology to LAN routers means

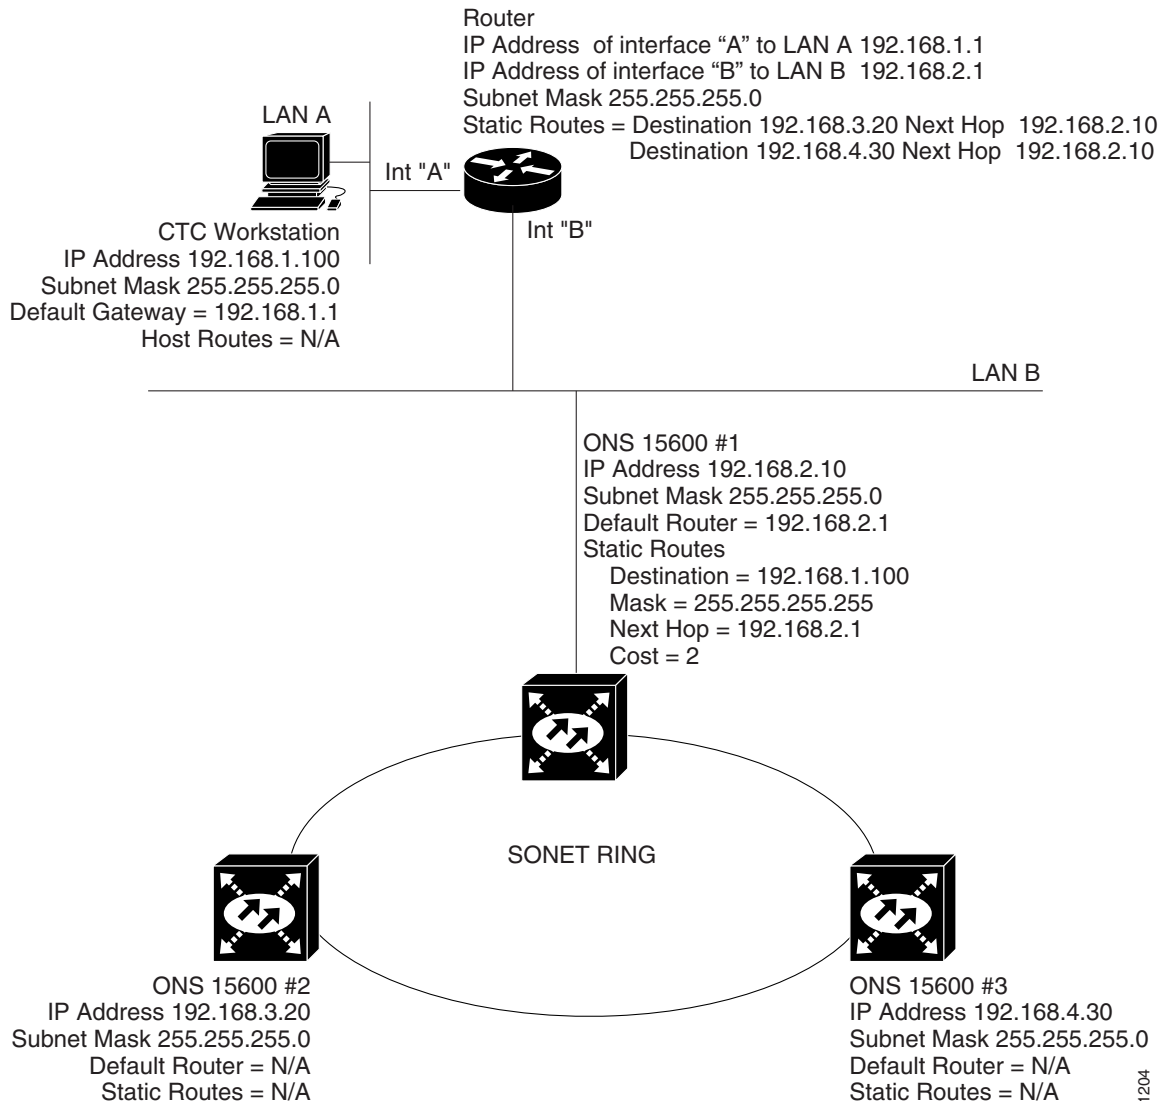
you do not need to manually enter static routes for ONS 15600 subnetworks. [Figure 9-7](#) shows the same network enabled for OSPF. When you are logged into a ONS 15600 node, CTC does not allow both a DCC interface and a LAN interface in the same nonzero OSPF area.

Figure 9-7 Scenario 6: OSPF Enabled



[Figure 9-8](#) shows the same network without OSPF. Static routes must be manually added to the router in order for CTC computers on LAN A to communicate with Nodes 2 and 3 because these nodes reside on different subnets.

Figure 9-8 Scenario 6: OSPF Not Enabled



81204

OSPF divides networks into smaller regions, called areas. An area is a collection of networked end systems, routers, and transmission facilities organized by traffic patterns. Each OSPF area has a unique ID number, known as the area ID, that can range from 0 to 4,294,967,295. Every OSPF network has one backbone area called "area 0." All other OSPF areas must connect to area 0.

When you enable ONS 15600 OSPF topology for advertising to an OSPF network, you must assign an OSPF area ID in decimal format to the ONS 15600 network. Coordinate the area ID number assignment with your LAN administrator. All DCC-connected ONS 15600s should be assigned the same OSPF area ID.

The ONS 15600 supports the multiple OSPF area feature, which allows the ability to configure and support multiple OSPF areas in each DCC-connected topology. A node is in a single OSPF area if all of its DCC or LAN interfaces are in the same OSPF area, while a node is in multiple OSPF areas if it has DCC or LAN interfaces in two or more OSPF areas. If the ONS 15600 has interfaces (DCC or LAN) in multiple OSPF areas, at least one ONS 15600 interface (DCC or LAN) must be in the backbone area 0.

If multiple ONS 15600 nodes and routers are connected to the same LAN in OSPF backbone area 0 and a link between two routers breaks, the backbone OSPF area 0 could divide into multiple gateway network elements (GNEs). If this occurs, the CTC session connected to Router 1 will not be able to communicate with the ONS 15600 connected to Router 2. To resolve, you must repair the link between the routers or provide another form of redundancy in the network. This is standard behavior for an OSPF network.

**Note**

To create OSPF virtual links, OSPF must be enabled on the LAN.

**Note**

Cisco recommends limiting the number of link-state packets (LSPs) that will be forwarded over the DCC interfaces.

9.2.7 Scenario 7: Provisioning the ONS 15600 Proxy Server

The ONS 15600 proxy server is a set of functions that allows you to configure ONS 15600s in environments where visibility and accessibility between ONS 15600s and CTC computers must be restricted. For example, you can set up a network so that field technicians and network operations center (NOC) personnel can both access the same ONS 15600s while preventing the field technicians from accessing the NOC LAN. To do this, one ONS 15600 is provisioned as a GNE and the other ONS 15600s are provisioned as ENEs. The GNE ONS 15600 tunnels connections between CTC computers and ENE ONS 15600s, providing management capability while preventing access for purposes other than ONS 15600 management.

The ONS 15600 proxy server performs the following tasks:

- Isolates DCC IP traffic from Ethernet (craft port) traffic and accepts packets based on filtering rules. The filtering rules (see [Table 9-3 on page 9-16](#) and [Table 9-4 on page 9-16](#)) depend on whether the packet arrives at the ONS 15600 DCC or TSC Ethernet interface.
- Processes Simple Network Time Protocol/Network Time Protocol (SNTP/NTP) requests. ONS 15600 ENEs can derive time-of-day from an SNTP/NTP LAN server through the ONS node GNE.
- Process SNMPv1 traps. The GNE ONS 15600 receives SNMPv1 traps from the ONS node ENEs and forwards them to all provisioned SNMPv1 trap destinations.

The ONS 15600 proxy server is provisioned using the Enable SOCKS proxy on port check box on the Provisioning > Network > General tab (see [Figure 9-9](#)). If checked, the ONS 15600 serves as a proxy for connections between CTC clients and ONS 15600s that are DCC-connected to the proxy ONS 15600. The CTC client establishes connections to DCC-connected nodes through the proxy node. The CTC client can connect to nodes that it cannot directly reach from the host on which it runs. If not selected, the node does not proxy for any CTC clients, although any established proxy connections continue until the CTC client exits. If set as a GNE, the CTC computer is visible to other DCC-connected nodes and firewall is enabled. If Proxy-only is selected, the firewall is not enabled. CTC can communicate with any other DCC-connected ONS 15600s.

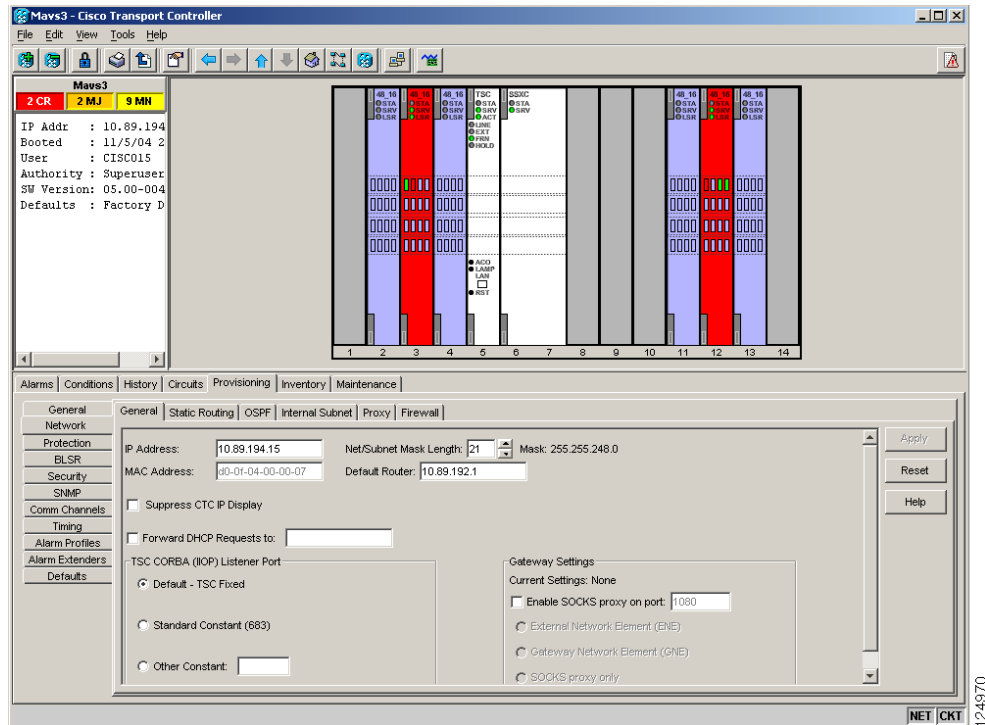
**Note**

The ONS 15600 ENE option on the Provisioning > Network > General tab behaves the same as the GNE option.

**Note**

If you launch CTC against a node through a NAT (Network Address Translation) or PAT (Port Address Translation) router and that node does not have proxy enabled, your CTC session starts and initially appears to be fine. However CTC never receives alarm updates and disconnects and reconnects every two minutes. If the proxy is accidentally disabled, it is still possible to enable the proxy during a reconnect cycle and recover your ability to manage the node, even through a NAT/PAT firewall.

Figure 9-9 Proxy Server Gateway Settings



9.2.7.1 Firewall Not Enabled

Figure 9-10 shows an ONS 15600 proxy server implementation. A ONS 15600 GNE is connected to a central office LAN and to ONS 15600 ENEs. The central office LAN is connected to a NOC LAN, which has CTC computers. The NOC CTC computer and craft technicians must both be able to access the ONS 15600 ENEs. However, the craft technicians must be prevented from accessing or seeing the NOC or central office LANs.

In the example, the ONS 15600 GNE is assigned an IP address within the central office LAN and is physically connected to the LAN through its LAN port. ONS 15600 ENEs are assigned IP addresses that are outside the central office LAN and given private network IP addresses. If the ONS 15600 ENEs are collocated, the craft LAN ports could be connected to a hub. However, the hub should have no other network connections.

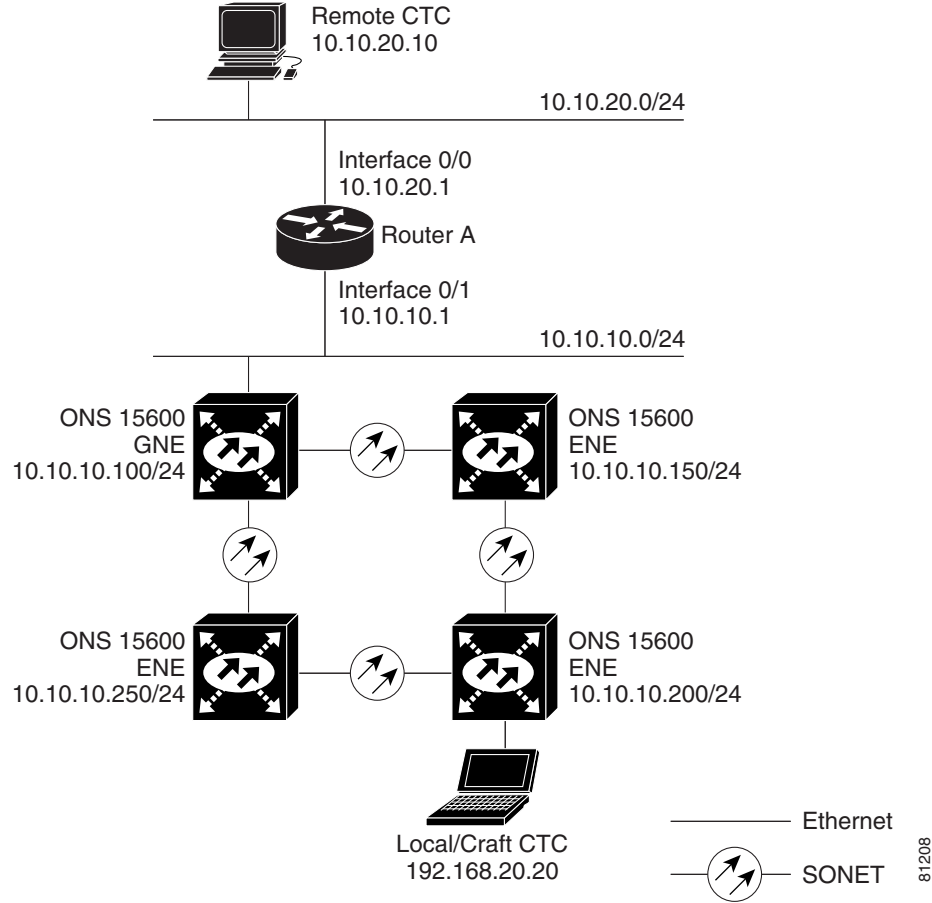
Figure 9-10 ONS 15600 Proxy Server with GNE and ENEs on the Same Subnet

Table 9-2 shows recommended settings for ONS 15600 GNEs and ENEs in the configuration shown in Figure 9-10.

Table 9-2 ONS 15600 GNE and ENE Settings

Setting	ONS 15600 GNE	ONS 15600 ENE
Enable proxy server on port	On	On
GNE	On	Off
ENE	Off	On
Proxy only	Off	Off
OSPF (LAN)	Off	Off
SNTP server (if used)	SNTP server IP address	Set to the ONS 15600 GNE IP address
SNMP (if used)	SNMPv1 trap destinations	Set SNMPv1 trap destinations to ONS 15600 GNE, port 391

Figure 9-11 shows the same proxy server implementation with ONS 15600 ENEs on different subnets. The ONS 15600 GNEs and ENEs are provisioned with the settings shown in Table 9-2.

Figure 9-11 Scenario 7: ONS 15600 Proxy Server with GNE and ENEs on Different Subnets

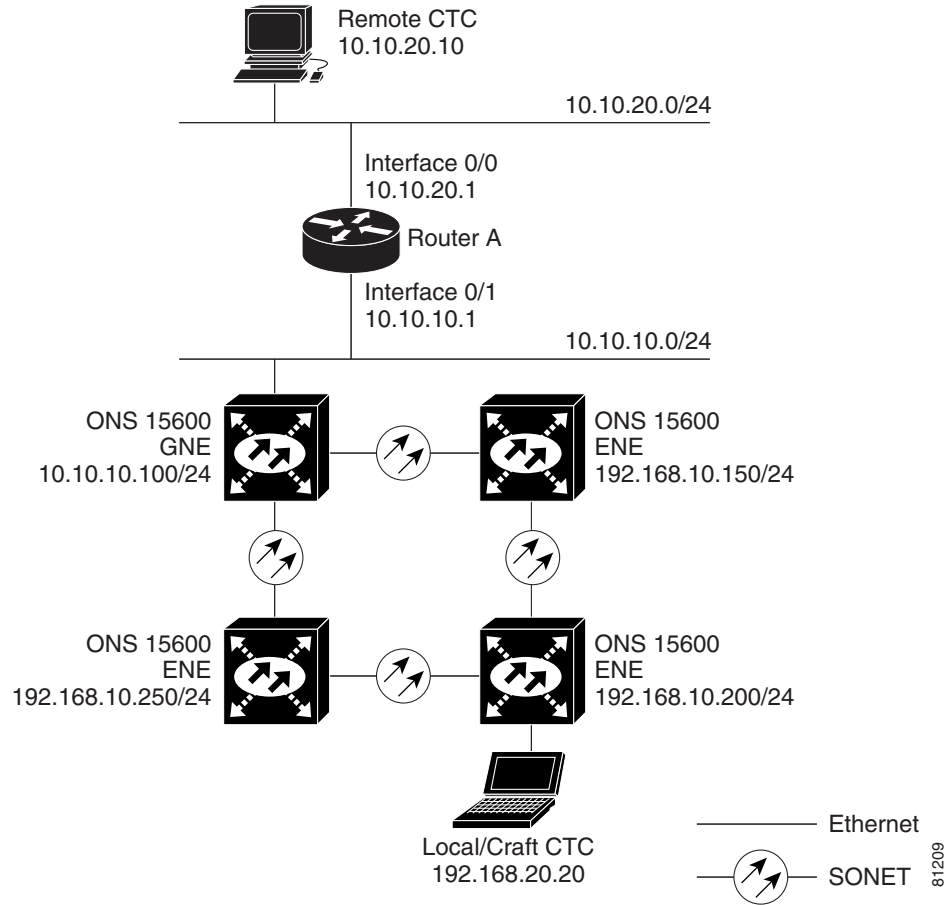
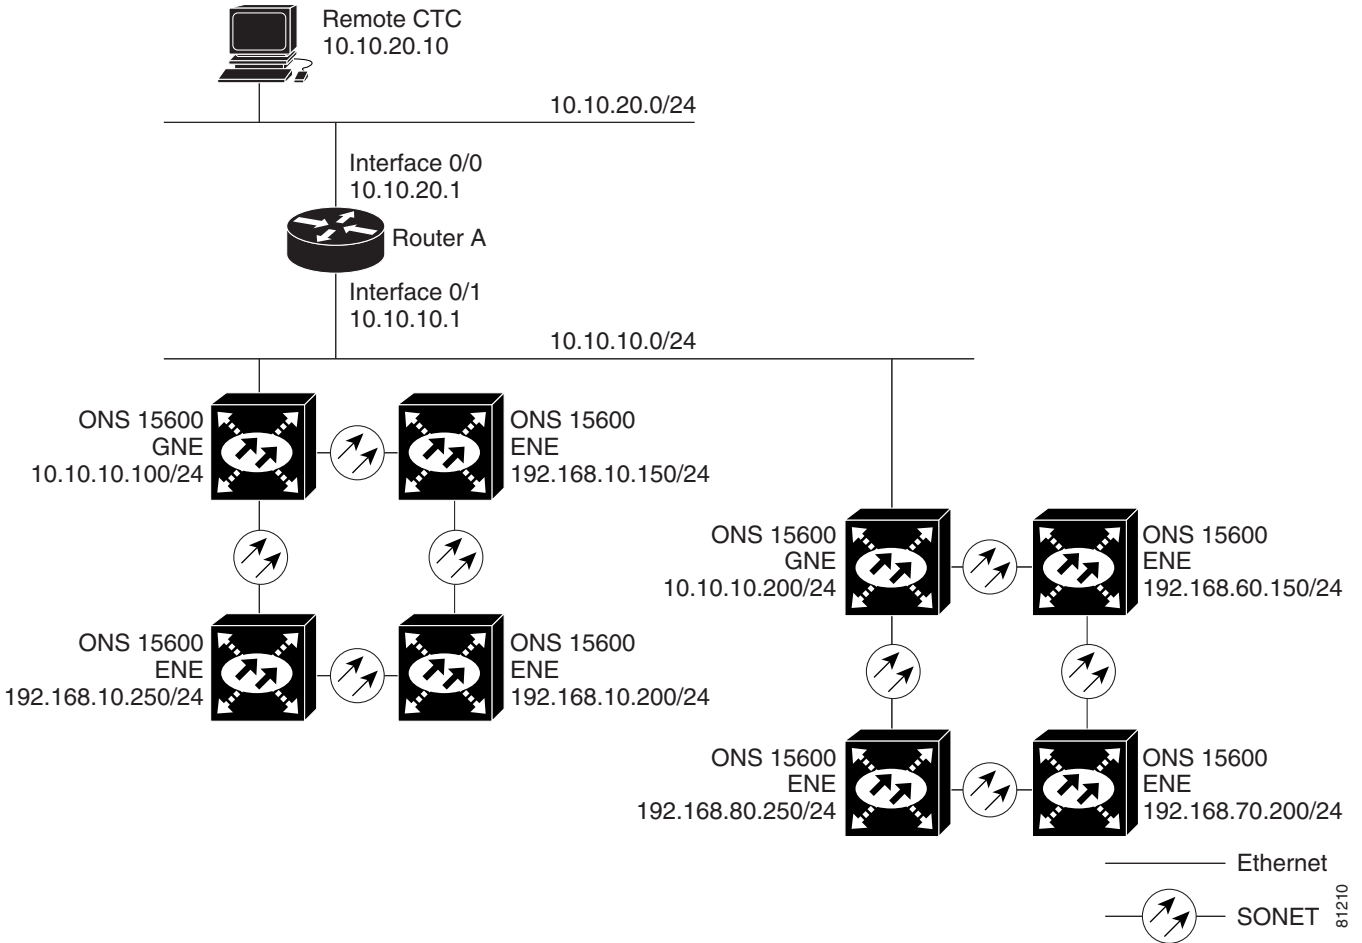


Figure 9-12 shows the Figure 9-11 implementation with ONS 15600 ENEs in multiple rings. The ONS 15600 GNEs and ENEs are provisioned with the settings shown in Table 9-2.

Figure 9-12 Scenario 7: ONS 15600 Proxy Server With ENEs on Multiple Rings



9.2.7.2 Firewall Enabled

Table 9-3 shows the rules the ONS 15600 uses to filter packets when the firewall is enabled.

Table 9-3 Proxy Server Firewall Filtering Rules

Packets Arriving At:	Are Accepted if the IP Destination Address Is:
TSC Ethernet interface	<ul style="list-style-type: none"> The ONS 15600 itself The ONS 15600 subnet broadcast address Within the 224.0.0.0/8 network (reserved network used for standard multicast messages)
DCC interface	<ul style="list-style-type: none"> The ONS 15600 itself Any destination connected through another DCC interface Within the 224.0.0.0/8 network

The rules in [Table 9-4](#) are applied if a packet is addressed to the ONS 15600. Rejected packets are discarded.

Table 9-4 Proxy Server Firewall Filtering Rules When Packet Addressed to ONS 15600

Packets Arriving At:	Accepts	Rejects
TSC Ethernet interface	<ul style="list-style-type: none"> All IP protocols except user datagram protocol (UDP) All UDP packets except packets address to the SNMP trap relay port 	<ul style="list-style-type: none"> UDP packets addressed to the SNMP trap relay port (391)
DCC interface	<ul style="list-style-type: none"> All ICMP, OSPF, RSVP, and LMP packets All TCP packets except packets addressed to the Telnet and proxy server ports 	<ul style="list-style-type: none"> TCP packets addressed to the Telnet port TCP packets addressed to the proxy server port Protocols not listed in the Accepted column

If an ONS 15600 or CTC computer resides behind a firewall that uses port filtering, you must enable an Internet Inter-ORB Protocol (IIOP) port on the ONS 15600 and/or CTC computer, depending on whether one or both devices reside behind a firewall. You can enable an IIOP port on the Provisioning > Network > General tabs in CTC.

[Figure 9-13](#) shows ONS 15600s in a protected network and the CTC computer in an external network. For the computer to access the ONS 15600s, you must provision the IIOP listener port specified by your firewall administrator on the ONS 15600. The ONS 15600 sends the port number to the CTC computer during the initial contact between the devices using Hyper-Text Transfer Protocol (HTTP). After the CTC computer obtains the ONS 15600 IIOP port, the computer opens a direct session with the node using the specified IIOP port.

Figure 9-13 Nodes Behind a Firewall


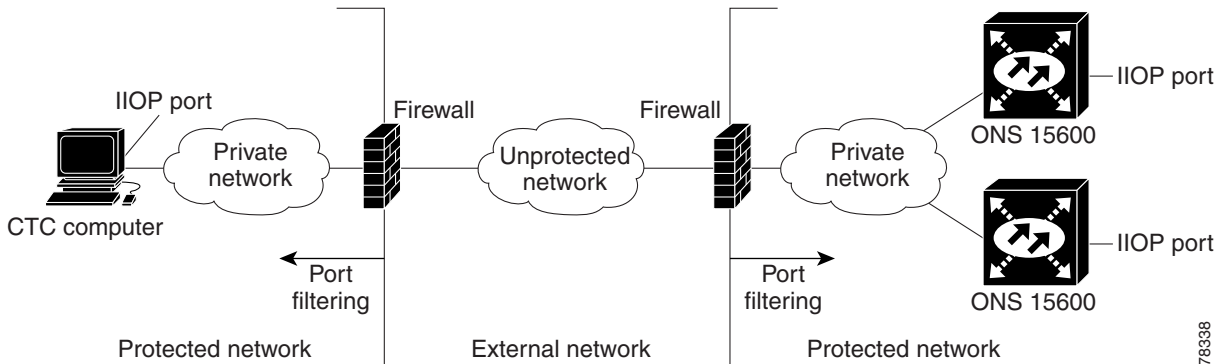


Figure 9-14 shows a CTC computer and ONS 15600s behind firewalls. For the computer to access the ONS 15600, you must provision the IIO port on the CTC computer and on the ONS 15600. Each firewall can use a different IIO port. For example, if the CTC computer firewall uses IIO port 4000, and the ONS 15600 firewall uses IIO port 5000, 4000 is the IIO port you provision for the CTC computer and 5000 is the IIO port you provision for the ONS 15600.

Figure 9-14 CTC Computer and ONS 15600s Residing Behind Firewalls



If you implement the proxy server, note that all DCC-connected ONS 15600s on the same Ethernet segment must have the same gateway setting. Mixed values produce unpredictable results, and might leave some nodes unreachable through the shared Ethernet segment.

If nodes become unreachable, correct the setting by performing one of the following actions:

- Disconnect the craft computer from the unreachable ONS 15600. Connect to the ONS 15600 through another network ONS 15600 that has a DCC connection to the unreachable ONS 15600.
- Disconnect all DCCs to the node by disabling them on neighboring nodes. Connect a CTC computer directly to the ONS 15600 and change its provisioning.

9.2.8 Scenario 8: Dual GNEs on a Subnet

The ONS 15600 provides GNE load balancing, which allows CTC to reach ENEs over multiple GNEs without the ENEs being advertised over OSPF. This feature allows a network to quickly recover from the loss of a GNE, even if the GNE is on a different subnet. If a GNE fails, all connections through that GNE fail. CTC disconnects from the failed GNE and from all ENEs for which the GNE was a proxy and

then reconnects through the remaining GNEs. GNE load balancing reduces the dependency on the launch GNE and DCC bandwidth, which enhances CTC performance. Figure 9-15 shows a network with dual GNEs on the same subnet.

Figure 9-15 Scenario 8: Dual GNEs on the Same Subnet


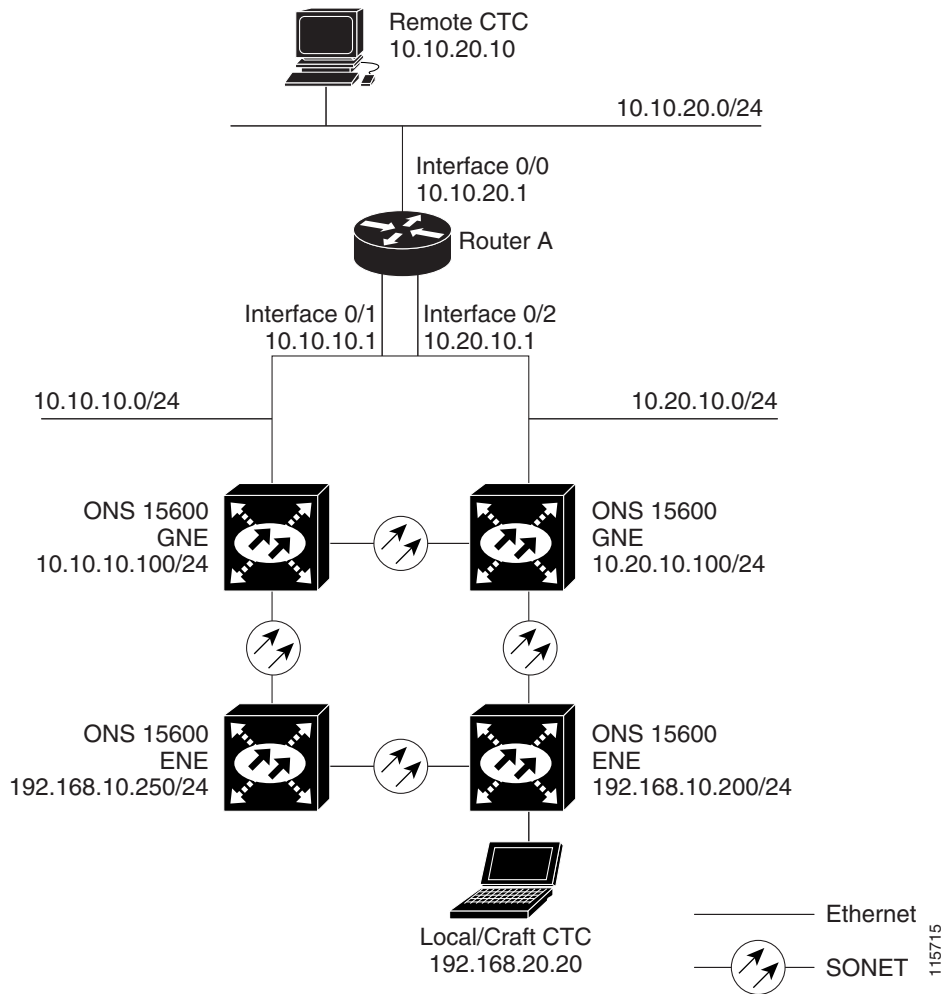


Figure 9-16 shows a network with dual GNEs on different subnets.

Figure 9-16 Scenario 8: Dual GNEs on Different Subnets



9.3 Provisionable Patchcords

A provisionable patchcord is a user-provisioned link that is advertised by OSPF throughout the network. Provisionable patchcords, also called virtual links, are needed if an ONS 15600 optical port is connected to an ONS 15454 transponder or muxponder client port provisioned in transparent mode. Provisionable patchcords are required on both ends of a physical link. The provisioning at each end includes a local patchcord ID, slot and port information, remote IP address, and remote patchcord ID. Patchcords appear as dashed lines in CTC network view.

Table 9-5 lists the supported combinations for ONS 15600 optical cards and the ONS 15454 transponder/muxponder cards used in a provisionable patchcord. For more information about the ONS 15454 transponder and muxponder cards, refer to the *Cisco ONS 15454 DWDM Reference Manual*.

Table 9-5 ONS 15600 and ONS 15454 Client and Trunk Card Combinations in Provisionable Patchcords

ONS 15600 Trunk Cards	ONS 15454 Client Cards		
	MXP_2.5G_10G/ TXP_MR_10G	TXP(P)_MR_2.5G	MXP_2.5G_10E/ TXP_MR_10E
OC-192	Yes	—	Yes
OC-48	Yes	Yes	Yes
ASAP (OC-48 PPM)	Yes	Yes	Yes

Optical ports have the following requirements when used in a provisionable patchcord:

- An optical port connected to an ONS 15454 transponder or muxponder port requires an SDCC or LDCC termination.
- If the optical port is the protection port in a 1+1 group, the working port must have an SDCC or LDCC termination provisioned.
- If a remote end (ONS 15454) of a provisionable patchcord is Y-cable protected, an optical port requires two patchcords.

9.4 Routing Table

ONS 15600 routing information appears on the Maintenance > Routing Table tabs (Figure 9-17). The routing table provides the following information:

- Destination—Displays the IP address of the destination network or host.
- Mask—Displays the subnet mask used to reach the destination host or network.
- Gateway—Displays the IP address of the gateway used to reach the destination network or host.
- Usage—Shows the number of times this route has been used.
- Interface—Shows the ONS 15600 interface used to access the destination.
 - pend0—The Ethernet management interface.
 - pdcc—An SDCC interface, that is, an OC-N trunk (span) card identified as the SDCC termination (0 to 128).
 - lo0—A loopback interface.
 - pend2—The craft-only RJ-45 jack on the front of the TSC.
 - motfcc0—Interface on the TSC that connect the TSC to all other cards except the other TSC.
 - hdlc0—Connects the two TSC cards together; traffic cards forward DCC packets over the motfcc0 Ethernet interface.

Figure 9-17 Viewing the ONS 15600 Routing Table

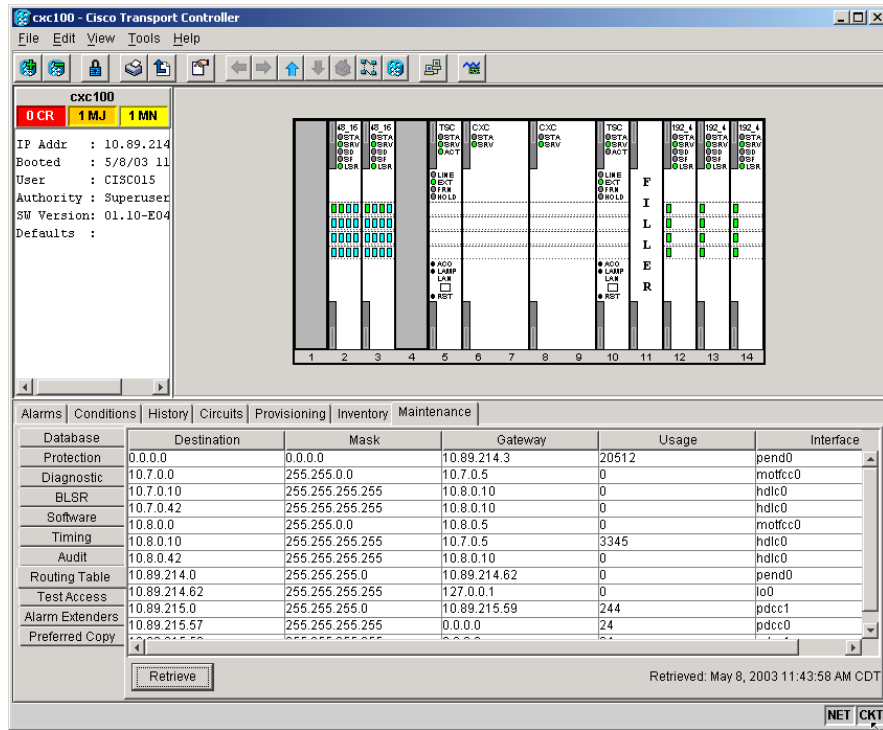


Table 9-6 shows sample routing entries for an ONS 15600.

Table 9-6 Sample Routing Table Entries

Entry	Destination	Mask	Gateway	Interface
1	0.0.0.0	0.0.0.0	172.20.214.1	cpm0
2	172.20.214.0	255.255.255.0	172.20.214.92	cpm0
3	172.20.214.92	255.255.255.255	127.0.0.1	lo0
4	172.20.214.93	255.255.255.255	0.0.0.0	pdcc0
5	172.20.214.94	255.255.255.255	172.20.214.93	pdcc0

Entry 1 shows the following:

- Destination (0.0.0.0) is the default route entry. All undefined destination network or host entries on this routing table will be mapped to the default route entry.
- Mask (0.0.0.0) is always 0 for the default route.
- Gateway (172.20.214.1) is the default gateway address. All outbound traffic that cannot be found in this routing table or is not on the node’s local subnet will be sent to this gateway.
- Interface (cpm0) indicates that the ONS 15600 Ethernet interface is used to reach the gateway.

Entry 2 shows the following:

- Destination (172.20.214.0) is the destination network IP address.
- Mask (255.255.255.0) is a 24-bit mask, meaning all addresses within the 172.20.214.0 subnet can be destinations.

- Gateway (172.20.214.92) is the gateway address. All outbound traffic belonging to this network is sent to this gateway.
- Interface (cpm0) indicates that the ONS 15600 Ethernet interface is used to reach the gateway.

Entry 3 shows the following:

- Destination (172.20.214.92) is the destination host IP address.
- Mask (255.255.255.255) is a 32-bit mask, meaning only the 172.20.214.92 address is a destination.
- Gateway (127.0.0.1) is a loopback address. The host directs network traffic to itself using this address.
- Interface (lo0) indicates that the local loopback interface is used to reach the gateway.

Entry 4 shows the following:

- Destination (172.20.214.93) is the destination host IP address.
- Mask (255.255.255.255) is a 32-bit mask, meaning only the 172.20.214.93 address is a destination.
- Gateway (0.0.0.0) means the destination host is directly attached to the node.
- Interface (pdcc0) indicates that a SONET DCC interface is used to reach the destination host.

Entry 5 shows a DCC-connected node that is accessible through a node that is not directly connected:

- Destination (172.20.214.94) is the destination host IP address.
- Mask (255.255.255.255) is a 32-bit mask, meaning only the 172.20.214.94 address is a destination.
- Gateway (172.20.214.93) indicates that the destination host is accessed through a node with the IP address 172.20.214.93.
- Interface (pdcc0) indicates that a SONET DCC interface is used to reach the gateway.

9.5 External Firewalls

This section provides sample access control lists for external firewalls. [Table 9-7](#) lists the ports that are used by the TSC.

Table 9-7 *Ports Used by the TSC*

Port	Function	Action ¹
0	Never used	D
20	FTP	D
21	FTP control	D
22	SSH (Secure Shell)	D
23	Telnet	D
80	HTTP	D
111	SUNRPC (Sun Remote Procedure Call)	D
161	SNMP traps destinations	D
162	SNMP traps destinations	D
513	rlogin	D
683	CORBA IIOP	OK

Table 9-7 Ports Used by the TSC (continued)

Port	Function	Action ¹
1080	Proxy server (socks)	D
2001-2017	I/O card Telnet	NA
2018	DCC processor on active TCC2/TCC2P	D
2361	TL1	D
3082	Raw TL1	D
3083	TL1	D
5001	BLSR server port	D
5002	BLSR client port	D
7200	SNMP alarm input port	D
9100	EQM port	D
9401	TCC boot port	D
9999	Flash manager	NA
10240-12287	Proxy client	D
57790	Default TCC listener port	OK

1. D = deny, NA = not applicable, OK = do not deny

The following ACL (access control list) example shows a firewall configuration when the SOCKS proxy server gateway setting is not enabled. In the example, the CTC workstation's address is 192.168.10.10, and the ONS 15600 address is 10.10.10.100. The firewall is attached to the GNE, so the inbound direction is from CTC to the GNE and the outbound direction is from the GNE to CTC. The CTC Common Object Request Broker Architecture (CORBA) Standard constant is 683 and the TCC CORBA Default is TCC Fixed (57790).

```
access-list 100 remark *** Inbound ACL, CTC -> NE ***
access-list 100 remark
access-list 100 permit tcp host 192.168.10.10 host 10.10.10.100 eq www
access-list 100 remark *** allows initial contact with ONS 15600 using http (port 80) ***
access-list 100 remark
access-list 100 permit tcp host 192.168.10.10 host 10.10.10.100 eq 57790
access-list 100 remark *** allows CTC communication with ONS 15600 GNE (port 57790) ***
access-list 100 remark
access-list 100 permit tcp host 192.168.10.10 host 10.10.10.100 established
access-list 100 remark *** allows ACKs back from CTC to ONS 15600 GNE ***

access-list 101 remark *** Outbound ACL, NE -> CTC ***
access-list 101 remark
access-list 101 permit tcp host 10.10.10.100 host 192.168.10.10 eq 683
access-list 101 remark *** allows alarms etc., from the 15600 (random port) to the CTC
workstation (port 683) ***
access-list 100 remark
access-list 101 permit tcp host 10.10.10.100 host 192.168.10.10 established
access-list 101 remark *** allows ACKs from the 15600 GNE to CTC ***
```

The following ACL (access control list) example shows a firewall configuration when the SOCKS proxy server gateway setting is enabled. As with the first example, the CTC workstation address is 192.168.10.10 and the ONS 15600 address is 10.10.10.100. The firewall is attached to the GNE, so inbound is CTC to the GNE and outbound is from the GNE to CTC. CTC CORBA Standard constant (683) and TCC CORBA Default is TCC Fixed (57790).

```

access-list 100 remark *** Inbound ACL, CTC -> NE ***
access-list 100 remark
access-list 100 permit tcp host 192.168.10.10 host 10.10.10.100 eq www
access-list 100 remark *** allows initial contact with the 15600 using http (port 80) ***
access-list 100 remark
access-list 100 permit tcp host 192.168.10.10 host 10.10.10.100 eq 1080
access-list 100 remark *** allows CTC communication with the 15600 GNE (port 1080) ***
access-list 100 remark

access-list 101 remark *** Outbound ACL, NE -> CTC ***
access-list 101 remark
access-list 101 permit tcp host 10.10.10.100 host 192.168.10.10 established
access-list 101 remark *** allows ACKs from the 15600 GNE to CTC ***

```

9.6 Open GNE

The ONS 15600 can communicate with non-ONS nodes that do not support point-to-point protocol (PPP) vendor extensions or OSPF type 10 opaque link-state advertisements (LSA), both of which are necessary for automatic node and link discovery. An open GNE configuration allows the DCC-based network to function as an IP network for non-ONS nodes.

To configure an open GNE network, you can provision SDCC and LDCC terminations to include a far-end, non-ONS node using either the default IP address of 0.0.0.0 or a specified IP address. You provision a far-end, non-ONS node by checking the “Far End is Foreign” check box during SDCC and LDCC creation. The default 0.0.0.0 IP address allows the far-end, non-ONS node to provide the IP address; if you set an IP address other than 0.0.0.0, a link is established only if the far-end node identifies itself with that IP address, providing an extra level of security.

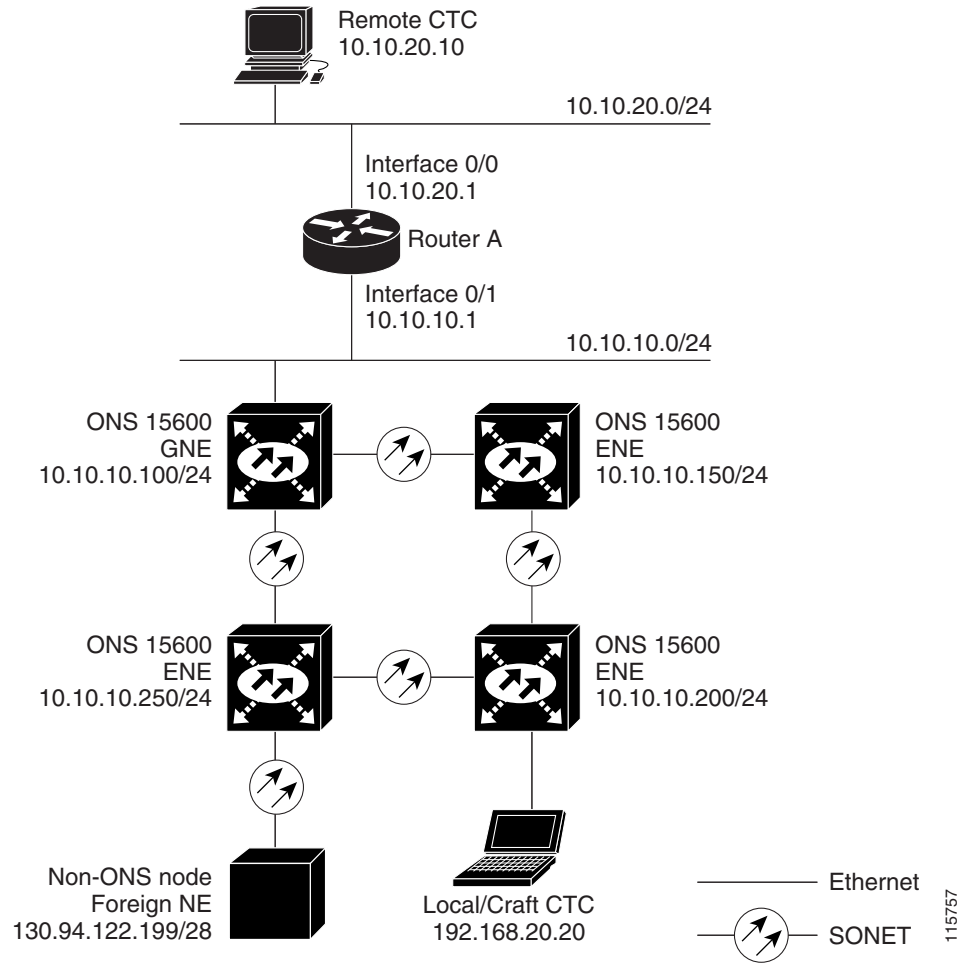
By default, the proxy server only allows connections to discovered ONS peers and the firewall blocks all IP traffic between the DCC network and LAN. You can, however, provision proxy tunnels to allow up to 12 additional destinations for SOCKS version 5 connections to non-ONS nodes. You can also provision firewall tunnels to allow up to 12 additional destinations for direct IP connectivity between the DCC network and LAN. Proxy and firewall tunnels include both a source and destination subnet. The connection must originate within the source subnet and terminate within the destination subnet before either the SOCKS connection or IP packet flow is allowed.

To set up proxy and firewall subnets in CTC, use the Provisioning > Network > Proxy and Firewalls subtabs. The availability of proxy and/or firewall tunnels depends on the network access settings of the node:

- If the node is configured with the proxy server enabled in GNE or ENE mode, you must set up a proxy tunnel and/or a firewall tunnel.
- If the node is configured with the proxy server enabled in proxy-only mode, you can set up proxy tunnels. Firewall tunnels are not allowed.
- If the node is configured with the proxy server disabled, neither proxy tunnels or firewall tunnels are allowed.

Figure 9-18 shows an example of a foreign node connected to the DCC network. Proxy and firewall tunnels are useful in this example because the GNE would otherwise block IP access between the PC and the foreign node.

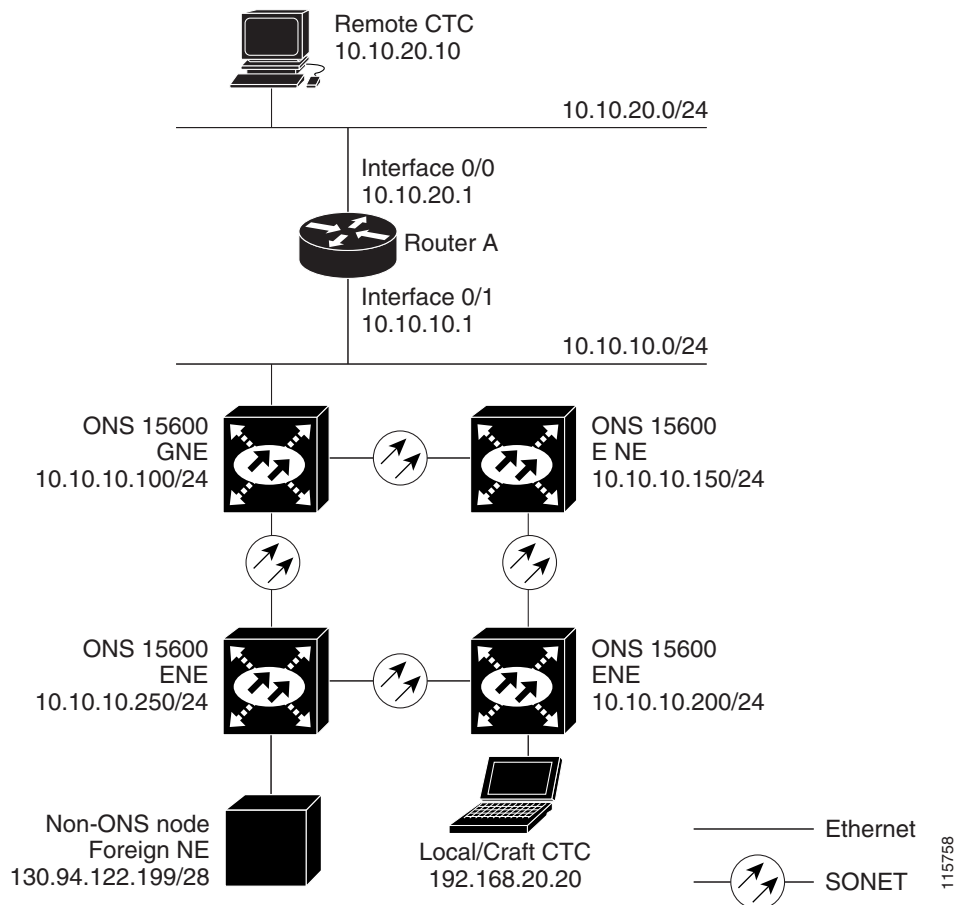
Figure 9-18 Proxy and Firewall Tunnels for Foreign Terminations



115757

Figure 9-19 shows a remote node connected to an ENE Ethernet port. Proxy and firewall tunnels are useful in this example because the GNE would otherwise block IP access between the PC and foreign node. This configuration also requires a firewall tunnel on the ENE.

Figure 9-19 Foreign Node Connection to an ENE Ethernet Port



9.7 TCP/IP and OSI Networking

ONS 15600 DCN communication is based on the TCP/IP protocol suite. However, ONS 15600s can also be networked with equipment that uses the OSI protocol suite. While TCP/IP and OSI protocols are not directly compatible, they do have the same objectives and occupy similar layers of the OSI reference model. Table 9-8 shows the protocols and mediation processes that are involved when TCP/IP-based NEs are networked with OSI-based NEs.

Table 9-8 TCP/IP and OSI Protocols

OSI Model	IP Protocols	OSI Protocols	IP-OSI Mediation	
Layer 7 Application	<ul style="list-style-type: none"> • TL1 • FTP • HTTP • Telnet • IOP 	<ul style="list-style-type: none"> • TARP¹ 	<ul style="list-style-type: none"> • TL1 (over OSI) • FTAM² • ACSE³ 	<ul style="list-style-type: none"> • T-TD⁴ • FT-TD⁵
Layer 6 Presentation			<ul style="list-style-type: none"> • PST⁶ 	—
Layer 5 Session			<ul style="list-style-type: none"> • Session 	—
Layer 4 Transport			<ul style="list-style-type: none"> • TCP • UDP 	<ul style="list-style-type: none"> • TP (Transport Class 4)
Layer 3 Network	<ul style="list-style-type: none"> • IP • OSPF 	<ul style="list-style-type: none"> • CLNP⁸ • ES-IS⁹ • IS-IS¹⁰ 		
Layer 2 Data link	<ul style="list-style-type: none"> • PPP 	<ul style="list-style-type: none"> • PPP • LAP-D¹¹ 		
Layer 1 Physical	DCC, LAN, fiber, electrical	DCC, LAN, fiber, electrical	—	

1. TARP = TID Address Resolution Protocol
2. FTAM = File Transfer and Access Management
3. ACSE = Association-control service element
4. T-TD = TL1-Translation Device
5. FT-TD = File Transfer-Translation Device
6. PST = Presentation layer
7. CLNS = Connectionless Network Layer Service
8. CLNP = Connectionless Network Layer Protocol
9. ES-IS = End System-to-Intermediate System
10. IS-IS = Intermediate System-to-Intermediate System
11. LAP-D = Link Access Protocol on the D Channel

9.7.1 Point-to-Point Protocol

PPP is a data link (Layer 2) encapsulation protocol that transports datagrams over point-to-point links. Although PPP was developed to transport IP traffic, it can carry other protocols including the OSI CLNP. PPP components used in the transport of OSI include:

- High-level data link control (HDLC)—Performs the datagram encapsulation for transport across point-to-point links.
- Link control protocol (LCP)—Establishes, configures, and tests the point-to-point connections.

CTC automatically enables IP over PPP whenever you create an SDCC or LDCC. The SDCC or LDCC can be provisioned to support OSI over PPP.

9.7.2 Link Access Protocol on the D Channel

LAP-D is a data link protocol used in the OSI protocol stack. LAP-D is assigned when you provision an ONS 15600 SDCC as OSI-only. Provisionable LAP-D parameters include:

- Transfer Service—One of the following transfer services must be assigned:
 - Acknowledged Information Transfer Service (AITS)—(Default) Does not exchange data until a logical connection between two LAP-D users is established. This service provides reliable data transfer, flow control, and error control mechanisms.
 - Unacknowledged Information Transfer Service (UITS)—Transfers frames containing user data with no acknowledgement. The service does not guarantee that the data presented by one user will be delivered to another user, nor does it inform the user if the delivery attempt fails. It does not provide any flow control or error control mechanisms.
- Mode—LAP-D is set to either Network or User mode. This parameter sets the LAP-D frame command/response (C/R) value, which indicates whether the frame is a command or a response.
- Maximum transmission unit (MTU)—The LAP-D N201 parameter sets the maximum number of octets in a LAP-D information frame. The range is 512 to 1500 octets.



Note The MTU must be the same size for all NEs on the network.

- Transmission Timers—The following LAP-D timers can be provisioned:
 - The T200 timer sets the timeout period for initiating retries or declaring failures.
 - The T203 timer provisions the maximum time between frame exchanges, that is, the trigger for transmission of the LAP-D “keep-alive” Receive Ready (RR) frames.

Fixed values are assigned to the following LAP-D parameters:

- Terminal Endpoint Identifier (TEI)—A fixed value of 0 is assigned.
- Service Access Point Identifier (SAPI)—A fixed value of 62 is assigned.
- N200 supervisory frame retransmissions—A fixed value of 3 is assigned.

9.7.3 OSI Connectionless Network Service

OSI connectionless network service is implemented by using the Connectionless Network Protocol (CLNP) and Connectionless Network Service (CLNS). CLNP and CLNS are described in the ISO 8473 standard. CLNS provides network layer services to the transport layer through CLNP. CLNS does not perform connection setup or termination because paths are determined independently for each packet that is transmitted through a network. CLNS relies on transport layer protocols to perform error detection and correction.

CLNP is an OSI network layer protocol that carries upper-layer data and error indications over connectionless links. CLNP provides the interface between the CLNS and upper layers. CLNP performs many of the same services for the transport layer as IP. The CLNP datagram is very similar to the IP datagram. It provides mechanisms for fragmentation (data unit identification, fragment/total length, and offset). Like IP, a checksum computed on the CLNP header verifies that the information used to process the CLNP datagram is transmitted correctly, and a lifetime control mechanism (Time to Live) limits the amount of time a datagram is allowed to remain in the system.

CLNP uses network service access points (NSAPs) to identify network devices. The CLNP source and destination addresses are NSAPs. In addition, CLNP uses a network element title (NET) to identify a network-entity in an end system (ES) or intermediate system (IS). NETs are allocated from the same name space as NSAP addresses. Whether an address is an NSAP address or a NET depends on the network selector value in the NSAP.

The ONS 15600 supports the ISO Data Country Code (ISO-DCC) NSAP address format as specified in ISO 8348. The NSAP address is divided into an initial domain part (IDP) and a domain-specific part (DSP). NSAP fields are shown in Table 9-9. NSAP field values are in hexadecimal format. All NSAPs are editable. Shorter NSAPs can be used. However NSAPs for all NEs residing within the same OSI network area usually have the same NSAP format.

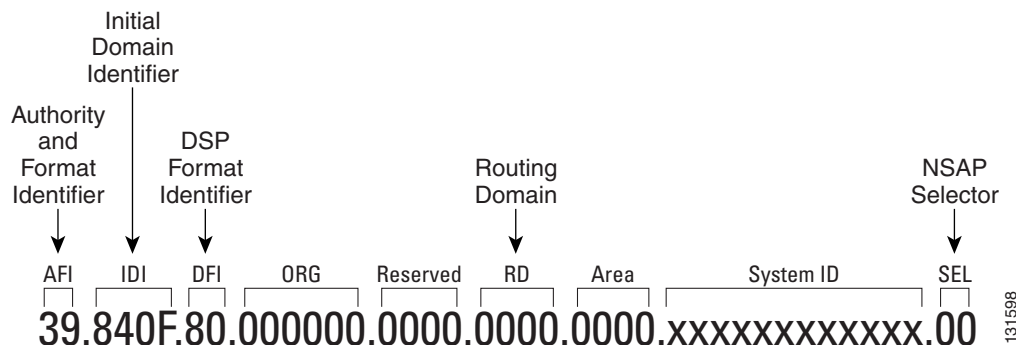
Table 9-9 NSAP Fields

Field	Definition	Description
IDP		
AFI	Authority and format identifier	Specifies the NSAP address format. The initial value is 39 for the ISO-DCC address format.
IDI	Initial domain identifier	Specifies the country code. The initial value is 840F, the United States country code padded with an F.
DSP		
DFI	DSP format identifier	Specifies the DSP format. The initial value is 80, indicating the DSP format follows American National Standards Institute (ANSI) standards.
ORG	Organization	Organization identifier. The initial value is 000000.
Reserved	Reserved	Reserved NSAP field. The Reserved field is normally all zeros (0000).
RD	Routing domain	Defines the routing domain. The initial value is 0000.
AREA	Area	Identifies the OSI routing area to which the node belongs. The initial value is 0000.

Table 9-9 NSAP Fields (continued)

Field	Definition	Description
System	System identifier	The ONS 15600 system identifier is set to its IEEE 802.3 MAC address. Each ONS 15600 supports twelve OSI virtual routers. Each router NSAP system identifier is the ONS 15600 IEEE 802.3 MAC address + n , where $n = 0$ to 2. For the primary virtual router, $n = 0$.
SEL	Selector	<p>The selector field directs the protocol data units (PDUs) to the correct destination using the CLNP network layer service. Selector values supported by the ONS 15600 include:</p> <ul style="list-style-type: none"> 00—Network Entity Title (NET). Used to exchange PDUs in the ES-IS and IS-IS routing exchange protocols. (See the “9.7.4.1 End System-to-Intermediate System Protocol” section on page 9-32, and “9.7.4.2 Intermediate System-to-Intermediate System” section on page 9-32.) 1D—Selector for Transport Class 4 (and for FTAM and TL1 applications (Telcordia GR-253-CORE standard) AF—Selector for the TARP protocol (Telcordia GR-253-CORE standard) 2F—Selector for the GRE IP-over-CLNS tunnel (ITU/RFC standard) CC—Selector for the Cisco IP-over-CLNS tunnels (Cisco specific) E0—Selector for the OSI ping application (Cisco specific) <p>NSELs are only advertised when the node is configured as an ES. They are not advertised when a node is configured as an IS. Tunnel NSELs are not advertised until a tunnel is created.</p>

Figure 9-20 shows the ISO-DCC NSAP address with the default values delivered with the ONS 15600. The System ID is automatically populated with the node MAC address.

Figure 9-20 ISO-DCC NSAP Address

The ONS 15600 main NSAP address is shown on the node view Provisioning > OSI > Main Setup subtab. This address is also the Router 1 primary manual area address, which is viewed and edited on Provisioning > OSI > Routers subtab. See the “[9.7.7 OSI Virtual Routers](#)” section on page 9-37 for information about the OSI router and manual area addresses in CTC.

9.7.4 OSI Routing

OSI architecture includes ESs and ISs. The OSI routing scheme includes:

- A set of routing protocols that allow ESs and ISs to collect and distribute the information necessary to determine routes. Protocols include the ES-IS and IS-IS protocols. ES-IS routing establishes connectivity and reach ability among ESs and ISs attached to the same (single) subnetwork.
- A routing information base (RIB) containing this information, from which routes between ESs can be computed. The RIB consists of a table of entries that identify a destination (for example, an NSAP), the subnetwork over which packets should be forwarded to reach that destination, and a routing metric. The routing metric communicates characteristics of the route (such as delay properties or expected error rate) that are used to evaluate the suitability of a route compared to another route with different properties, for transporting a particular packet or class of packets.
- A routing algorithm, Shortest Path First (SPF), that uses information contained in the RIB to derive routes between ESs.


In OSI networking, discovery is based on announcements. An ES uses the ES-IS protocol end system hello (ESH) message to announce its presence to ISs and ESs connected to the same network. Any ES or IS that is listening for ESHs gets a copy. ISs store the NSAP address and the corresponding subnetwork address pair in routing tables. ESs might store the address, or they might wait to be informed by ISs when they need such information.

An IS composes intermediate system hello (ISH) messages to announce its configuration information to ISs and ESs that are connected to the same broadcast subnetwork. Like the ESHs, the ISH contains the addressing information for the IS (the NET and the subnetwork point-of-attachment address [SNPA]) and a holding time. ISHs might also communicate a suggested ES configuration time recommending a configuration timer to ESs.

The exchange of ISHs is called neighbor greeting or initialization. Each router learns about the other routers with which they share direct connectivity. After the initialization, each router constructs a link-state packet (LSP). The LSP contains a list of the names of the IS's neighbors and the cost to reach each of the neighbors. Routers then distribute the LSPs to all of the other routers. When all LSPs are propagated to all routers, each router has a complete map of the network topology (in the form of LSPs). Routers use the LSPs and the SPF algorithm to compute routes to every destination in the network.

OSI networks are divided into areas and domains. An area is a group of contiguous networks and attached hosts that is designated as an area by a network administrator. A domain is a collection of connected areas. Routing domains provide full connectivity to all ESs within them. Routing within the same area is known as Level 1 routing. Routing between two areas is known as Level 2 routing. LSPs that are exchanged within a Level 1 area are called L1 LSPs. LSPs that are exchanged across Level 2 areas are called L2 LSPs. [Figure 9-21](#) shows an example of Level 1 and Level 2 routing.

Figure 9-21 Level 1 and Level 2 OSI Routing



When you provision an ONS 15600 for a network with NEs that use both the TCP/IP and OSI protocol stacks, you will provision it as one of the following:

- Intermediate System Level 1—The ONS 15600 performs OSI IS functions. It communicates with IS and ES nodes that reside within its OSI area. It depends upon an IS L1/L2 node to communicate with IS and ES nodes that reside outside its OSI area.
- Intermediate System Level 1/Level 2—The ONS 15600 performs IS functions. It communicates with IS and ES nodes that reside within its OSI area. It also communicates with IS L1/L2 nodes that reside in other OSI areas. This option should not be provisioned unless the node is connected to another IS L1/L2 node that resides in a different OSI area. The node must also be connected to all nodes within its area that are provisioned as IS L1/L2.

9.7.4.1 End System-to-Intermediate System Protocol

ES-IS is an OSI protocol that defines how ESs (hosts) and ISs (routers) learn about each other. ES-IS configuration information is transmitted at regular intervals through the ES and IS hello messages. The hello messages contain the subnetwork and network layer addresses of the systems that generate them.

The ES-IS configuration protocol communicates both OSI network layer addresses and OSI subnetwork addresses. OSI network layer addresses identify either the NSAP, which is the interface between OSI Layer 3 and Layer 4, or the NET, which is the network layer entity in an OSI IS. OSI SNPAs are the points at which an ES or IS is physically attached to a subnetwork. The SNPA address uniquely identifies each system attached to the subnetwork. In an Ethernet network, for example, the SNPA is the 48-bit MAC address. Part of the configuration information transmitted by ES-IS is the NSAP-to-SNPA or NET-to-SNPA mapping.

9.7.4.2 Intermediate System-to-Intermediate System

IS-IS is an OSI link-state hierarchical routing protocol that floods the network with link-state information to build a complete, consistent picture of a network topology. IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Level 1 ISs need to know only how to get to the nearest Level 2 IS. The backbone routing protocol can change without impacting the intra-area routing protocol.

OSI routing begins when the ESs discover the nearest IS by listening to ISH packets. When an ES wants to send a packet to another ES, it sends the packet to one of the ISs on its directly attached network. The router then looks up the destination address and forwards the packet along the best route. If the destination ES is on the same subnetwork, the local IS knows this from listening to ESHs and forwards the packet appropriately. The IS also might provide a redirect (RD) message back to the source to tell it that a more direct route is available. If the destination address is an ES on another subnetwork in the same area, the IS knows the correct route and forwards the packet appropriately. If the destination address is an ES in another area, the Level 1 IS sends the packet to the nearest Level 2 IS. Forwarding through Level 2 ISs continues until the packet reaches a Level 2 IS in the destination area. Within the destination area, the ISs forward the packet along the best path until the destination ES is reached.

Link-state update messages help ISs learn about the network topology. Each IS generates an update specifying the ESs and ISs to which it is connected, as well as the associated metrics. The update is then sent to all neighboring ISs, which forward (flood) it to their neighbors, and so on. (Sequence numbers terminate the flood and distinguish old updates from new ones.) Using these updates, each IS can build a complete topology of the network. When the topology changes, new updates are sent.

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation is reasonably efficient. Three optional IS-IS metrics (costs)—delay, expense, and error—are not supported by the ONS 15600. IS-IS maintains a mapping of the metrics to the quality of service (QoS) option in the CLNP packet header. IS-IS uses the mappings to compute routes through the internetwork.

9.7.5 TARP

TARP is used when TL1 target identifiers (TIDs) must be translated to NSAP addresses. The TID-to-NSAP translation occurs by mapping TIDs to the NETs, then deriving NSAPs from the NETs by using the NSAP selector values (Table 9-9 on page 9-29).

TARP uses a selective PDU propagation methodology in conjunction with a distributed database (that resides within the NEs) of TID-to-NET mappings. TARP allows NEs to translate between TID and NET by automatically exchanging mapping information with other NEs. The TARP PDU is carried by the standard CLNP Data PDU. TARP PDU fields are shown in Table 9-10.

Table 9-10 TARP PDU Fields

Field	Abbreviation	Size (bytes)	Description
TARP Lifetime	tar-lif	2	The TARP time-to-live in hops.
TARP Sequence Number	tar-seq	2	The TARP sequence number used for loop detection.
Protocol Address Type	tar-pro	1	Used to identify the type of protocol address that the TID must be mapped to. The value FE is used to identify the CLNP address type.
TARP Type Code	tar-tcd	1	The TARP Type Code identifies the TARP type of PDU. Five TARP types, shown in Table 9-11, are defined.
TID Target Length	tar-tln	1	The number of octets that are in the tar-ttg field.

Table 9-10 TARP PDU Fields (continued)

Field	Abbreviation	Size (bytes)	Description
TID Originator Length	tar-oln	1	The number of octets that are in the tar-tor field.
Protocol Address Length	tar-pln	1	The number of octets that are in the tar-por field.
TID of Target	tar-ttg	$n = 0, 1, 2...$	TID value for the target NE.
TID of Originator	tar-tor	$n = 0, 1, 2...$	TID value of the TARP PDU originator.
Protocol Address of Originator	tar-por	$n = 0, 1, 2...$	Protocol address (for the protocol type identified in the tar-pro field) of the TARP PDU originator. When the tar-pro field is set to FE (hex), tar-por will contain a CLNP address (that is, the NET).

Table 9-11 shows the TARP PDUs types that govern TARP interaction and routing.

Table 9-11 TARP PDU Types

Type	Description	Procedure
1	Sent when a device has a TID for which it has no matching NSAP.	After an NE originates a TARP Type 1 PDU, the PDU is sent to all adjacencies within the NE's routing area.
2	Sent when a device has a TID for which it has no matching NSAP and no response was received from the Type 1 PDU.	After an NE originates a TARP Type 2 PDU, the PDU is sent to all Level 1 and Level 2 neighbors.
3	Sent as a response to Type 1, Type 2, or Type 5 PDUs.	After a TARP Request (Type 1 or 2) PDU is received, a TARP Type 3 PDU is sent to the request originator. Type 3 PDUs do not use the TARP propagation procedures.
4	Sent as a notification when a change occurs locally, for example, a TID or NSAP change. It might also be sent when an NE initializes.	A Type 4 PDU is a notification of a TID or Protocol Address change at the NE that originates the notification. The PDU is sent to all adjacencies inside and outside the NE's routing area.
5	Sent when a device needs a TID that corresponds to a specific NSAP.	When a Type 5 PDU is sent, the CLNP destination address is known, so the PDU is sent to only that address. Type 5 PDUs do not use the TARP propagation procedures.

9.7.5.1 TARP Processing

A TARP data cache (TDC) is created at each NE to facilitate TARP processing. In CTC, the TDC is displayed and managed on the node view Maintenance > OSI > TDC subtab. The TDC subtab contains the following TARP PDU fields:

- TID—TID of the originating NE (tar-tor).
- NSAP—NSAP of the originating NE.

- **Type**— Indicates whether the TARP PDU was created through the TARP propagation process (dynamic) or manually created (static).

Provisionable timers, shown in [Table 9-12](#), control TARP processing.

Table 9-12 TARP Timers

Timer	Description	Default (seconds)	Range (seconds)
T1	Waiting for response to TARP Type 1 Request PDU	15	0–3600
T2	Waiting for response to TARP Type 2 Request PDU	25	0–3600
T3	Waiting for response to address resolution request	40	0–3600
T4	Timer starts when T2 expires (used during error recovery)	20	0–3600

[Table 9-13](#) shows the main TARP processes and the general sequence of events that occurs in each process.

Table 9-13 TARP Processing Flow

Process	General TARP Flow
Find a NET that matches a TID	<ol style="list-style-type: none"> 1. TARP checks its TDC for a match. If a match is found, TARP returns the result to the requesting application. 2. If no match is found, a TARP Type 1 PDU is generated and Timer T1 is started. 3. If Timer T1 expires before a match is found, a Type 2 PDU is generated and Timer T2 is started. 4. If Timer T2 expires before a match is found, Timer T4 is started. 5. If Timer T4 expires before a match is found, a Type 2 PDU is generated and Timer T2 is started.
Find a TID that matches a NET	A Type 5 PDU is generated. Timer T3 is used. However, if the timer expires, no error recovery procedure occurs, and a status message is provided to indicate that the TID cannot be found.
Send a notification of TID or protocol address change	TARP generates a Type 4 PDU in which the tar-ttg field contains the NE's TID value that existed prior to the change of TID or protocol address. Confirmation that other NEs successfully received the address change is not sent.

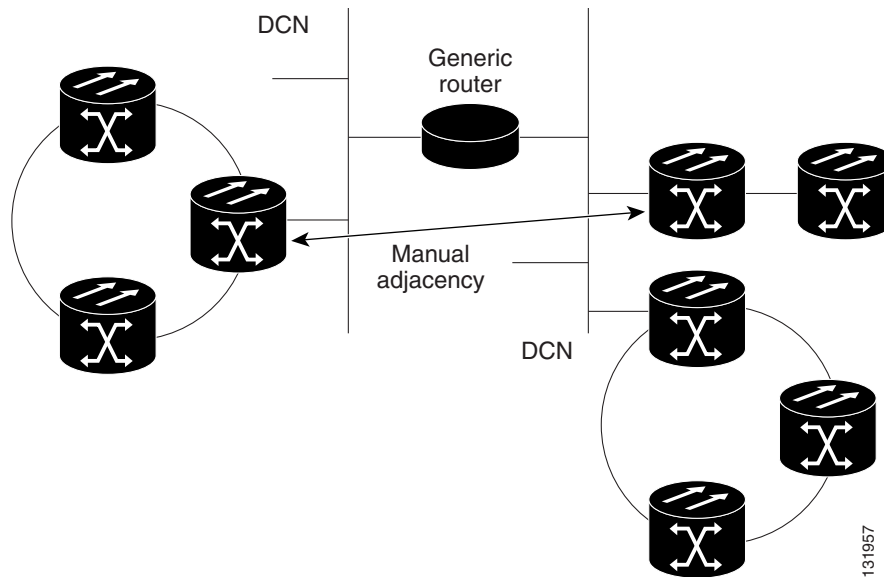
9.7.5.2 TARP Loop Detection Buffer

The TARP loop detection buffer (LDB) can be enabled to prevent duplicate TARP PDUs from entering the TDC. When a TARP Type 1, 2, or 4 PDU arrives, TARP checks its LDB for a NET address (tar-por) of the PDU originator match. If no match is found, TARP processes the PDU and assigns a tar-por, tar-seq (sequence) entry for the PDU to the LDB. If the tar-seq is zero, a timer associated with the LDB entry is started using the provisionable LDB entry timer on the node view OSI > TARP > Config tab. If a match exists, the tar-seq is compared to the LDB entry. If the tar-seq is not zero and is less than or equal to the LDB entry, the PDU is discarded. If the tar-seq is greater than the LDB entry, the PDU is processed and the tar-seq field in the LDB entry is updated with the new value. The ONS 15600 LDB holds approximately 500 entries. The LDB is flushed periodically based on the time set in the LDB Flush timer on the node view OSI > TARP > Config tab.

9.7.5.3 Manual TARP Adjacencies

TARP adjacencies can be manually provisioned in networks where ONS 15600s must communicate across routers or non-SONET NEs that lack TARP capability. In CTC, manual TARP adjacencies are provisioned on the node view Provisioning > OSI > TARP > MAT (Manual Area Table) subtab. The manual adjacency causes a TARP request to hop through the general router or non-SONET NE, as shown in Figure 9-22.

Figure 9-22 Manual TARP Adjacencies



9.7.5.4 Manual TID to NSAP Provisioning


TIDs can be manually linked to NSAPs and added to the TDC. Static TDC entries are similar to static routes. For a specific TID, you force a specific NSAP. Resolution requests for that TID always return that NSAP. No TARP network propagation or instantaneous replies are involved. Static entries allow you to forward TL1 commands to NEs that do not support TARP. However, static TDC entries are not dynamically updated, so outdated entries are not removed after the TID or the NSAP changes on the target node.

9.7.6 TCP/IP and OSI Mediation

Two mediation processes facilitate TL1 networking and file transfers between NEs and ONS client computers running TCP/IP and OSI protocol suites:


- T-TD—Performs a TL1-over-IP to TL1-over-OSI gateway mediation to enable an IP-based OSS to manage OSI-only NEs subtended from a GNE. Figure 9-23 shows the T-TD protocol flow.

Figure 9-23 T-TD Protocol Flow



- FT-TD—Performs an FTP conversion between FTAM and FTP. The FT-TD gateway entity includes an FTAM responder (server) and an FTP client, allowing FTAM initiators (clients) to store, retrieve, or delete files from an FTP server. The FT-TD gateway is unidirectional and is driven by the FTAM initiator. The FT-TD FTAM responder exchanges messages with the FTAM initiator over the full OSI stack. [Figure 9-24](#) shows the FT-TD protocol flow.

Figure 9-24 FT-TD Protocol Flow



The ONS 15600 uses FT-TD for the following file transfer processes:

- Software downloads
- Database backups and restores

9.7.7 OSI Virtual Routers

The ONS 15600 supports twelve OSI virtual routers. The routers are provisioned on the Provisioning > OSI > Routers tab. Each router has an editable manual area address and a unique NSAP System ID that is set to the node MAC address + *n*. For Router 1, *n* = 0. For Router 2, *n* = 1. For Router 3, *n* = 2, and for Router 12, *n* = 11. Each router can be enabled and connected to different OSI routing areas. However, Router 1 is the primary router, and it must be enabled before Routers 2 through 12 can be enabled. The Router 1 manual area address and System ID create the NSAP address assigned to the node’s TID. In addition, Router 1 supports OSI TARP, mediation, and tunneling functions that are not supported by Routers 2 through 12. These include:

- TID-to-NSAP resolution
- TARP data cache
- IP-over-CLNS tunnels
- FTAM

- FT-TD
- T-TD
- LAN subnet

OSI virtual router constraints depend on the routing mode provisioned for the node. [Table 9-14](#) shows the number of IS L1s, IS L1/L2s, and DCCs that are supported by each router. An IS L1 and IS L1/L2 support one ES per DCC subnet and up to 100 ESs per LAN subnet.

Table 9-14 *OSI Virtual Router Constraints*

Routing Mode	Router 1	Router 2 — 12	IS L1 per area	IS L1/L2 per area	DCC per IS
IS L1	Yes	Yes	250	—	60
IS L1/L2	Yes	Yes	250	50	60

Each OSI virtual router has a primary manual area address. You can also create two additional manual area addresses. These manual area addresses can be used to:

- Split up an area—Nodes within a given area can accumulate to a point that they are difficult to manage, cause excessive traffic, or threaten to exceed the usable address space for an area. Additional manual area addresses can be assigned so that you can smoothly partition a network into separate areas without disrupting service.
- Merge areas—Use transitional area addresses to merge as many as three separate areas into a single area that shares a common area address.
- Change to a different address—You might need to change an area address for a particular group of nodes. Use multiple manual area addresses to allow incoming traffic intended for an old area address to continue being routed to associated nodes.

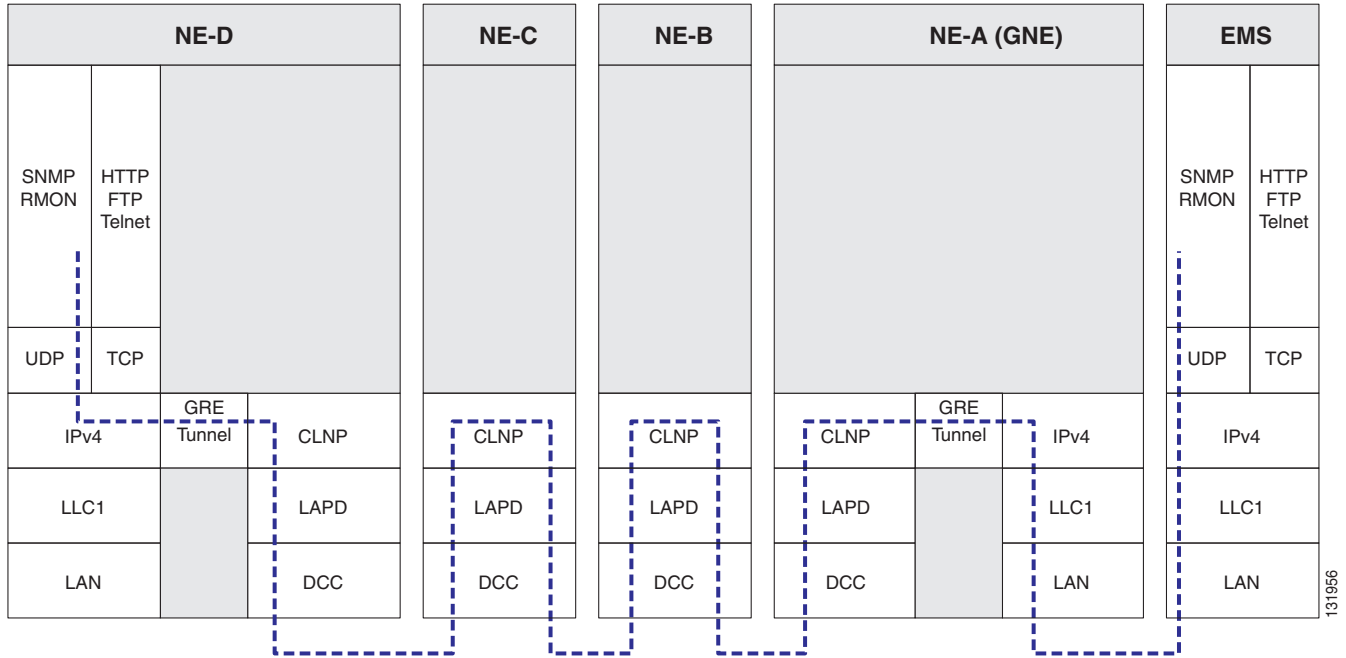
9.7.8 IP-over-CLNS Tunnels

IP-over-CLNS tunnels are used to encapsulate IP for transport across OSI NEs. The ONS 15600 supports two tunnel types:

- GRE—Generic Routing Encapsulation is a tunneling protocol that encapsulates one network layer for transport across another. GRE tunnels add both a CLNS header and a GRE header to the tunnel frames. GRE tunnels are supported by Cisco routers and some other vendor NEs.
- Cisco IP—The Cisco IP tunnel directly encapsulates the IP packet with no intermediate header. Cisco IP is supported by most Cisco routers.

[Figure 9-25](#) shows the protocol flow when an IP-over-CLNS tunnel is created through four NEs (A, B, C, and D). The tunnel ends are configured on NEs A and D, which support both IP and OSI. NEs B and C only support OSI, so they only route the OSI packets.

Figure 9-25 IP-over-CLNS Tunnel Flow



9.7.8.1 Provisioning IP-over-CLNS Tunnels

IP-over-CLNS tunnels must be carefully planned to prevent nodes from losing visibility or connectivity. Before you begin a tunnel, verify that the tunnel type, either Cisco IP or GRE, is supported by the equipment at the other end. Always verify IP and NSAP addresses. Provisioning of IP-over-CLNS tunnels in CTC is performed on the node view Provisioning > OSI > IP over CLNS Tunnels tab. For procedures, refer to the “Turn Up Node” chapter in the *ONS 15600 Procedures Guide*.

Provisioning IP-over-CLNS tunnels on Cisco routers requires the following prerequisite tasks, as well as other OSI provisioning:

- (Required) Enable IS-IS
- (Optional) Enable routing for an area on an interface
- (Optional) Assign multiple area addresses
- (Optional) Configure IS-IS interface parameters
- (Optional) Configure miscellaneous IS-IS parameters

The Cisco IOS commands used to create IP-over-CLNS tunnels (CTunnels) are shown in [Table 9-15](#).

Table 9-15 IP Over CLNS Tunnel IOS Commands

Step	Step	Purpose
1	Router (config) # interface ctunnel <i>interface-number</i>	Creates a virtual interface to transport IP over a CLNS tunnel and enters interface configuration mode. The interface number must be unique for each CTunnel interface.
2	Router (config-if) # ctunnel destination <i>remote-nsap-address</i>	Configures the destination parameter for the CTunnel. Specifies the destination NSAPI address of the CTunnel, where the IP packets are extracted.
3	Router (config-if) # ip address <i>ip-address mask</i>	Sets the primary or secondary IP address for an interface.

If you are provisioning an IP-over-CLNS tunnel on a Cisco router, always follow procedures provided in the Cisco IOS documentation for the router you are provisioning. For information about ISO CLNS provisioning including IP-over-CLNS tunnels, see the “Configuring ISO CLNS” chapter in the *Cisco IOS Apollo Domain, Banyon VINES, DECnet, ISO CLNS, and XNS Configuration Guide*.

9.7.8.2 IP-Over-CLNS Tunnel Scenario 1: ONS Node to Other Vendor GNE

Figure 9-26 shows an IP-over-CLNS tunnel created from an ONS node to another vendor GNE. The other vendor NE has an IP connection to an IP DCN to which a CTC computer is attached. An OSI-only (LAP-D) SDCC and a GRE tunnel are created between the ONS NE 1 to the other vendor GNE.

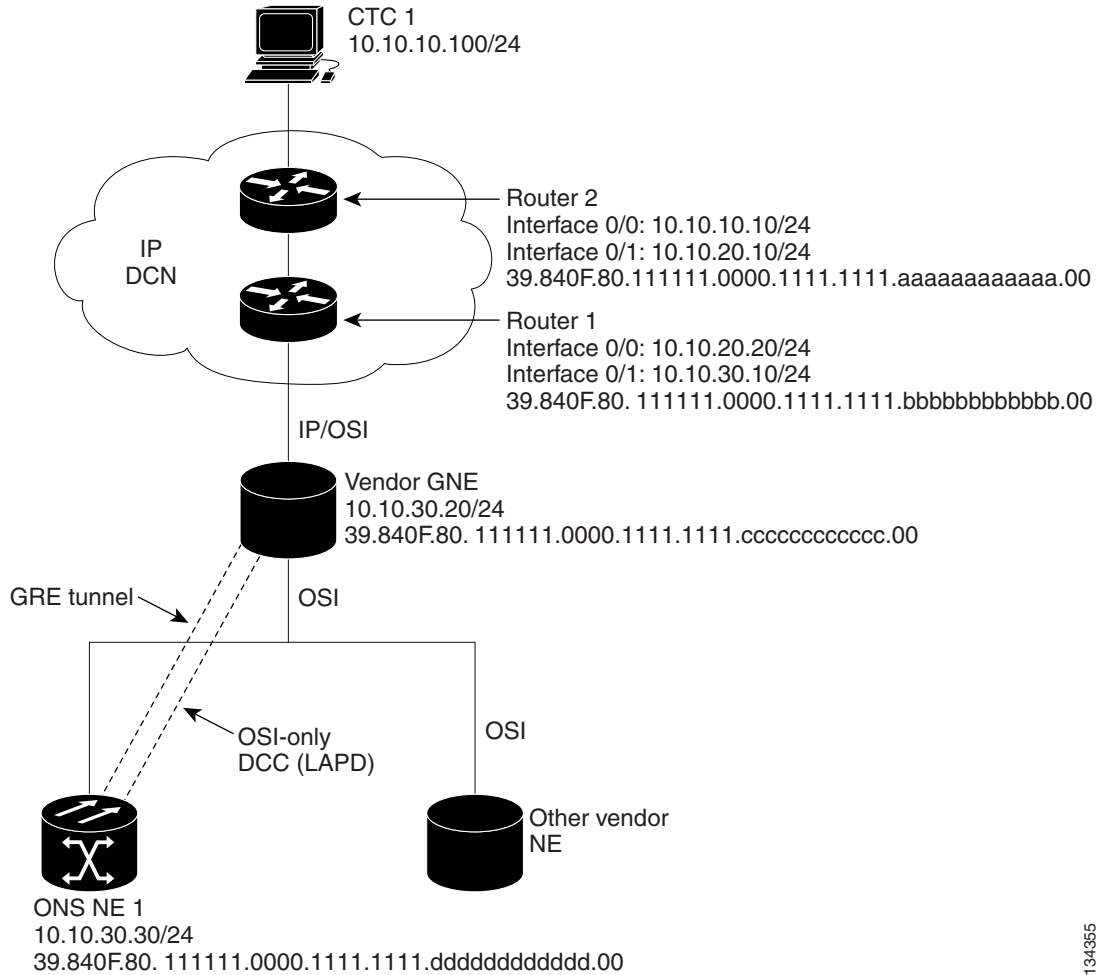
IP-over-CLNS tunnel on the ONS NE 1:

- Destination: 10.10.10.100 (CTC 1)
- Mask: 255.255.255.255 for host route (CTC 1 only), or 255.255.255.0 for subnet route (all CTC computers residing on the 10.10.10.0 subnet)
- NSAP: 39.840F.80.1111.0000.1111.1111.cccccccccc.00 (other vendor GNE)
- Metric: 110
- Tunnel Type: GRE

IP-over-CLNS tunnel on the other vendor GNE:

- Destination: 10.20.30.30 (ONS NE 1)
- Mask: 255.255.255.255 for host route (ONS NE 1 only), or 255.255.255.0 for subnet route (all ONS nodes residing on the 10.30.30.0 subnet)
- NSAP: 39.840F.80.1111.0000.1111.1111.dddddddddd.00 (ONS NE 1)
- Metric: 110
- Tunnel Type: GRE

Figure 9-26 IP Over CLNS Tunnel Scenario 1: ONS NE to Other Vendor GNE



134355

9.7.8.3 IP-Over-CLNS Tunnel Scenario 2: ONS Node to Router

Figure 9-27 shows an IP-over-CLNS tunnel from an ONS node to a router. The other vendor NE has an OSI connection to a router on an IP DCN, to which a CTC computer is attached. An OSI-only (LAP-D) SDCC is created between the ONS NE 1 and the other vendor GNE. The OSI over IP tunnel can be either the Cisco IP tunnel or a GRE tunnel, depending on the tunnel types supported by the router.

IP-over-CLNS tunnel on ONS NE 1:

- Destination: 10.10.30.10 (Router 1, Interface 0/1)
- Mask: 255.255.255.255 for host route (Router 1 only), or 255.255.255.0 for subnet route (all routers on the same subnet)
- NSAP: 39.840F.80.1111.0000.1111.1111.bbbbbbbbbbbb.00 (Router 1)
- Metric: 110
- Tunnel Type: Cisco IP

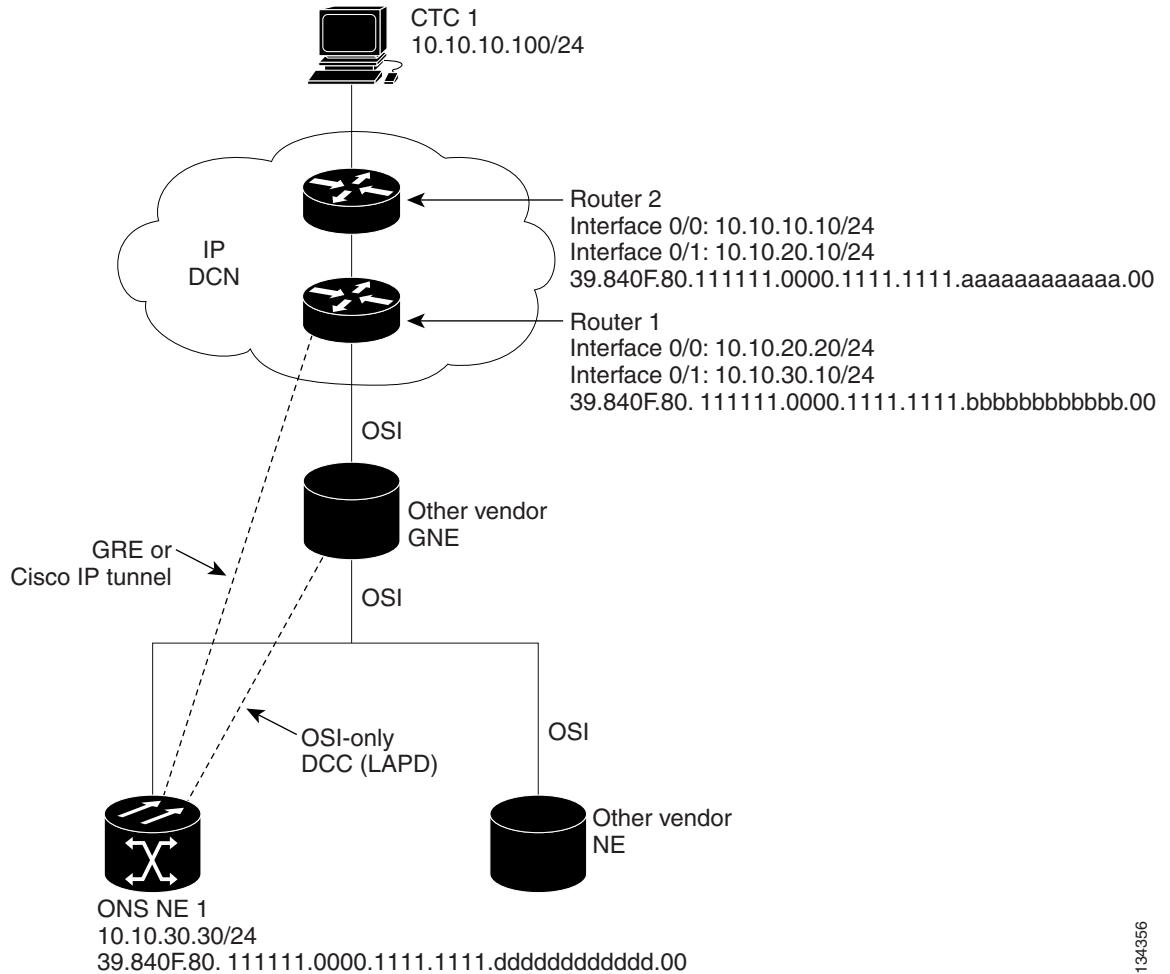
CTunnel (IP over CLNS) on Router 1:

```

ip routing
clns routing
interface ctunnel 102
  ip address 10.10.30.30 255.255.255.0
  ctunnel destination 39.840F.80.1111.0000.1111.1111.ddddddddddd.00
interface Ethernet0/1
  clns router isis
router isis
  net 39.840F.80.1111.0000.1111.1111.bbbbbbbbbbbb.00

```

Figure 9-27 IP Over CLNS Tunnel Scenario 2: ONS Node to Router



134356

9.7.8.4 IP-Over-CLNS Tunnel Scenario 3: ONS Node to Router Across an OSI DCN

Figure 9-28 shows an IP-over-CLNS tunnel from an ONS node to a router across an OSI DCN. The other vendor NE has an OSI connection to an IP DCN to which a CTC computer is attached. An OSI-only (LAP-D) SDCC is created between the ONS NE 1 and the other vendor GNE. The OSI over IP tunnel can be either the Cisco IP tunnel or a GRE tunnel, depending on the tunnel types supported by the router.

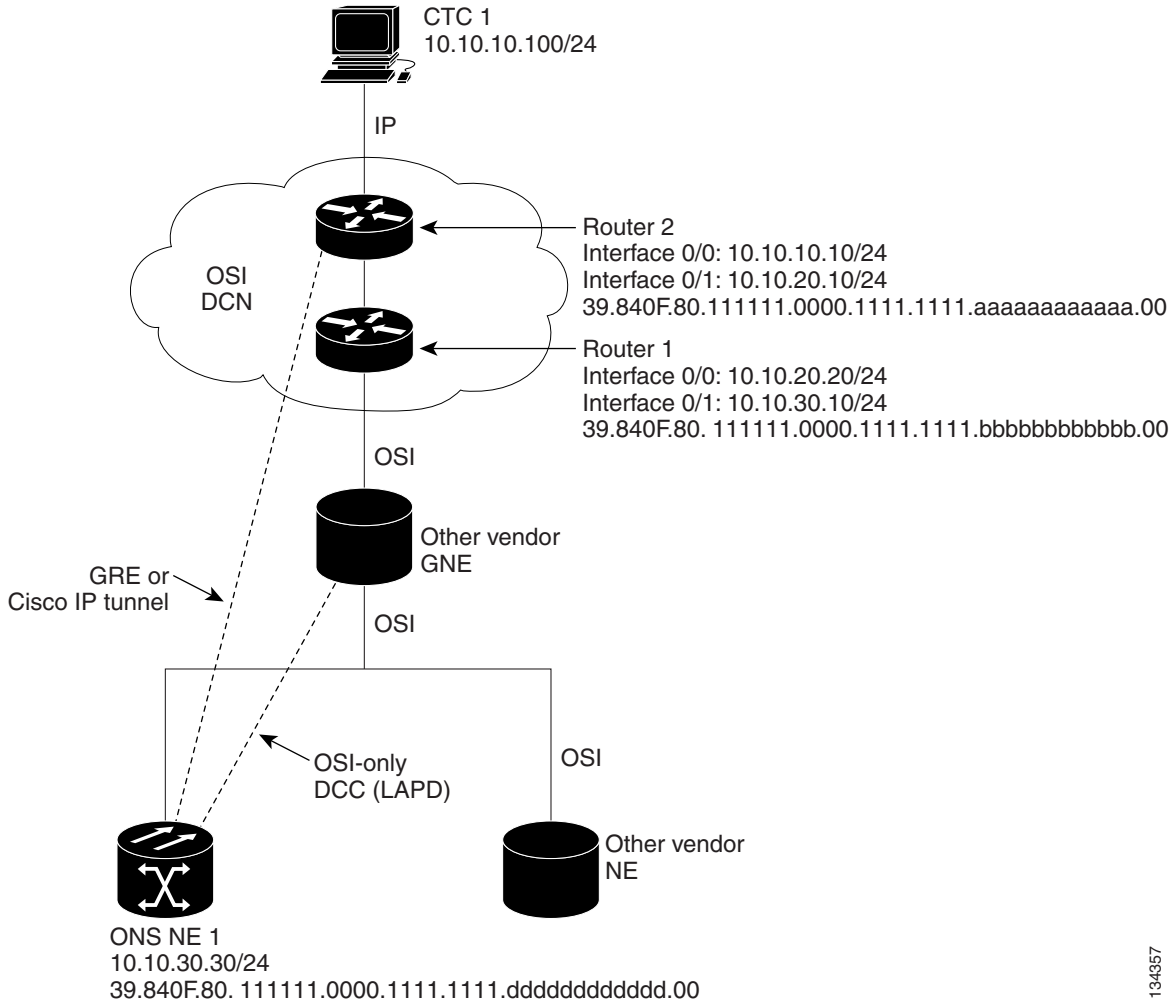
IP-over-CLNS tunnel on ONS NE 1:

- Destination: Router 2 IP address
- Mask: 255.255.255.255 for host route (CTC 1 only), or 255.255.255.0 for subnet route (all CTC computers on the same subnet)
- NSAP: Other vendor GNE NSAP address
- Metric: 110
- Tunnel Type: Cisco IP

IP over OSI tunnel on Router 2 (sample Cisco IOS provisioning):

```
ip routing
clns routing
interface ctunnel 102
    ip address 10.10.30.30 255.255.255.0
    ctunnel destination 39.840F.80.1111.0000.1111.1111.aaaaaaaaaaaa.00
interface Ethernet0/1
    clns router isis
router isis
    net 39.840F.80.1111.0000.1111.1111.aaaaaaaaaaaa.00
```

Figure 9-28 IP-Over-CLNS Tunnel Scenario 3: ONS Node to Router Across an OSI DCN



9.7.9 OSI/IP Networking Scenarios

The following eight scenarios show examples of ONS 15600s in networks with OSI-based NEs. The scenarios show ONS 15600s in a variety of roles. The scenarios assume the following:


- ONS 15600 NEs are configured as dual OSI and IP nodes with both IP and NSAP addresses. They run both OSPF and OSI (IS-IS or ES-IS) routing protocols as “Ships-In-The-Night,” with no route redistribution.
- ONS 15600 NEs run TARP, which allows them to resolve a TL1 TID to a NSAP address. A TID might resolve to both an IP and an NSAP address when the destination TID is an ONS 15600 NE that has both IP and NSAP address.
- DCC links between ONS 15600 NEs and OSI-only NEs run the full OSI stack over LAP-D, which includes IS-IS, ES-IS, and TARP.
- DCC links between ONS 15600 NEs run the full OSI stack and IP (OSPF) over PPP.

- All ONS 15600 NEs participating in an OSI network run OSI over PPP between themselves. This is needed so that other vendor GNEs can route TL1 commands to all ONS 15600 NEs participating in the OSI network.

9.7.9.1 OSI/IP Scenario 1: IP OSS, IP DCN, ONS GNE, IP DCC, and ONS ENE

Figure 9-29 shows OSI/IP Scenario 1, the current ONS 15600 IP-based implementation, with an IP DCN, IP-over-PPP DCC, and OSPF routing.

Figure 9-29 OSI/IP Scenario 1: IP OSS, IP DCN, ONS GNE, IP DCC, and ONS ENE




1	IP OSS manages ONS 15600 using TL1 and FTP.
2	DCCs carry IP over the PPP protocol.
3	The ONS 15600 network is managed by IP over OSPF.

9.7.9.2 OSI/IP Scenario 2: IP OSS, IP DCN, ONS GNE, OSI DCC, and Other Vendor ENE

OSI/IP Scenario 2 (Figure 9-30) shows an ONS 15600 GNE in a multivendor OSI network. Both the ONS 15600 GNE and the other vendor NEs are managed by an IP OSS using TL1 and FTP. The ONS 15600 is also managed by CTC and Cisco Transport Manager (CTM). Because the other vendor NE only supports TL1 and FTAM over the full OSI stack, the ONS 15600 GNE provides T-TD and FT-TD mediation to convert TL1/IP to TL1/OSI and FTAM/OSI to FTP/IP.

Figure 9-30 OSI/IP Scenario 2: IP OSS, IP DCN, ONS GNE, OSI DCC, and Other Vendor ENE



1	The IP OSS manages ONS 15600 and other vendor NEs using TL1 and FTP.
2	The ONS 15600 GNE performs mediation for other vendor NEs.
3	DCCs between the ONS 15600 GNE and ONS 15600 NEs are provisioned for IP and OSI over PPP.
4	DCCs between the ONS 15600 GNE and other vendor NEs are provisioned for OSI over LAP-D.
5	The ONS 15600 and the other vendor NE network include IP over OSPF and OSI over the IS-IS protocol.

The ONS 15600 GNE routes TL1 traffic to the correct NE by resolving the TL1 TID to either an IP or NSAP address. For TL1 traffic to other vendor NEs (OSI-only nodes), the TID is resolved to an NSAP address. The ONS 15600 GNE passes the TL1 to the mediation function, which encapsulates it over the full OSI stack and routes it to the destination using the IS-IS protocol.

For TL1 traffic to ONS 15600 NEs, the TID is resolved to both an IP and an NSAP address. The ONS 15600 GNE follows the current TL1 processing model and forwards the request to the destination NE using the TCP/IP stack and OSPF routing.

OSS-initiated software downloads consist of two parts: the OSS to destination NE TL1 download request and the file transfer. The TL1 request is handled the same as described earlier. The ONS 15600 NEs use FTP for file transfers. OSI-only NEs use FTAM to perform file transfers. The FTAM protocol is carried over OSI between the OSI NE and the ONS 15600 GNE. The GNE mediation translates between FTAM to FTP.

9.7.9.3 OSI/IP Scenario 3: IP OSS, IP DCN, Other Vendor GNE, OSI DCC, and ONS ENE


In OSI/IP Scenario 3 (Figure 9-31), all TL1 traffic between the OSS and GNE is exchanged over the IP DCN. TL1 traffic targeted for the GNE is processed locally. All other TL1 traffic is forwarded to the OSI stack, which performs IP-to-OSI TL1 translation. The TL1 is encapsulated in the full OSI stack and sent to the target NE over the DCC. The GNE can route to any node within the IS-IS domain because all NEs, ONS 15600 and non-ONS 15600, have NSAP addresses and support IS-IS routing.

TL1 traffic received by an ONS 15600 NE and not addressed to its NSAP address is forwarded by IS-IS routing to the correct destination. TL1 traffic received by an ONS 15600 NE and addressed to its NSAP is sent up the OSI stack to the mediation function, which extracts the TL1 and passes it to the ONS 15600 TL1 processor.

An OSS initiated software download includes the OSS to destination node TL1 download request and the file transfer. The TL1 request is handled as described earlier. The target node uses FTAM for file transfers because the GNE does not support IP on the DCC and cannot forward FTP. The ONS 15600 NEs therefore must support an FTAM client and initiate file transfer using FTAM when subtended to an OSI GNE.

In this scenario, the GNE has both IP and OSI DCN connections. The GNE only supports TL1 and FTP over IP. Both are translated and then carried over OSI to the destination ENE (ONS 15600 or OSI-only NE). All other IP traffic is discarded by the GNE. The CTC/CTM IP traffic is carried over an IP-over-OSI tunnel to an ONS 15600 NE. The tunnel is created between an external router and an ONS 15600 NE. The traffic is sent to the ONS 15600 terminating the tunnel. That ONS 15600 then forwards the traffic over the tunnel to CTC/CTM by way of the external router.


Figure 9-31 *OSI/IP Scenario 3: IP OSS, IP DCN, Other Vendor GNE, OSI DCC, and ONS ENE*



1	The IP OSS manages the ONS 15600 and other vendor NEs using TL1 and FTP.
2	The other vendor GNE performs mediation for TL1 and FTP, so the DCCs to the ONS 15600 and other vendor NEs are OSI-only.
3	CTC/CTM communicates with ONS 15600 NEs over a IP-over-CLNS tunnel. The tunnel is created from the ONS 15600 node to the external router.
4	The ONS 15600 NE exchanges TL1 over the full OSI stack using FTAM for file transfer.

Figure 9-32 shows the same scenario, except the IP-over-CLNS tunnel endpoint is the GNE and not the DCN router.

Figure 9-32 OSI/IP Scenario 3 with OSI/IP-over-CLNS Tunnel Endpoint at the GNE




1	The IP OSS manages ONS and other vendor NEs using TL1 and FTP.
2	The router routes requests to the other vendor GNE.
3	The other vendor GNE performs mediation for TL1 and FTP, so the DCCs to ONS 15600 and other vendor NEs are OSI-only.
4	CTC/CTM communicates with ONS 15600 NEs over an IP-over-CLNS tunnel between the ONS 15600 and the GNE.
5	ONS 15600 NEs exchange TL1 over the full OSI stack. FTAM is used for file transfer.

9.7.9.4 OSI/IP Scenario 4: Multiple ONS DCC Areas

OSI/IP Scenario 4 (Figure 9-33) is similar to OSI/IP Scenario 3 except that the OSI GNE is subtended by multiple isolated ONS 15600 areas. A separate IP-over-CLNS tunnel is required to each isolated ONS 15600 OSPF area. An alternate approach is to create a single IP-over-CLNS tunnel from CTC/CTM to an ONS 15600 NE, and then to configure a tunnel from that NE to an NE in each isolated OSPF area. This approach requires additional static routes.

Figure 9-33 OSI/IP Scenario 4: Multiple ONS DCC Areas




1	The IP OSS manages ONS 15600 and other vendor NEs using TL1 and FTP.
2	A separate tunnel is created for each isolated ONS 15600 DCC area.

9.7.9.5 OSI/IP Scenario 5: GNE Without an OSI DCC Connection

OSI/IP Scenario 5 (Figure 9-34) is similar to OSI/IP Scenario 3 except that the OSI GNE only has an IP connection to the DCN. It does not have an OSI DCN connection to carry CTC/CTM IP traffic through an IP-over-OSI tunnel. A separate DCN to ONS 15600 NE connection is created to provide CTC/CTM access.

Figure 9-34 OSI/IP Scenario 5: GNE Without an OSI DCC Connection




1	The IP OSS manages ONS 15600 and other vendor NEs using TL1 and FTP.
2	The other vendor GNE performs mediation on TL1 and FTP, so DCCs are OSI-only.
3	CTC/CTM communicates with ONS 15600 NEs over a separate IP DCN connection.
4	ONS 15600 NE exchanges TL1 over the full OSI stack. FTAM is used for file transfers.

9.7.9.6 OSI/IP Scenario 6: IP OSS, OSI DCN, ONS GNE, OSI DCC, and Other Vendor ENE

OSI/IP Scenario 6 (Figure 9-35) shows how the ONS 15600 supports OSI DCNs. The OSI DCN has no impact on the ONS 15600 because all IP traffic (CTC/CTM, FTP, and TL1) is tunneled through the OSI DCN.

Figure 9-35 OSI/IP Scenario 6: IP OSS, OSI DCN, ONS GNE, OSI DCC, and Other Vendor ENE

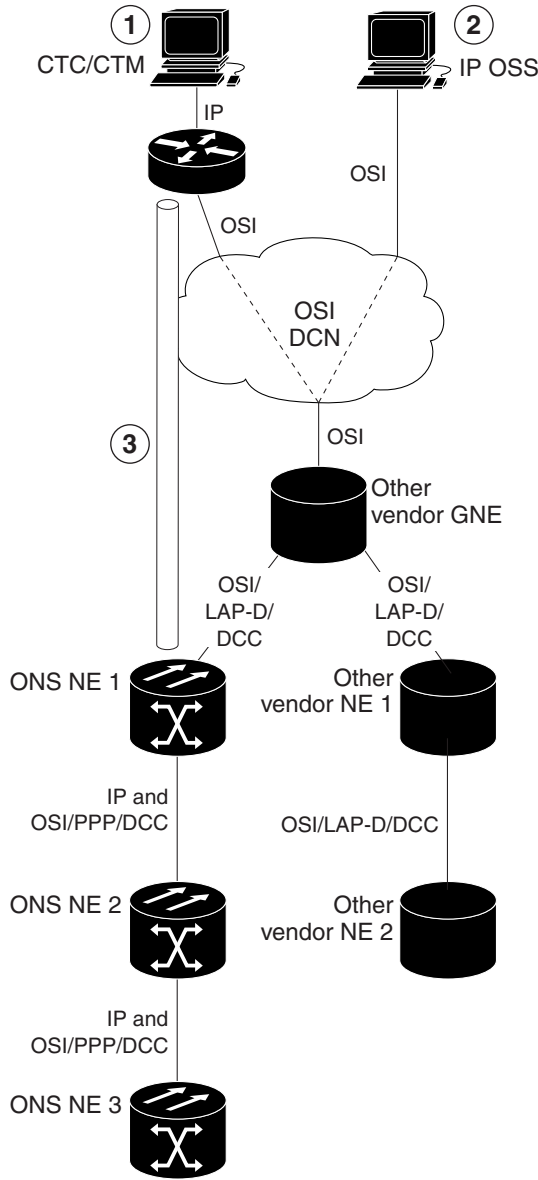


1	The IP OSS manages ONS 15600 and other vendor NEs using TL1 and FTP.
2	OSS IP traffic is tunneled through the DCN to the ONS 15600 GNE.
3	CTC/CTM IP traffic is tunneled through the DCN to the ONS 15600 GNE.
4	The GNE performs mediation for other vendor NEs.

9.7.9.7 OSI/IP Scenario 7: OSI OSS, OSI DCN, Other Vendor GNE, OSI DCC, and ONS NEs

OSI/IP Scenario 7 (Figure 9-36) shows an example of a European network.

Figure 9-36 OSI/IP Scenario 7: OSI OSS, OSI DCN, Other Vendor GNE, OSI DCC, and ONS NEs



131937

1	ONS 15600 NEs are managed by CTC/CTM only (TL1/FTP is not used).
2	The OSI OSS manages other vendor NEs only.
3	CTC/CTM communicates with the ONS 15600 over a IP-over-CLNS tunnel between the ONS 15600 NE and external router.

In European networks:

- CTC and CTM are used for management only.
- IP-over-CLNS tunnels are widely accepted and deployed.

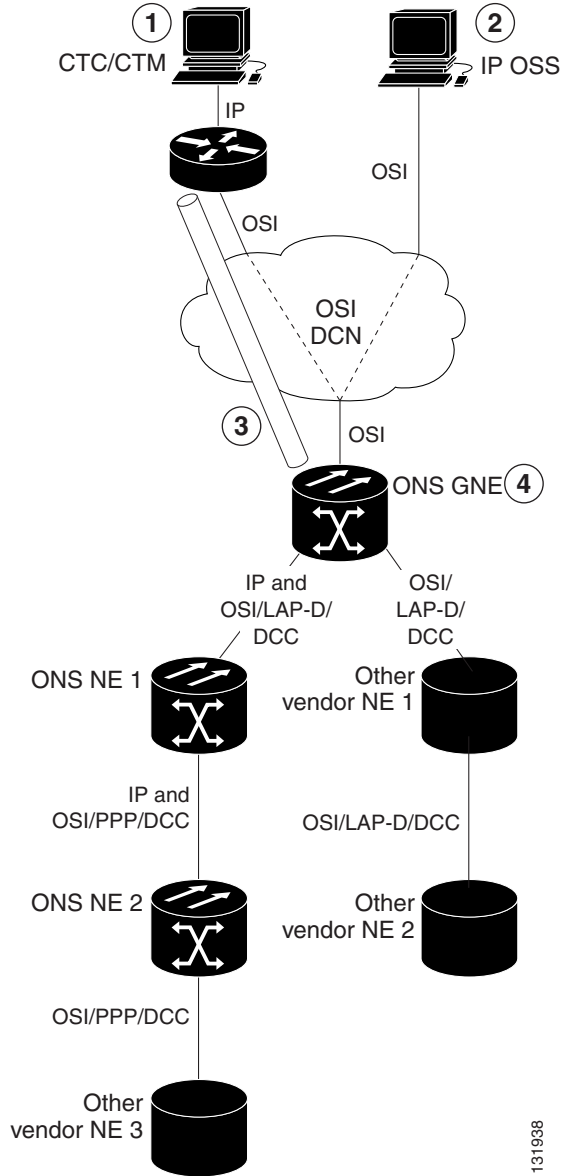
- TL1 management is not required.
- FTP file transfer is not required.
- TL1 and FTAM to FTP mediation is not required.

Management traffic between CTC/CTM and ONS 15600 NEs is carried over an IP-over-CLNS tunnel. A static route is configured on the ONS 15600 that terminates the tunnel (ONS 15600 NE 1) so that downstream ONS 15600 NEs (ONS 15600 NE 2 and 3) know how to reach CTC/CTM.

9.7.9.8 OSI/IP Scenario 8: OSI OSS, OSI DCN, ONS GNE, OSI DCC, and Other Vendor NEs

OSI/IP Scenario 8 ([Figure 9-37](#)) is another example of a European network. Similar to OSI/IP Scenario 7, the ONS 15600 NEs are solely managed by CTC/CTM. The CTC/CTM IP traffic is carried over a IP-over-OSI tunnel between an external router and the ONS 15600 GNE. The GNE extracts the IP from the tunnel and forwards it to the destination ONS 15600. Management traffic between the OSS and other vendor NEs is routed by the ONS 15600 GNE and NEs. This is possible because all ONS 15600 NEs run dual stacks (OSI and IP).

Figure 9-37 OSI/IP Scenario 8: OSI OSS, OSI DCN, ONS GNE, OSI DCC, and Other Vender NEs



1	The ONS NEs are managed by CTC/CTM only (TL1/FTP is not used).
2	The OSI OSS manages other vendor NEs only.
3	CTC/CTM communicates with the ONS 15600 over an IP-over-CLNS tunnel between the ONS 15600 NE and the external router. A static route is needed on the GNE.
4	The ONS 15600 GNE routes OSI traffic to other vendor NEs. No IP-over-CLNS tunnel is needed.

9.7.10 OSI Provisioning in CTC

Table 9-16 shows the OSI actions that are performed from the node view Provisioning tab. Refer to the *Cisco ONS 15600 Procedures Guide* for OSI procedures and tasks.

Table 9-16 OSI Actions from the CTC Provisioning Tab

Tab	Actions
OSI > Main Setup	<ul style="list-style-type: none"> View and edit Primary Area Address. Change OSI routing mode. Change LSP buffers.
OSI > TARP > Config	Configure the TARP parameters: <ul style="list-style-type: none"> PDU L1/L2 propagation and origination. TARP data cache and loop detection buffer. LAN storm suppression. Type 4 PDU on startup. TARP timers: LDB, T1, T2, T3, T4.
OSI > TARP > Static TDC	Add and delete static TARP data cache entries.
OSI > TARP > MAT	Add and delete static manual area table entries.
OSI > Routers > Setup	<ul style="list-style-type: none"> Enable and disable routers. Add, delete, and edit manual area addresses.
OSI > Routers > Subnets	Edit SDCC, LDCC, and LAN subnets that are provisioned for OSI.
OSI > Tunnels	Add, delete, and edit Cisco and IP-over-CLNS tunnels.
Comm Channels > SDCC	<ul style="list-style-type: none"> Add OSI configuration to an SDCC. Choose the data link layer protocol, PPP or LAP-D.
Comm Channels > LDCC	<ul style="list-style-type: none"> Add OSI configuration to an SDCC.

Table 9-17 shows the OSI actions that are performed from the node view Maintenance tab.

Table 9-17 OSI Actions from the CTC Maintenance Tab

Tab	Actions
OSI > ISIS RIB	View the IS-IS routing table.
OSI > ESIS RIB	View ESs that are attached to ISs.
OSI > TDC	<ul style="list-style-type: none"> View the TARP data cache and identify static and dynamic entries. Perform TID to NSAP resolutions. Flush the TDC.