CHAPTER 8

Management Network Connectivity

This chapter provides an overview of Cisco ONS 15310-MA SDH data communications network (DCN) connectivity. Cisco Optical Networking System (ONS) network communication is based on IP, including communication between Cisco Transport Controller (CTC) computers and ONS 15310-MA SDH nodes, and communication among networked ONS 15310-MA SDH nodes. The chapter provides scenarios showing ONS 15310-MA SDH nodes 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.

Although ONS 15310-MA SDH DCN communication is based on IP, ONS 15310-MA SDH nodes can be networked to equipment that is based on the Open System Interconnection (OSI) protocol suites. This chapter describes the OSI implementation and provides scenarios that show how the ONS 15310-MA SDH can be networked within a mixed IP and OSI environment.

Chapter topics include:

- 8.1 IP Networking Overview, page 8-2
- 8.2 IP Addressing Scenarios, page 8-2
- 8.3 Routing Table, page 8-16
- 8.4 External Firewalls, page 8-18
- 8.5 Open GNE, page 8-20
- 8.6 TCP/IP and OSI Networking, page 8-22
- 8.7 IPv6 Network Compatibility, page 8-40
- 8.8 IPv6 Native Support, page 8-40
- 8.9 FTP Support for ENE Database Backup, page 8-42

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 networking setup instructions, refer to the "Turn Up a Node" chapter of the *Cisco ONS 15310-MA SDH Procedure Guide*.

Note To connect ONS 15310-MA SDH nodes to an IP network, you must work with a LAN administrator or other individual at your site who has IP networking training and experience.

8.1 IP Networking Overview

ONS 15310-MA SDH nodes can be connected in many different ways within an IP environment:

- They can be connected to LANs through direct connections or a router.
- IP subnetting can create ONS 15310-MA SDH login node groups, which allow you to provision non-data communications channel (DCC) connected nodes in a network.
- Different IP functions and protocols can be used to achieve specific network goals. For example, Proxy Address Resolution Protocol (ARP) enables one LAN-connected ONS 15310-MA SDH to serve as a gateway for ONS 15310-MA SDH nodes that are not connected to the LAN.
- You can create static routes to enable connections among multiple Cisco Transport Controller (CTC) sessions with ONS 15310-MA SDH nodes that reside on the same subnet with multiple CTC sessions.
- If ONS 15310-MA SDH nodes are connected to Open Shortest Path First (OSPF) networks, ONS 15310-MA SDH network information is automatically communicated across multiple LANs and WANs.
- The ONS 15310-MA SDH proxy server controls the visibility and accessibility between CTC computers and ONS 15310-MA SDH element nodes.

8.2 IP Addressing Scenarios

ONS 15310-MA SDH IP addressing generally has seven common scenarios or configurations. Use the scenarios as building blocks for more complex network configurations. Table 8-1 provides a general list of items to check when setting up ONS 15310-MA SDH nodes in IP networks.

Table 8-1 General P Troubleshooting Checklist

Item	What to Check
Link integrity	Verify that link integrity exists between: • CTC computer and network hub/switch • ONS 15310-MA SDH nodes (RJ-45 ports labeled LAN) and network hub/switch • Router ports and hub/switch ports
Node hub/switch ports	Verify connectivity. If connectivity problems occur, set the hub or switch port that is connected to the ONS 15310-MA SDH to 10 Mbps half-duplex.
Ping	Ping the node to test connections between computers and ONS 15310-MA SDH nodes.
IP addresses/subnet masks	Verify that ONS 15310-MA SDH IP addresses and subnet masks are set up correctly.
Optical connectivity	Verify that ONS 15310-MA SDH optical trunk ports are in service and that a DCC is enabled on each trunk port.

8.2.1 Scenario 1: CTC and ONS 15310-MA SDH Nodes on the Same Subnet

Scenario 1 shows a basic ONS 15310-MA SDH LAN configuration (Figure 8-1). The ONS 15310-MA SDH nodes and CTC computer reside on the same subnet. All nodes connect to LAN A and have DCC connections.

Figure 8-1 Scenario 1: CTC and ONS 15310-MA SDH Nodes on the Same Subnet

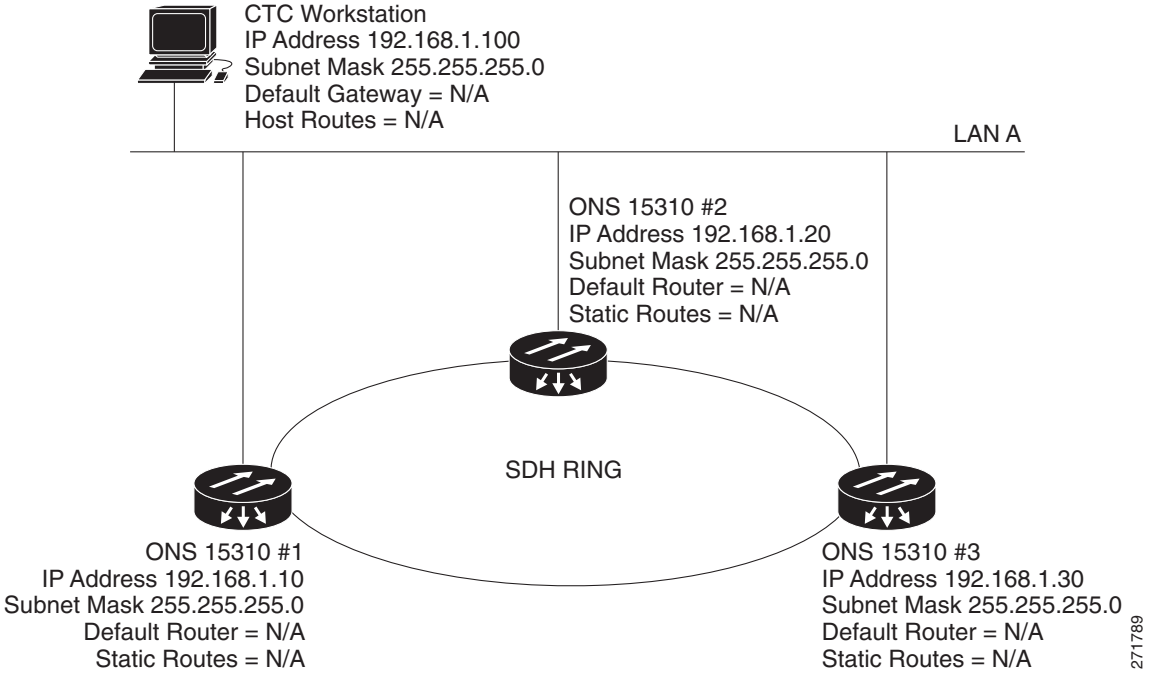

8.2.2 Scenario 2: CTC and ONS 15310-MA SDH Nodes Connected to a Router

In Scenario 2 the CTC computer resides on a subnet (192.168.1.0) and attaches to LAN A (Figure 8-2). The ONS 15310-MA SDH nodes 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 Figure 8-2, a DHCP server is not available.

Figure 8-2 Scenario 2: CTC and ONS 15310-MA SDH Nodes Connected to Router

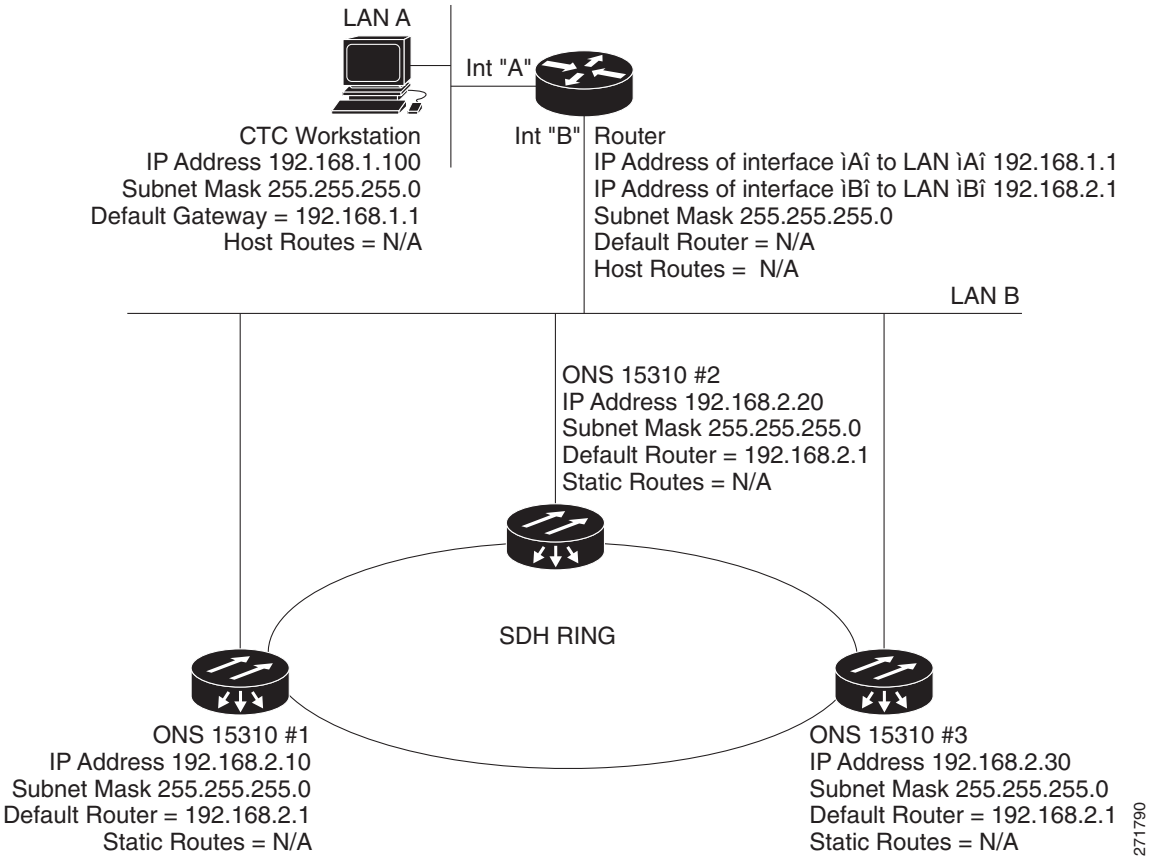

8.2.3 Scenario 3: Using Proxy ARP to Enable an ONS 15310-MA SDH Gateway

ARP matches higher-level IP addresses to the physical addresses of the destination host. It uses a lookup table (called 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 15310-MA SDH to respond to the ARP request for ONS 15310-MA SDH nodes not connected to the LAN. (Proxy ARP requires no user configuration.) For the proxy ARP node to require no user confirmation, the DCC-connected nodes must reside on the same subnet. When a LAN device sends an ARP request to an ONS 15310-MA SDH that is not connected to the LAN, the gateway ONS 15310-MA SDH returns its MAC address to the LAN device. The LAN device then sends the datagram for the remote ONS 15310-MA SDH to the MAC address of the proxy node. The proxy ONS 15310-MA SDH uses its routing table to forward the datagram to the non-LAN ONS 15310-MA SDH.

Scenario 3 is similar to Scenario 1, but only one ONS 15310-MA SDH node (#1) connects to the LAN (Figure 8-3). Two ONS 15310-MA SDH nodes (#2 and #3) connect to Node 1 through the SDH DCC. Because all three nodes are on the same subnet, Proxy ARP enables Node 1 to serve as a gateway for Nodes 2 and 3.

Note This scenario assumes all CTC connections are to Node 1. If you connect a laptop to either Node 2 or Node 3, network partitioning occurs, and neither the laptop or the CTC computer is able to see all nodes. If you want laptops to connect directly to end network elements, you need to create static routes (see Scenario 5) or enable the ONS 15310-MA SDH proxy server (see Scenario 7).

Figure 8-3 Scenario 3: Using Proxy ARP

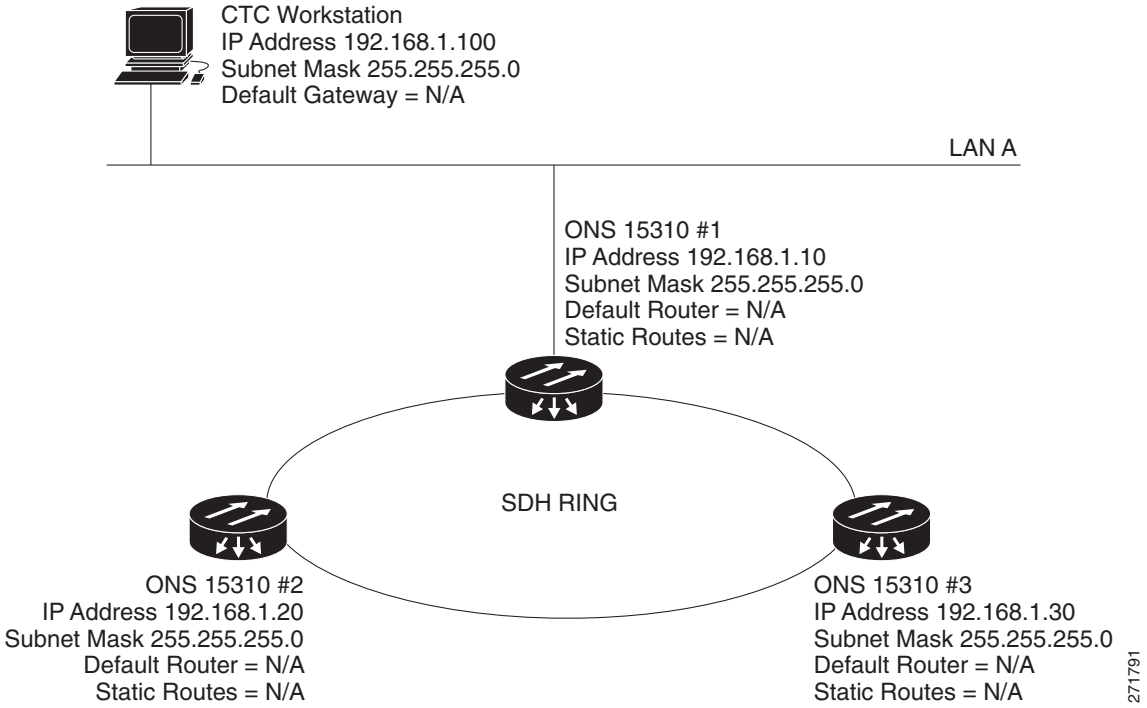

You can also use proxy ARP to communicate with hosts attached to the craft Ethernet ports of DCC-connected nodes (Figure 8-4). The node with an attached host must have a static route to the host. Static routes are propagated to all DCC peers using OSPF. The existing proxy ARP node is the gateway for additional hosts. Each node examines its routing table for routes to hosts that are not connected to the DCC network but are within the subnet. The existing proxy server replies to ARP requests for these additional hosts with the node MAC address. The existence of the host route in the routing table ensures that the IP packets addressed to the additional hosts are routed properly. Other than establishing a static route between a node and an additional host, no provisioning is necessary. The following restrictions apply:

- Only one node acts as the proxy ARP server for any given additional host.
- A node cannot be the proxy ARP server for a host connected to its Ethernet port.

In Figure 8-4, Node 1 announces to Node 2 and 3 that it can reach the CTC host. Similarly, Node 3 announces that it can reach the ONS 152xx. The ONS 152xx is shown as an example; any network element can be set up as an additional host.

Figure 8-4 Scenario 3: Using Proxy ARP with Static Routing

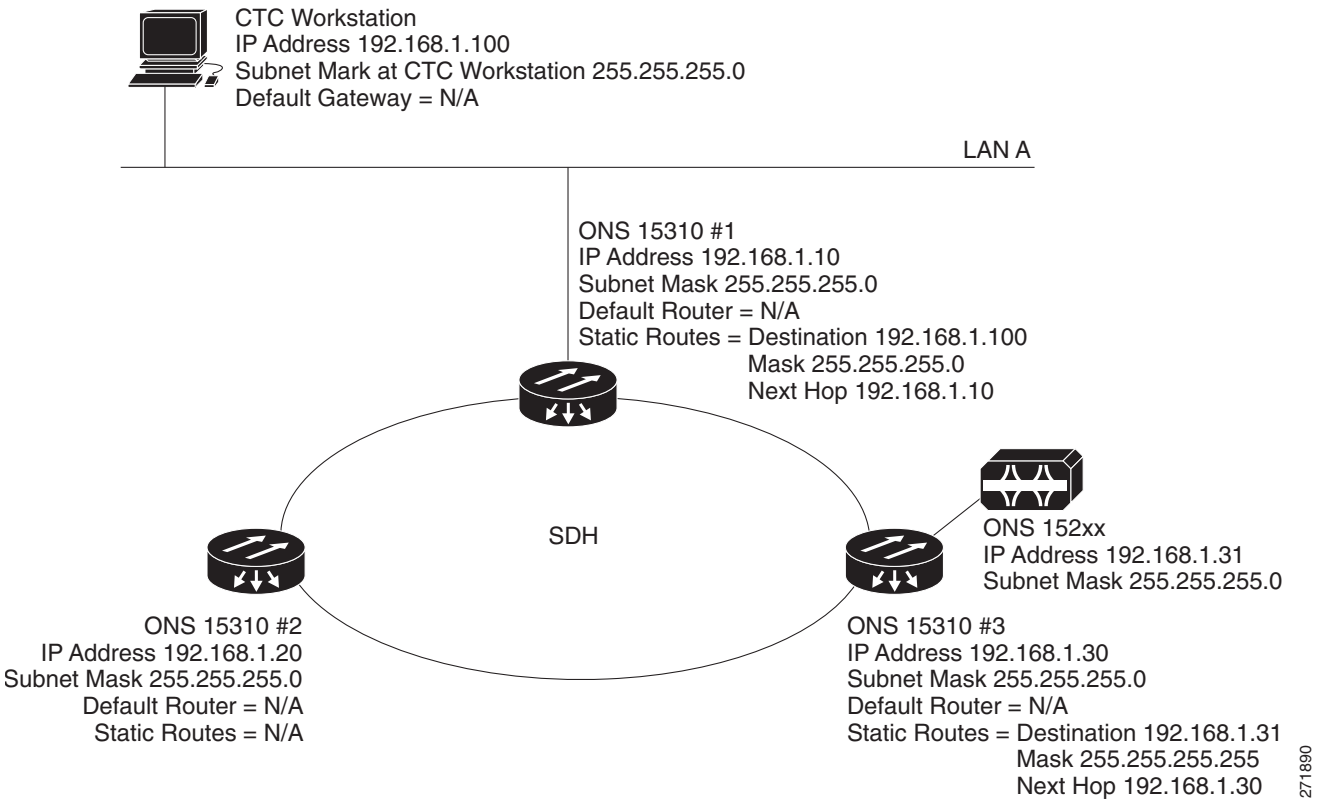

8.2.4 Scenario 4: Default Gateway on CTC Computer

Scenario 4 is similar to Scenario 3, but ONS 15310-MA SDH Node 2 and Node 3 reside on different subnets, 192.168.2.0 and 192.168.3.0, respectively (Figure 8-5). Node 1 and the CTC computer are on subnet 192.168.1.0. Proxy ARP is not used because the network includes different subnets. For the CTC computer to communicate with Nodes 2 and 3, Node 1 is entered as the default gateway on the CTC computer.

Figure 8-5 Scenario 4: Default Gateway on a CTC Computer

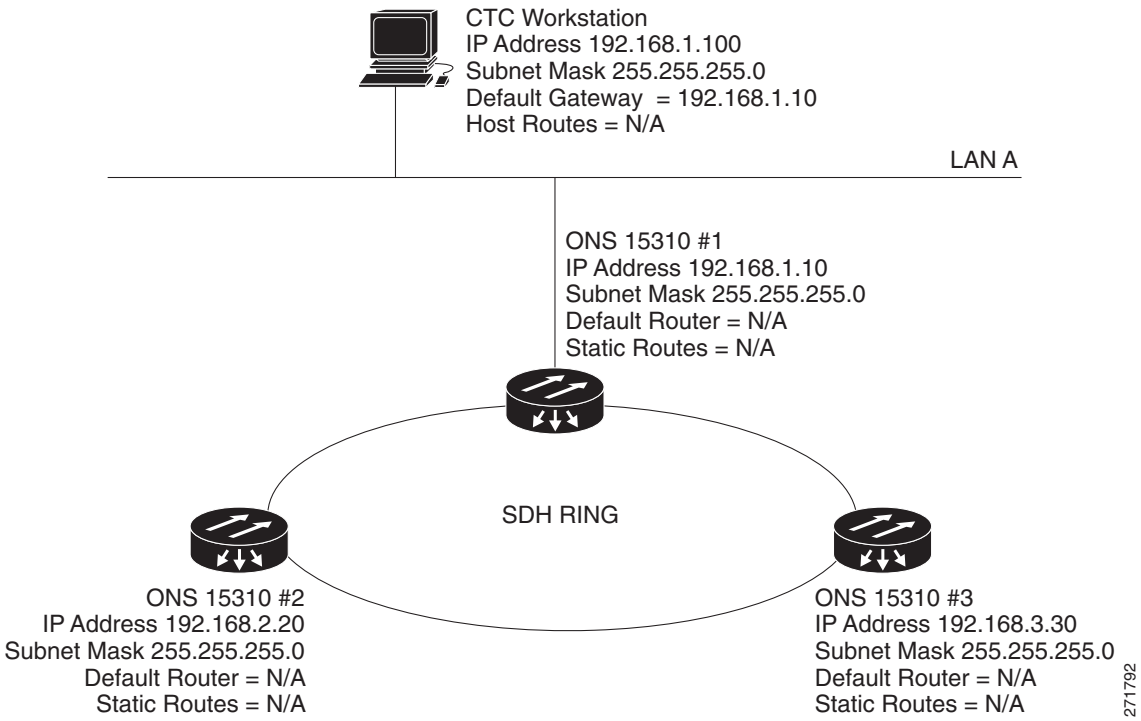

8.2.5 Scenario 5: Using Static Routes to Connect to LANs

Static routes are used for two purposes:

- To connect ONS 15310-MA SDH nodes to CTC sessions on one subnet that are connected by a router to ONS 15310-MA SDH nodes 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 15310-MA SDH nodes residing on the same subnet.

In Figure 8-6, 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 15310-MA SDH nodes residing on different subnets are connected through Node 1 to the router through interface B. Because Nodes 2 and 3 are on different subnets, proxy ARP does not enable Node 1 as a gateway. To connect to CTC computers on LAN A, a static route is created on Node 1.

Figure 8-6 Scenario 5: Static Route with One CTC Computer Used as a Destination

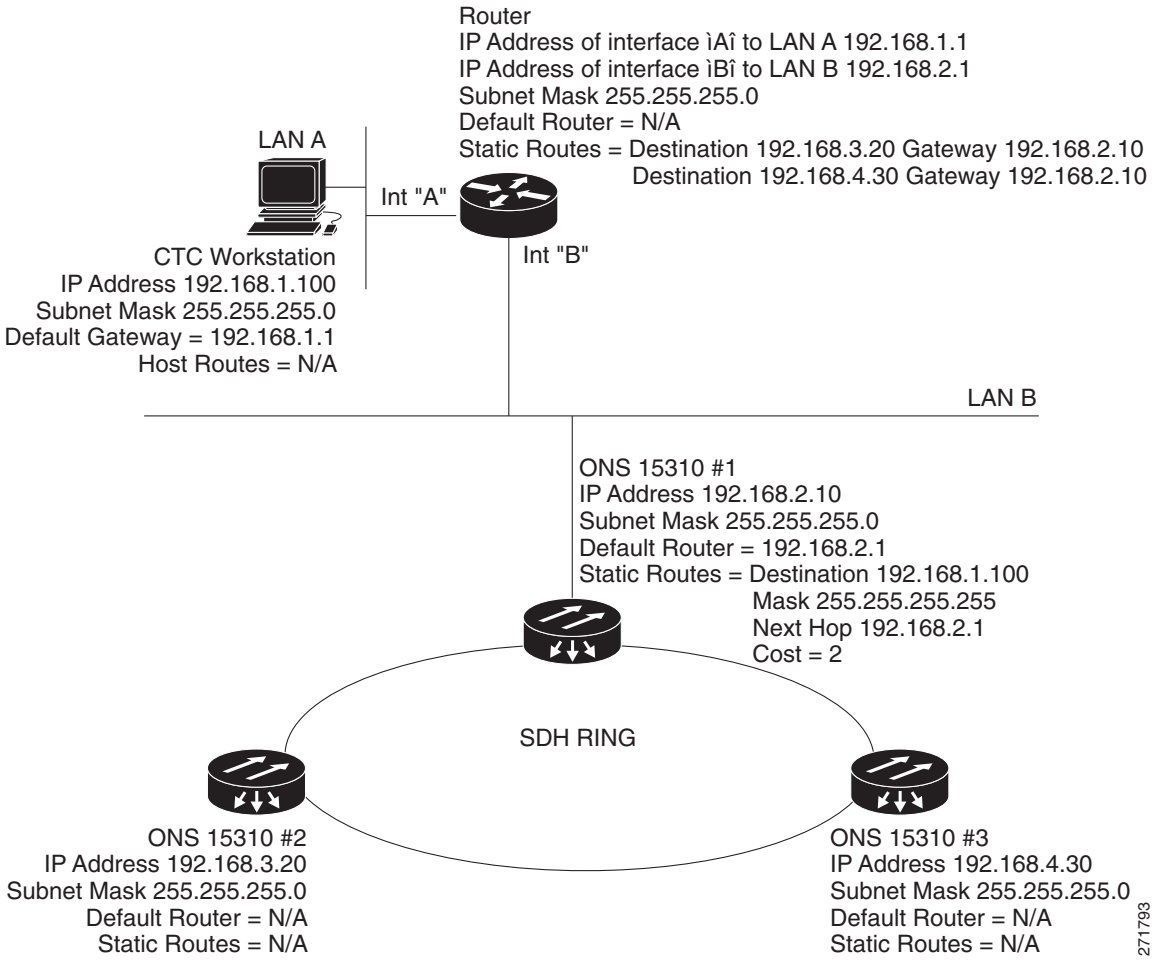

The destination and subnet mask entries control access to the ONS 15310-MA SDH nodes:

- If a single CTC computer is 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 a router, enter a destination of 0.0.0.0 and a subnet mask of 0.0.0.0. Figure 8-7 shows an example.

The IP address of router interface B is entered as the next hop, and the cost (number of hops from source to destination) is 2.

Figure 8-7 Scenario 5: Static Route with Multiple LAN Destinations

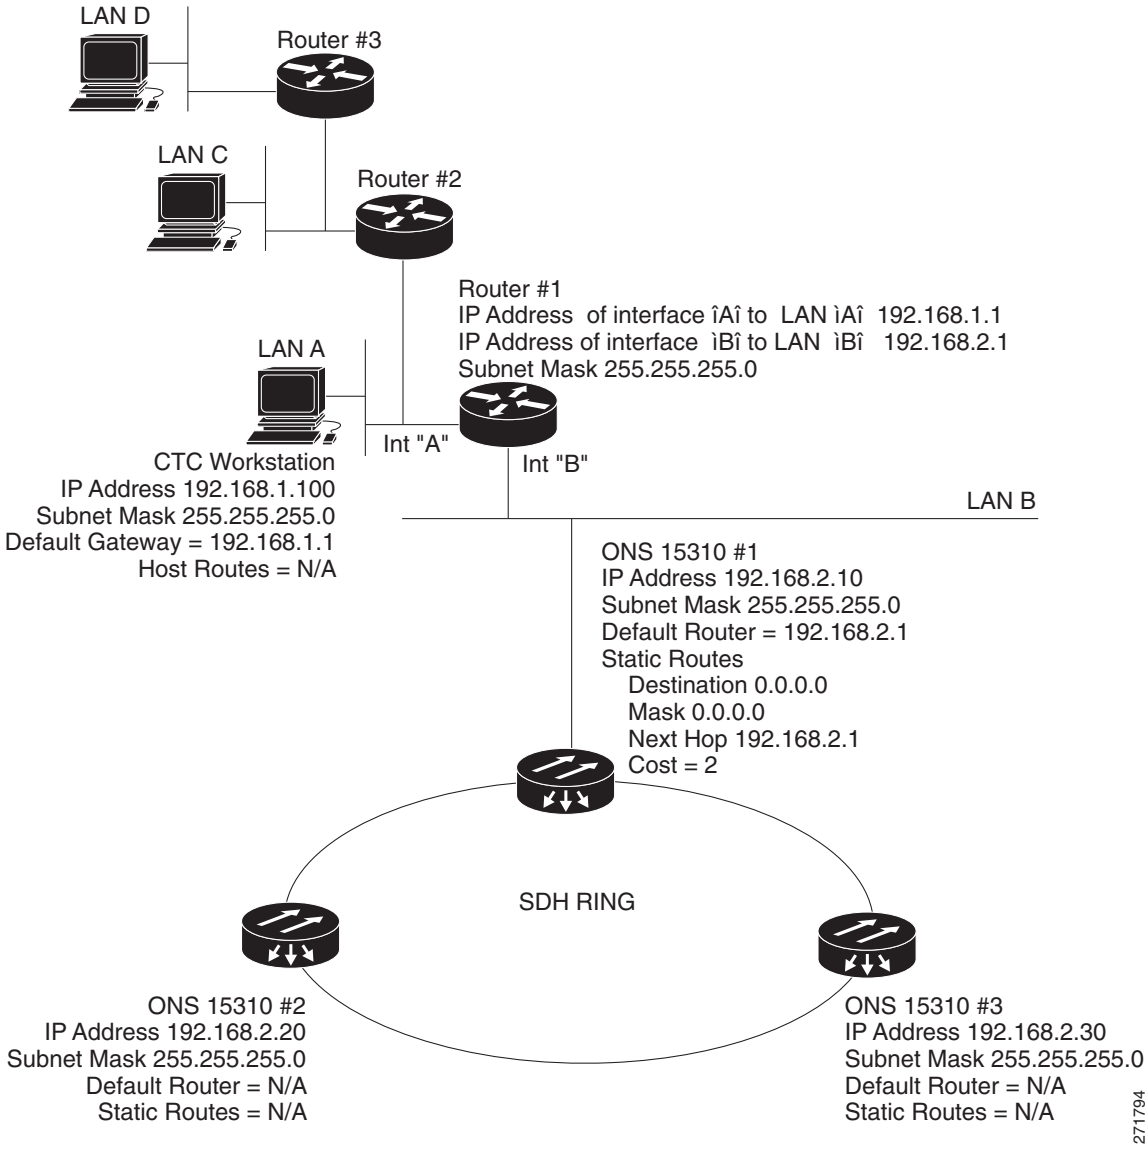

8.2.6 Scenario 6: Using OSPF

Open Shortest Path First (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 the status of their links to their neighbors. Link-state protocols advertise their directly connected networks and their active links. Each link state router captures the link state "advertisements" 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. Routes are recalculated when topology changes occur.

The ONS 15310-MA SDH uses OSPF protocol in internal ONS 15310-MA SDH networks for node discovery, circuit routing, and node management. You can enable OSPF on the ONS 15310-MA SDH so that the ONS 15310-MA SDH topology is sent to OSPF routers on a LAN. Advertising the ONS 15310-MA SDH network topology to LAN routers eliminates the need to enter static routes for ONS 15310-MA SDH subnetworks manually.

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. Every OSPF network has one backbone area called “area 0.” All other OSPF areas must connect to area 0.

When you enable an ONS 15310-MA SDH OSPF topology for advertising to an OSPF network, you must assign an OSPF area ID in decimal format to the network. Coordinate the area ID number assignment with your LAN administrator. All DCC-connected ONS 15310-MA SDH nodes should be assigned the same OSPF area ID.

Figure 8-8 shows a network enabled for OSPF.

Figure 8-8 Scenario 6: OSPF Enabled

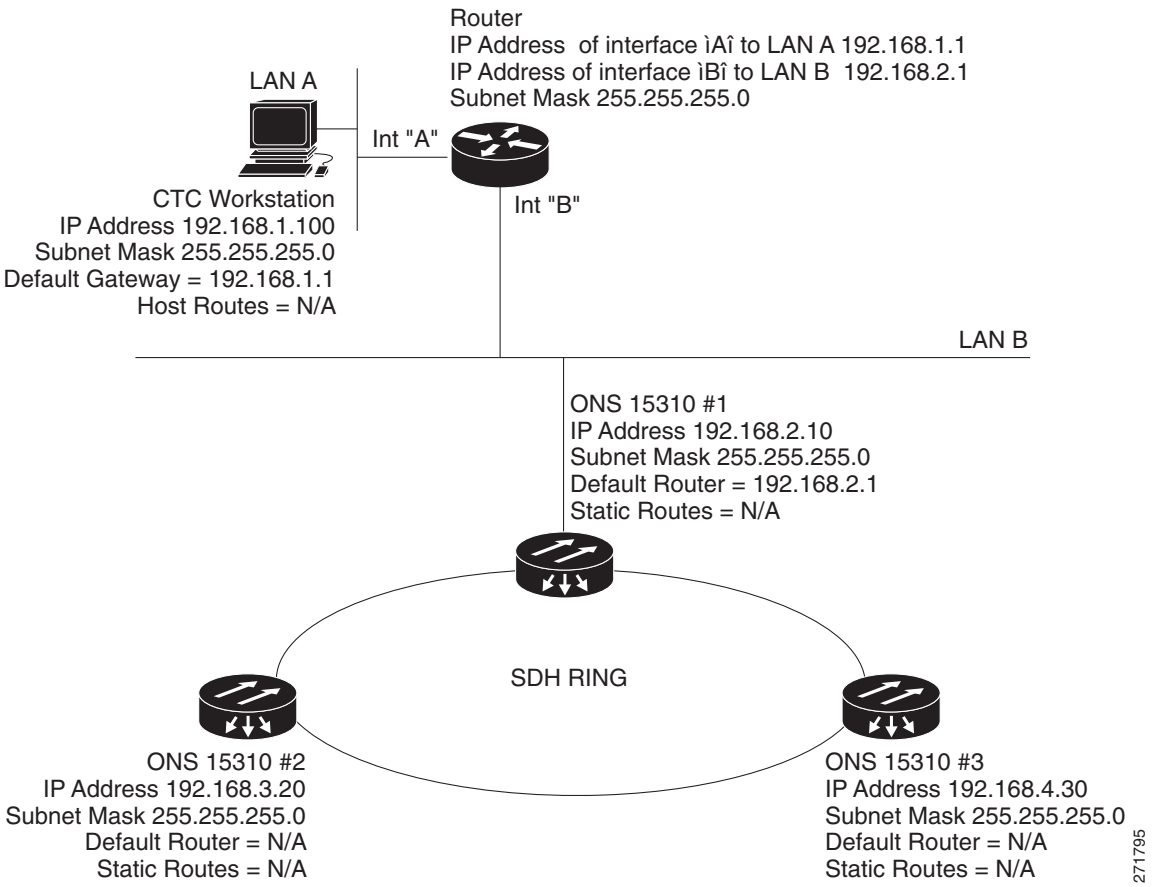

Figure 8-9 shows the same network without OSPF. Static routes must be manually added to the router for CTC computers on LAN A to communicate with Nodes 2 and 3 because these nodes reside on different subnets.

Figure 8-9 Scenario 6: OSPF Not Enabled

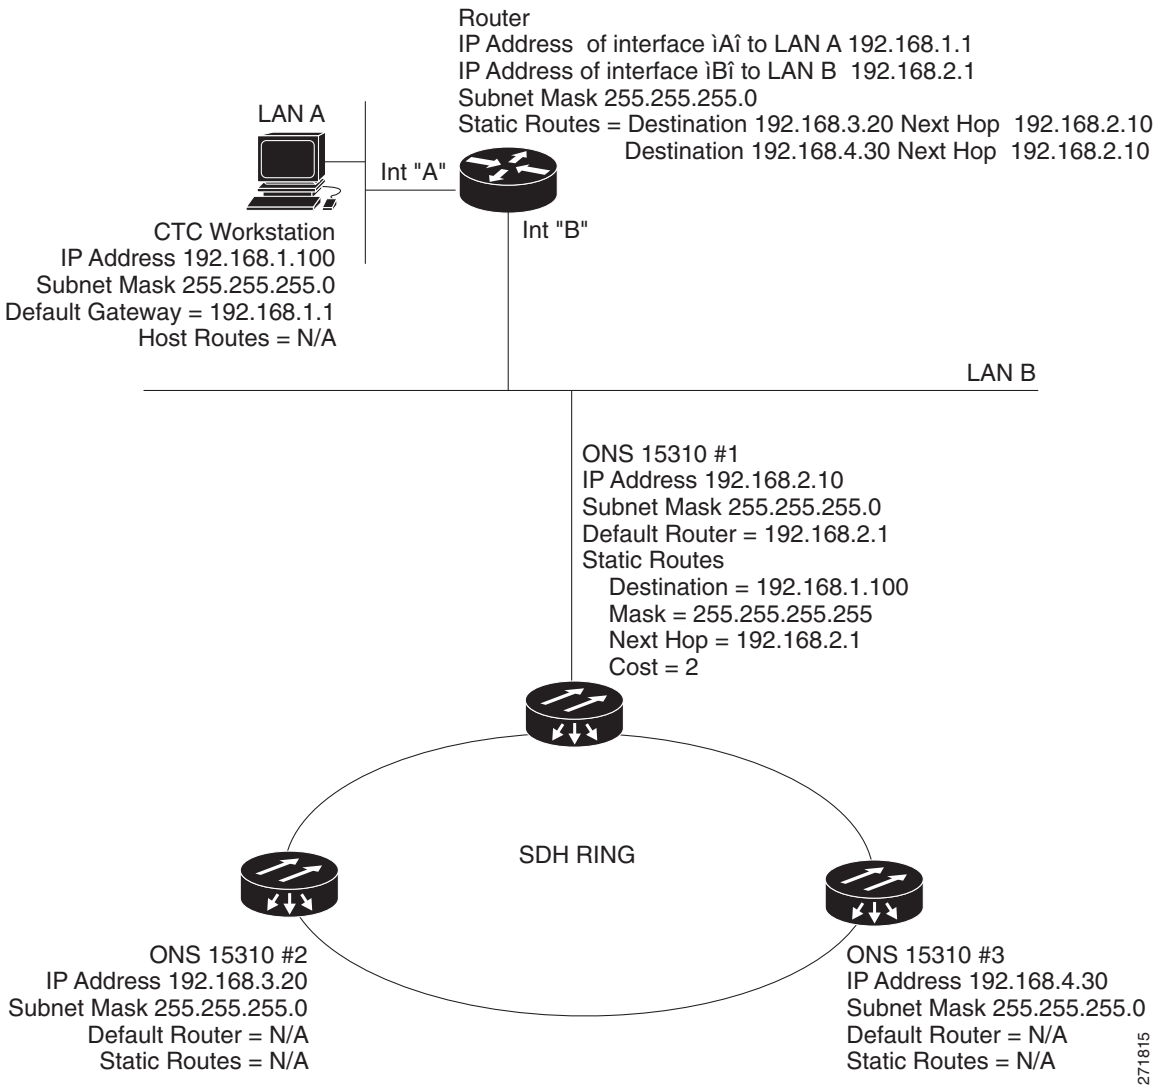

8.2.7 Scenario 7: Provisioning the ONS 15310-MA SDH Proxy Server

The ONS 15310-MA SDH proxy server is a set of functions that allows you to network ONS 15310-MA SDH nodes in environments where visibility and accessibility between nodes and CTC computers must be restricted. For example, you can set up a network so that field technicians and network operating center (NOC) personnel can both access the same nodes while preventing the field technicians from accessing the NOC LAN. To do this, one ONS 15310-MA SDH node is provisioned as a gateway network element (GNE) and the other nodes are provisioned as end network elements (ENEs). The GNE tunnels connections between CTC computers and ENEs, which provides management capability while preventing access for non-ONS 15310-MA SDH management purposes.

The ONS 15310-MA SDH 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 depend on whether the packet arrives at the DCC or CRAFT port Ethernet interface. Table 8-3 on page 8-15 and Table 8-4 on page 8-16 provide the filtering rules.
- Processes SNTP (Simple Network Timing Protocol) and NTP (Network Timing Protocol) requests. Element ONS 15310-MA SDH NEs can derive time-of-day from an SNTP/NTP LAN server through the GNE.
- Process SNMPv1 traps. The GNE receives SNMPv1 traps from the ENE and forwards them to all provisioned SNMPv1 trap destinations.

The ONS 15310-MA SDH proxy server is provisioned using the Enable proxy server on port check box on the Provisioning > Network > General tab. If checked, the ONS 15310-MA SDH serves as a proxy for connections between CTC clients and ONS 15310-MA SDH nodes that are DCC-connected to the proxy ONS 15310-MA SDH. 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 the Enable proxy server on port check box is not checked, the node does not proxy for any CTC clients, although any established proxy connections continue until the CTC client exits. In addition, you can set the proxy server as an ENE or a GNE:

- External Network Element (ENE)—If set as an ENE, the ONS 15310-MA SDH neither installs nor advertises default or static routes. CTC computers can communicate with the node using the craft port, but they cannot communicate directly with any other DCC-connected node.

  In addition, firewall is enabled, which means that the node prevents IP traffic from being routed between the DCC and the LAN port. The ONS 15310-MA SDH can communicate with machines connected to the LAN port or connected through the DCC. However, the DCC-connected machines cannot communicate with the LAN-connected machines, and the LAN-connected machines cannot communicate with the DCC-connected machines. A CTC client using the LAN to connect to the firewall-enabled node can use the proxy capability to manage the DCC-connected nodes that would otherwise be unreachable. A CTC client connected to a DCC-connected node can only manage other DCC-connected nodes and the firewall itself.
- Gateway Network Element (GNE)—If set as a GNE, the CTC computer is visible to other DCC-connected nodes and firewall is enabled.
- Proxy-only—If Proxy-only is selected, CTC cannot communicate with any other DCC-connected ONS 15310-MA SDH nodes and firewall is not enabled.

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.

Note ENEs that belong to different private subnetworks do not need to have unique IP addresses. Two ENEs that are connected to different GNEs can have the same IP address. However, ENEs that connect to the same GNE must always have unique IP addresses.

Figure 8-10 shows an ONS 15310-MA SDH proxy server implementation. A GNE is connected to a central office LAN and to 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 ENEs. However, the craft technicians must be prevented from accessing or seeing the NOC or central office LANs.

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

Figure 8-10 ONS 15310-MA SDH Proxy Server with GNE and ENEs on the Same Subnet

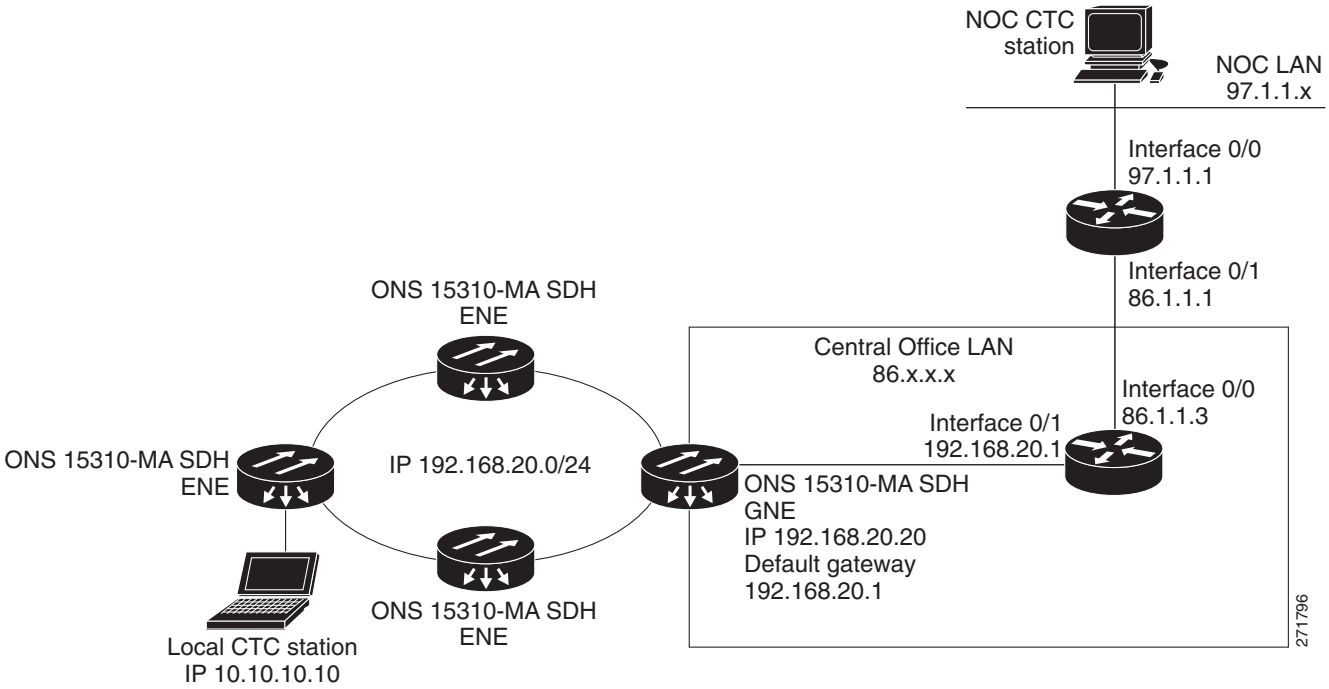

Table 8-2 shows recommended settings for ONS 15310-MA SDH GNEs and ENEs in the configuration shown in Figure 8-10.

Table 8-2 ONS 15310-MA SDH GNE and ENE Settings

Setting	ONS 15310-MA SDH GNE	ONS 15310-MA SDH ENE
OSPF	Off	Off
SNTP Server (if used)	SNTP server IP address	Set to node GNE IP address
SNMP (if used)	SNMPv1 trap destinations	Set SNMPv1 trap destinations to node GNE

Figure 8-11 shows the same proxy server implementation with ONS 15310-MA SDH ENEs on different subnets. In this example, GNEs and ENEs are provisioned with the settings shown in Table 8-2.

Figure 8-11 Scenario 7: Proxy Server with GNE and ENEs on Different Subnets

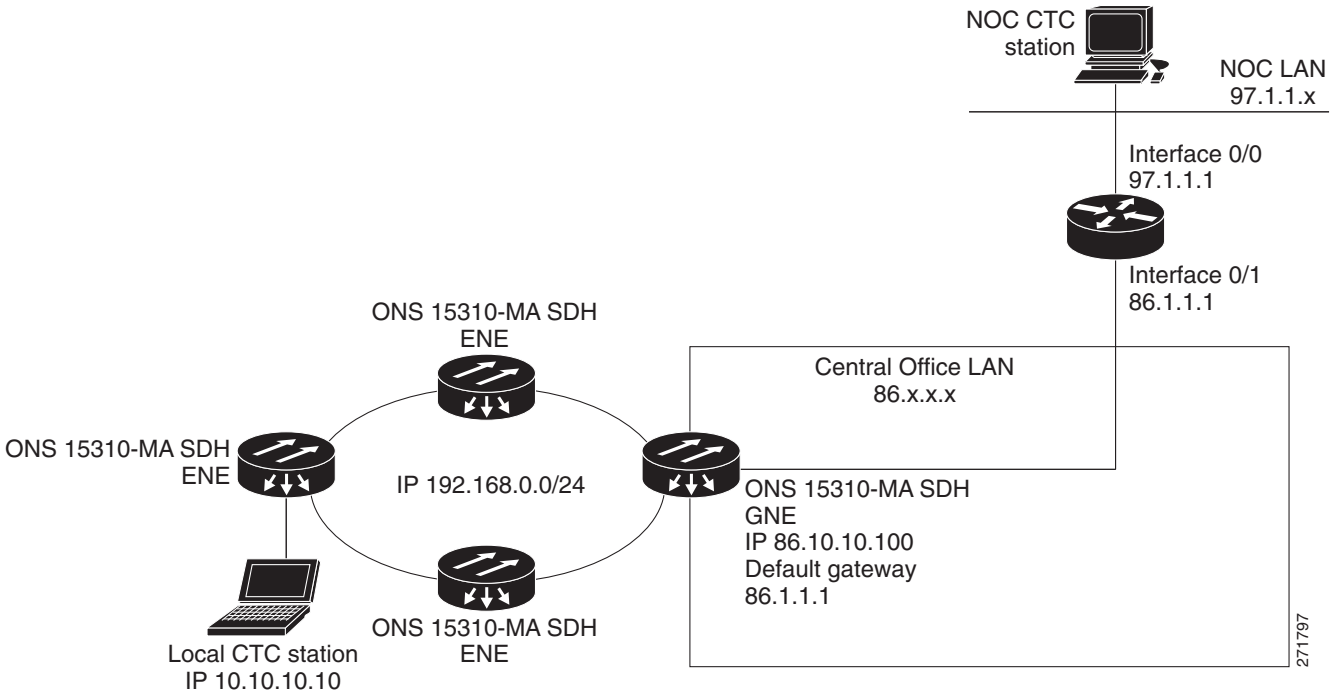

Figure 8-12 shows the implementation with ONS 15310-MA SDH ENEs in multiple rings. In this example, GNEs and ENEs are provisioned with the settings shown in Table 8-2.

Figure 8-12 Scenario 7: Proxy Server with ENEs on Multiple Rings

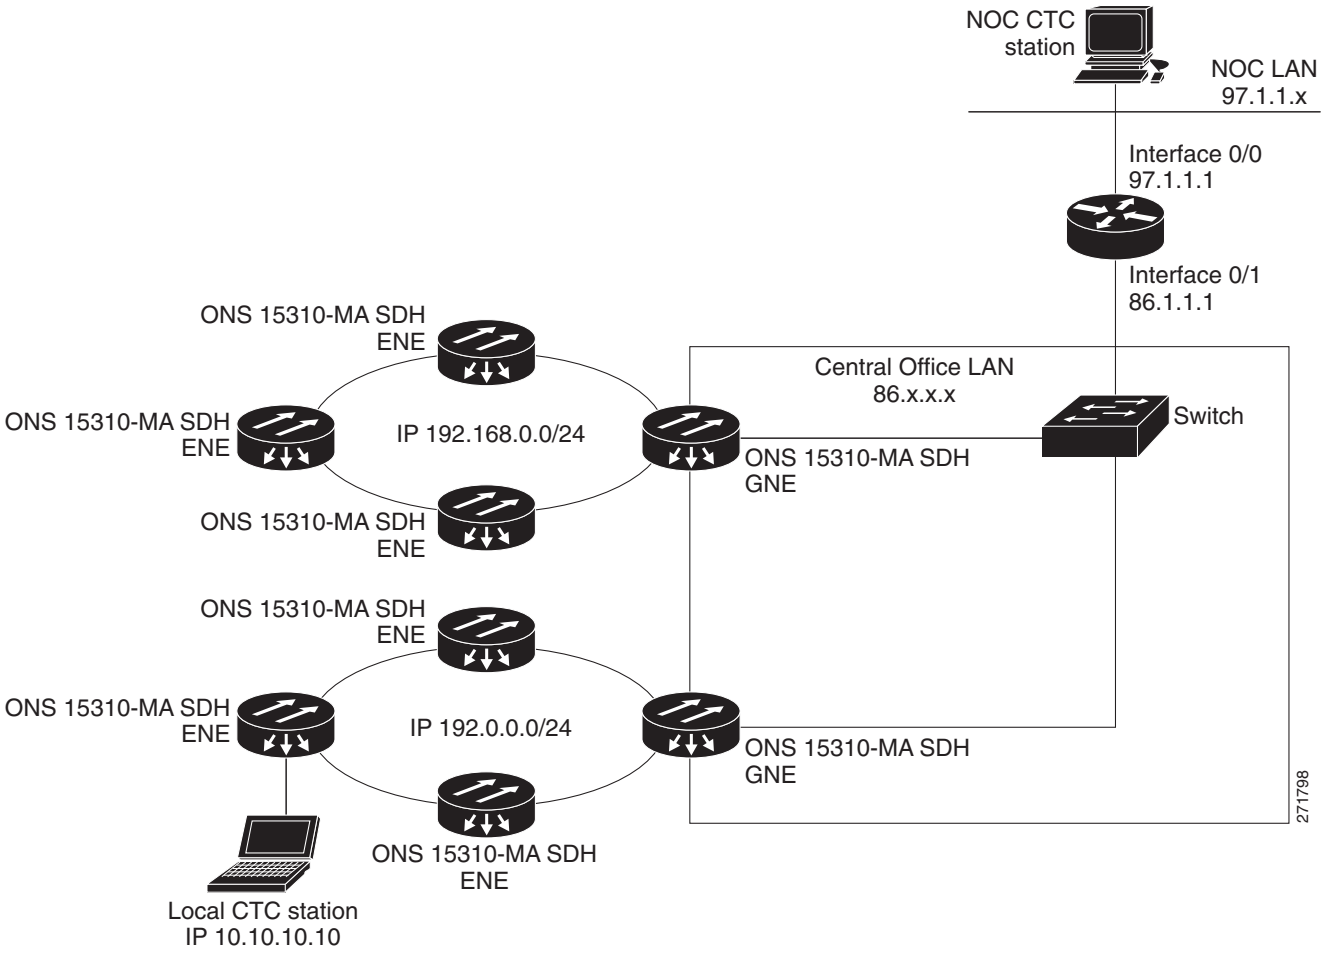

Table 8-3 shows the rules the ONS 15310-MA SDH follows to filter packets when Enable Firewall is enabled.

Table 8-3 Proxy Server Firewall Filtering Rules

Packets arriving at:	Are accepted if the IP destination address is:
15310E-CTX-K9 Ethernet interface	• The ONS 15310-MA SDH shelf itself • The ONS 15310-MA SDH's subnet broadcast address • Within the 224.0.0.0/8 network (reserved network used for standard multicast messages) • Subnet mask = 255.255.255.255
DCC interface	• The ONS 15310-MA SDH itself • Any destination that is connected through another DCC interface • Within the 224.0.0.0/8 network

Table 8-4 shows additional rules that apply if the packet addressed to the ONS 15310-MA SDH is discarded. Rejected packets are silently discarded.

Table 8-4 Proxy Server Firewall Filtering Rules When the Packet is Addressed to the ONS 15310-MA SDH

Packets Arrive At	Accepts	Rejects
15310E-CTX-K9 LAN port	• All User Datagram Protocol (UDP) packets except those in the Rejected column	• UDP packets addressed to the SNMP trap relay port (391)
DCC interface	• All UDP packets • All TCP packets except those packets addressed to the Telnet and SOCKS proxy server ports • OSPF packets • Internet Control Message Protocol (ICMP) packets	• TCP packets addressed to the Telnet port • TCP packets addressed to the proxy server port • All packets other than UDP, TCP, OSPF, ICMP

If you implement the proxy server, keep the following rules in mind:

1. All DCC-connected ONS 15310-MA SDH nodes on the same Ethernet segment must have the same Craft Access Only setting. Mixed values produce unpredictable results, and might leave some nodes unreachable through the shared Ethernet segment.
2. All DCC-connected ONS 15310-MA SDH nodes on the same Ethernet segment must have the same Enable Firewall setting. Mixed values produce unpredictable results. Some nodes might become unreachable.
3. If you check Enable Firewall, always check Enable Proxy. If Enable Proxy is unchecked, CTC is not able to see nodes on the DCC side of the ONS 15310-MA SDH.
4. If Craft Access Only is checked, check Enable Proxy. If Enable Proxy is not checked, CTC is not able to see nodes on the DCC side of the ONS 15310-MA SDH.

If nodes become unreachable in cases 1, 2, and 3, you can correct the setting with one of the following actions:

- Disconnect the craft computer from the unreachable ONS 15310-MA SDH. Connect to the ONS 15310-MA SDH through another ONS 15310-MA SDH in the network that has a DCC connection to the unreachable node.
- Disconnect the Ethernet cable from the unreachable ONS 15310-MA SDH. Connect a CTC computer directly to the ONS 15310-MA SDH.

8.3 Routing Table

ONS 15310-MA SDH routing information appears on the Maintenance > Routing Table tabs. 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 the listed route has been used.
- Interface—Shows the ONS 15310-MA SDH interface used to access the destination.

- cpm0—The ONS 15310-MA SDH Ethernet interface (RJ45 LAN jack)
- pdcc0—An RS-DCC interface, that is, an STMN trunk port identified as the RS-DCC termination
- lo0—A loopback interface

Table 8-5 shows sample routing entries for an ONS 15310-MA SDH.

Table 8-5 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 is 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 is sent to this gateway.
- Interface (cpm0) indicates that the ONS 15310-MA SDH 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 a destination.
- 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 15310-MA SDH 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 an SDH RS-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 IP address 172.20.214.93.
- Interface (pdcc0) indicates that an SDH RS-DCC interface is used to reach the gateway.

8.4 External Firewalls

Table 8-6 shows the ports that are used by the 15310E-CTX-K9 cards.

Table 8-6 Ports Used by the 15310E-CTX-K9

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)	NA
161	SNMP traps destinations	D
162	SNMP traps destinations	D
513	rlogin	NA
683	CORBA IIOP	OK
1080	Proxy server (socks)	D
2001-2017	I/O card Telnet	D
2018	DCC processor on active 15310-MA SDH-CTX	D
2361	TL1	D
3082	Raw TL1	D
3083	TL1	D
5001	Multiplex-section shared protection ring (MS-SPRing) server port	D
5002	MS-SPRing client port	D
7200	SNMP alarm input port	D
9100	EQM port	D
9401	TCC boot port	D
9999	Flash manager	D

Table 8-6 Ports Used by the 15310E-CTX-K9 (continued)

Port	Function	Action[1]
10240-12287	Proxy client	D
57790	Default TCC listener port	OK

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

The following access control list (ACL) examples show a firewall configuration when the proxy server gateway setting is not enabled. In the example, the CTC workstation address is 192.168.10.10 and the ONS 15310-MA SDH address is 10.10.10.100. The firewall is attached to the GNE, so the inbound path is CTC to the GNE and the outbound path is from the GNE to CTC. The CTC 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 the 15310-MA SDH 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 the 15310-MA SDH GNE (port 57790)
***
access-list 100 remark

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 15310-MA SDH (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 15310-MA SDH GNE to CTC ***
```

The following ACL examples show a firewall configuration when the proxy server gateway setting is enabled. As with the first example, the CTC workstation address is 192.168.10.10 and the ONS 15310-MA SDH address is 10.10.10.100. The firewall is attached to the GNE, so the inbound path is CTC to the GNE and the outbound path is from the GNE to CTC. The CTC 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 the 15310-MA SDH 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 15310-MA SDH GNE proxy server
(port 1080) ***
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 from CTC to the 15310-MA SDH GNE ***
access-list 101 remark *** Outbound ACL, NE -> CTC ***
access-list 101 remark
access-list 101 permit tcp host 10.10.10.100 eq 1080 host 192.168.10.10
access-list 101 remark *** allows alarms and other communications from the 15310-MA SDH
(proxy server) 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 15310-MA SDH GNE to CTC ***
```

8.5 Open GNE

The ONS 15310-MA SDH 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 RS-DCC and MS-DCC 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 RS-DCC and MS-DCC 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 8-13 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 8-13 Proxy and Firewall Tunnels for Foreign Terminations

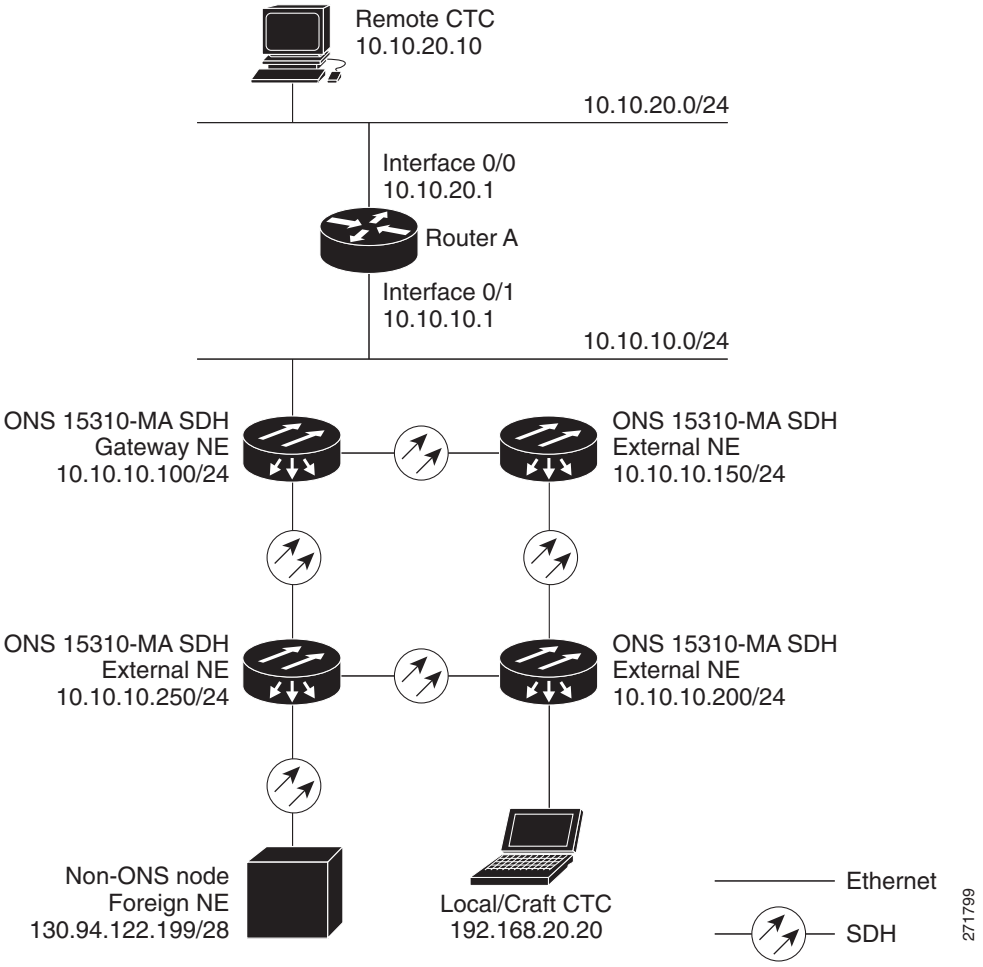

Figure 8-14 shows a remote node connected to an ENE Ethernet port. Proxy and firewall tunnels are useful in this example because the GNE would otherwise block IP access between the PC and foreign node. This configuration also requires a firewall tunnel on the ENE.

Figure 8-14 Foreign Node Connection to an ENE Ethernet Port

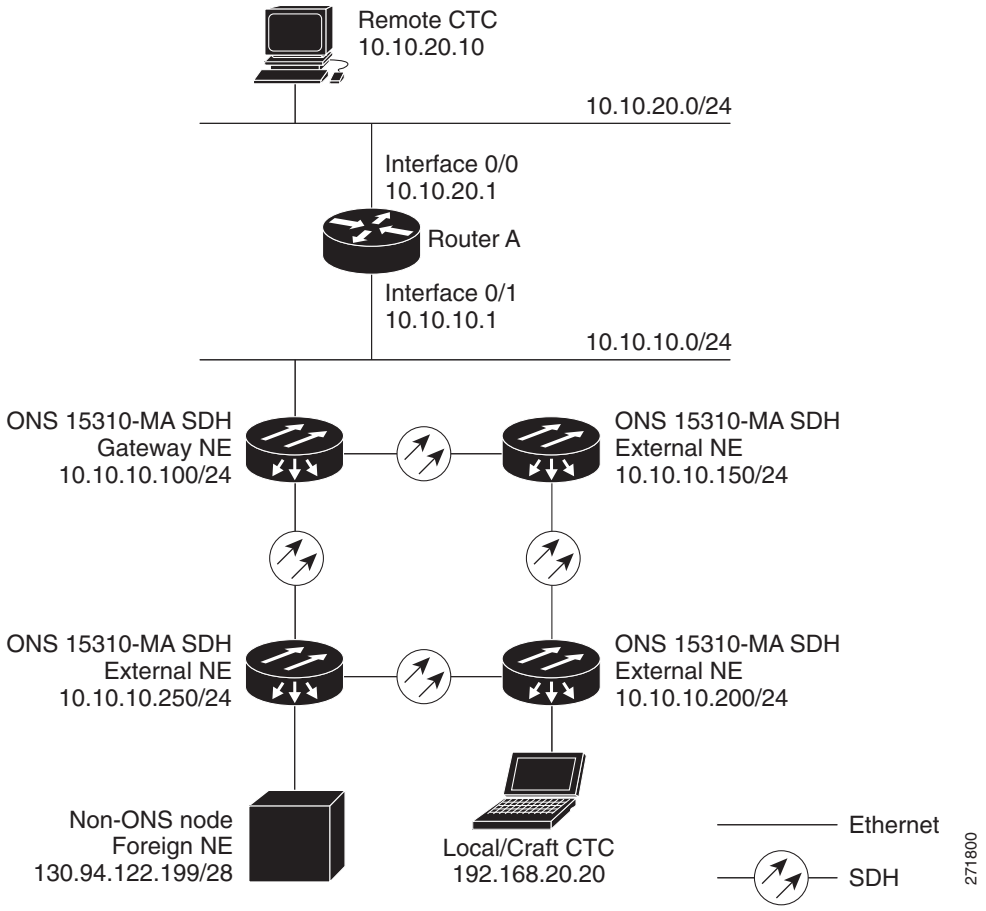

8.6 TCP/IP and OSI Networking

ONS 15310-MA SDH DCN communication is based on the TCP/IP protocol suite. However, ONS 15310-MA SDH nodes 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 8-7 shows the protocols that are involved when TCP/IP-based NEs are networked with OSI-based NEs.

Table 8-7 TCP/IP and OSI Protocols

OSI Model	IP Protocols	OSI Protocols		IP-OSI Tunnels
Layer 7 Application	• TL1 • FTP • HTTP • Telnet • IIOP	• TARP[1]	• TL1 (over OSI) • FTAM[2] • ACSE[3]	
Layer 6 Presentation			• Administrative State[4]	
Layer 5 Session			• Session	
Layer 4 Transport	• TCP • UDP		• TP (Transport) Class 4	• IP-over-CLNS[5] tunnels
Layer 3 Network	• IP • OSPF	• CLNP[6] • ES-IS[7] • IS-IS[8]		
Layer 2 Data link	• PPP	• PPP • LAP-D[9]		
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. Administrative State = Presentation layer
5. CLNS = Connectionless Network Layer Service
6. CLNP = Connectionless Network Layer Protocol
7. ES-IS = End System-to-Intermediate System
8. IS-IS = Intermediate System-to-Intermediate System
9. LAP-D = Link Access Protocol on the D Channel

8.6.1 Point-to-Point Protocol

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 Connectionless Network Protocol (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 point-to-point connections.

CTC automatically enables IP over PPP whenever you create an RS-DCC or MS-DCC. The RS-DCC or MS-DCC can be provisioned to support OSI over PPP.

8.6.2 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 15310-MA SDH RS-DCC 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.

8.6.3 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 15310-MA SDH support 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 8-8. NSAP field values are in hexadecimal format. All NSAPs are editable and shorter NSAPs can be used; however, NSAPs for all NEs residing within the same OSI network area usually have the same NSAP format.

Table 8-8 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.
System	System identifier	The ONS 15310-MA SDH system identifier is set to its IEEE 802.3 MAC address.

Table 8-8 NSAP Fields (continued)

Field	Definition	Description
SEL	Selector	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 15310-MA SDH include: • 00—Network Entity Title (NET). Used to exchange PDUs in the ES-IS and IS-IS routing exchange protocols. (See the "8.6.4.1 End System-to-Intermediate System Protocol" section on page 8-28, and the "8.6.4.2 Intermediate System-to-Intermediate System Protocol" section on page 8-28.) • 1D—Selector for Transport Class 4 (and for FTAM and TL1 applications • AF—Selector for the TARP protocol • 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) 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.

Figure 8-15 shows the default ISO-DCC NSAP address delivered with the ONS 15310-MA SDH. The System ID is automatically populated with the node's MAC address.

Figure 8-15 ISO-DCC NSAP Address

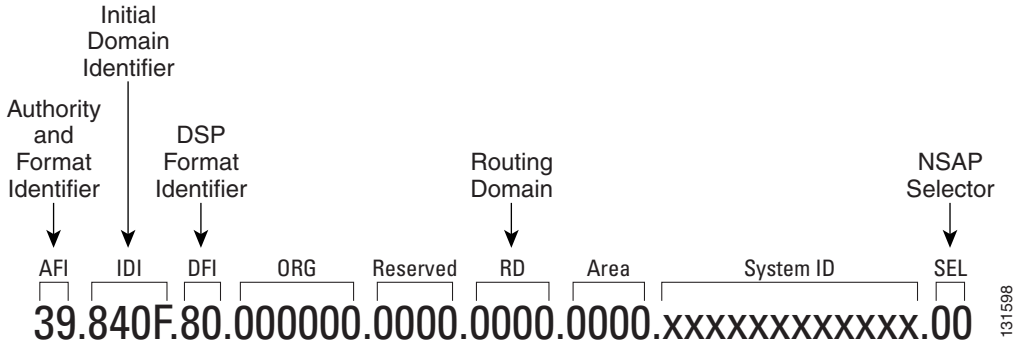

The ONS 15310-MA SDH main NSAP address is shown on the node view Provisioning > OSI > Main Setup subtab. This address is also the Router 1 primary manual area address, which is viewed and edited on the Provisioning > OSI > Routers subtab. See the "8.6.6 OSI Virtual Routers" section on page 8-32 for information about the OSI router and manual area addresses in CTC.

8.6.4 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 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 provides full connectivity to all ESs within the domains. 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 8-16 shows an example of Level 1 and Level 2 routing.

Note The ONS 15310-MA SDH do not support Level 1/Level 2 routing. Level 1/Level 2 routing is supported by the ONS 15454, ONS 15454 SDH, and the ONS 15600.

Figure 8-16 Level 1 and Level 2 OSI Routing

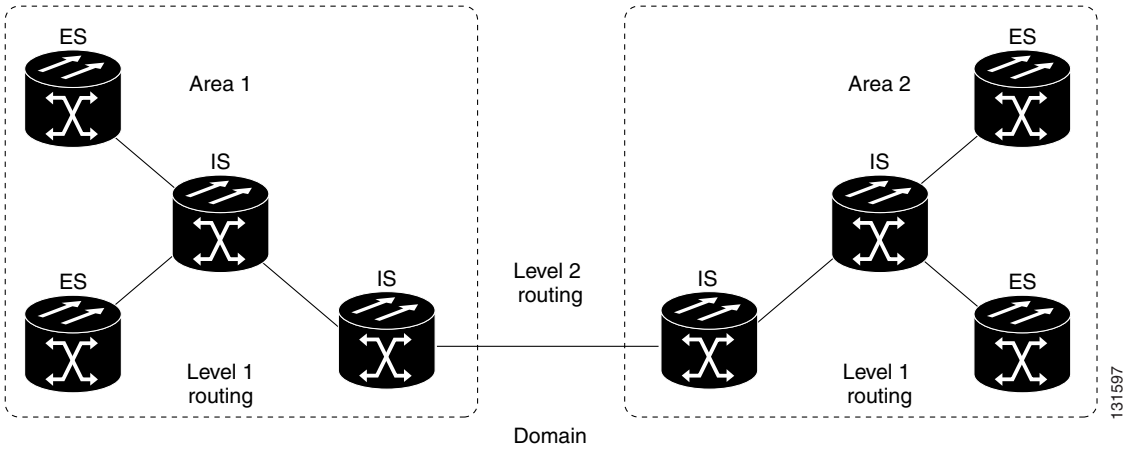

When you provision an ONS 15310-MA SDH for a network with NEs that use both the TCP/IP and OSI protocol stacks, you will provision it as one of the following:

- End System—The ONS 15310-MA SDH performs OSI ES functions and relies upon an IS for communication with nodes that reside within its OSI area.
- Intermediate System Level 1—The ONS 15310-MA SDH 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.

8.6.4.1 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.

8.6.4.2 Intermediate System-to-Intermediate System Protocol

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 15310-MA SDH. 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.

8.6.5 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 (see Table 8-8 on page 8-25).

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 8-9.

Table 8-9 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 8-10, are defined.
TID Target Length	tar-tln	1	The number of octets that are in the tar-ttg field.
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.

Table 8-9 TARP PDU Fields (continued)

Field	Abbreviation	Size (bytes)	Description
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 8-10 shows the TARP PDU types that govern TARP interaction and routing.

Table 8-10 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.

8.6.5.1 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 8-11, control TARP processing.

Table 8-11 TARP Timers

Timer	Description	Default (seconds)	Range (seconds)
E1	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
DS3/E3	Waiting for response to address resolution request	40	0–3600
T4	Timer starts when T2 expires (used during error recovery)	20	0–3600

Table 8-12 shows the main TARP processes and the general sequence of events that occurs in each process.

Table 8-12 TARP Processing Flow

Process	General TARP Flow
Find a NET that matches a TID	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 E1 is started. 3. If Timer E1 expires before a match if 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 DS3/E3 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.

8.6.5.2 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 the NET address 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 Cisco ONS 15310-MA SDH 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 tabs.

8.6.5.3 Manual TARP Adjacencies

TARP adjacencies can be manually provisioned in networks where ONS 15310-MA SDH nodes must communicate across routers or non-SDH 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-SDH NE, as shown in Figure 8-17.

Figure 8-17 Manual TARP Adjacencies

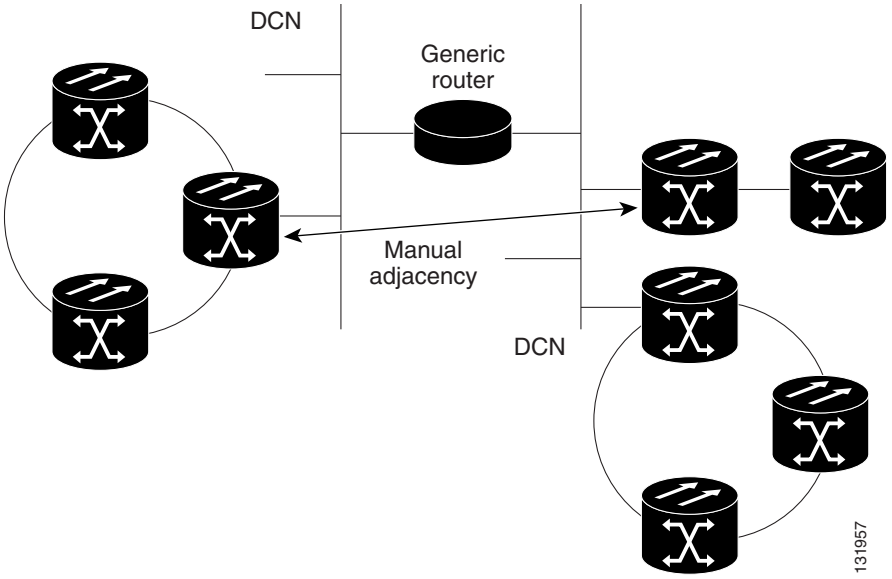

8.6.5.4 Manual TID to NSAP Provisioning

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.

8.6.6 OSI Virtual Routers

The ONS 15310-MA SDH support one OSI virtual router. The router is provisioned on the Provisioning > OSI > Routers tabs. The router has an editable manual area address and a unique NSAP System ID that is set to the node MAC address. The router can be enabled and connected to different OSI routing areas. The Router 1 manual area address and System ID create the NSAP address assigned to the node's TID. Router 1 supports OSI TARP and tunneling functions. These include:

- TARP data cache
- IP-over-CLNS tunnels
- LAN subnet

In addition to the 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.

8.6.7 IP-over-CLNS Tunnels

IP-over-CLNS tunnels are used to encapsulate IP for transport across OSI NEs. The ONS 15310-MA SDH 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 8-18 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 8-18 IP-over-CLNS Tunnel Flow

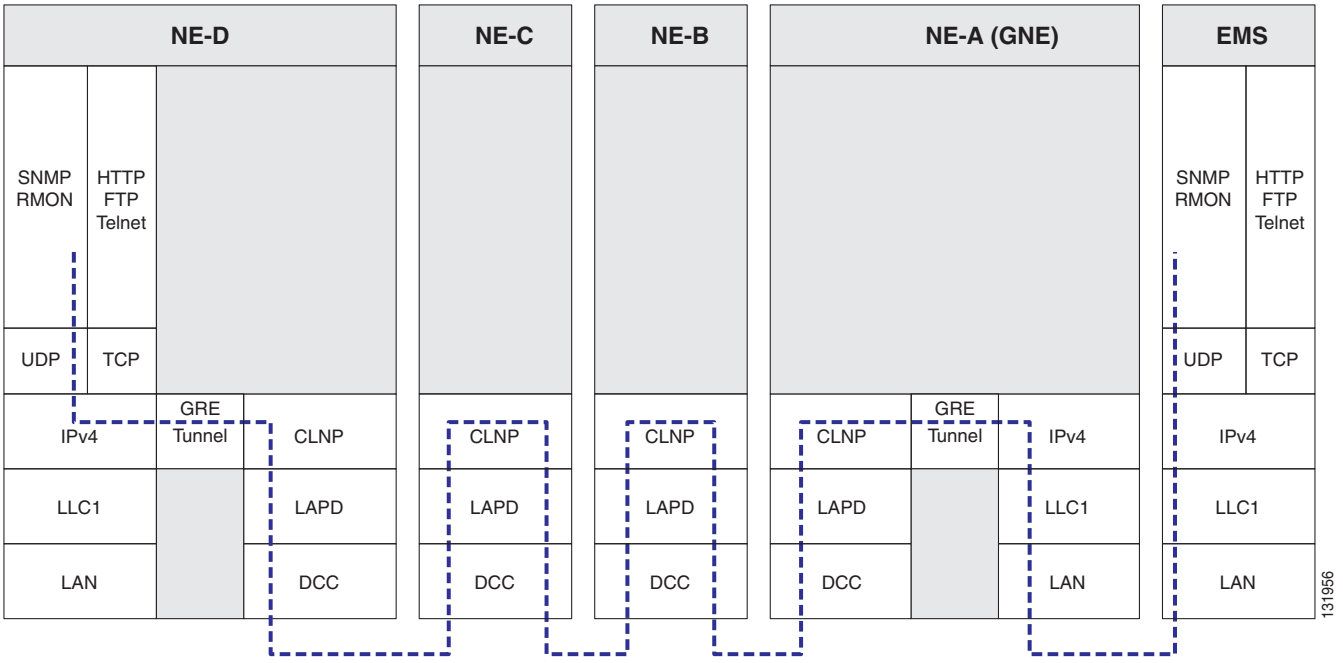

8.6.7.1 Provisioning IP-over-CLNS Tunnels

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, see the "Turn Up a Node" chapter in the *Cisco ONS 15310-MA SDH Procedure 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 8-13.

Table 8-13 IP Over CLNS Tunnel Cisco IOS Commands

Step	Step	Purpose
1	Router (config) # **interface ctunnel** *interface-number*	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** *remote-nsap-address*	Configures the destination parameter for the CTunnel. Specifies the destination NSAP1 address of the CTunnel, where the IP packets are extracted.
3	Router (config-if) # **ip address** *ip-address mask*	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, refer to the "Configuring ISO CLNS" chapter in the *Cisco IOS Apollo Domain, Banyon VINES, DECnet, ISO CLNS, and XNS Configuration Guide*.

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

Figure 8-19 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) RS-DCC and a GRE tunnel are created between the ONS NE 1 to the other vender GNE.

IP-over-CLNS tunnel provisioning 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.ccccccccccc.00 (other vendor GNE)
- Metric: 110

- Tunnel Type: GRE

IP-over-CLNS tunnel provisioning on the other vender 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.ddddddddddd.00 (ONS NE 1)
- Metric: 110
- Tunnel Type: GRE

Figure 8-19 IP Over CLNS Tunnel Scenario 1: ONS NE to Other Vender GNE

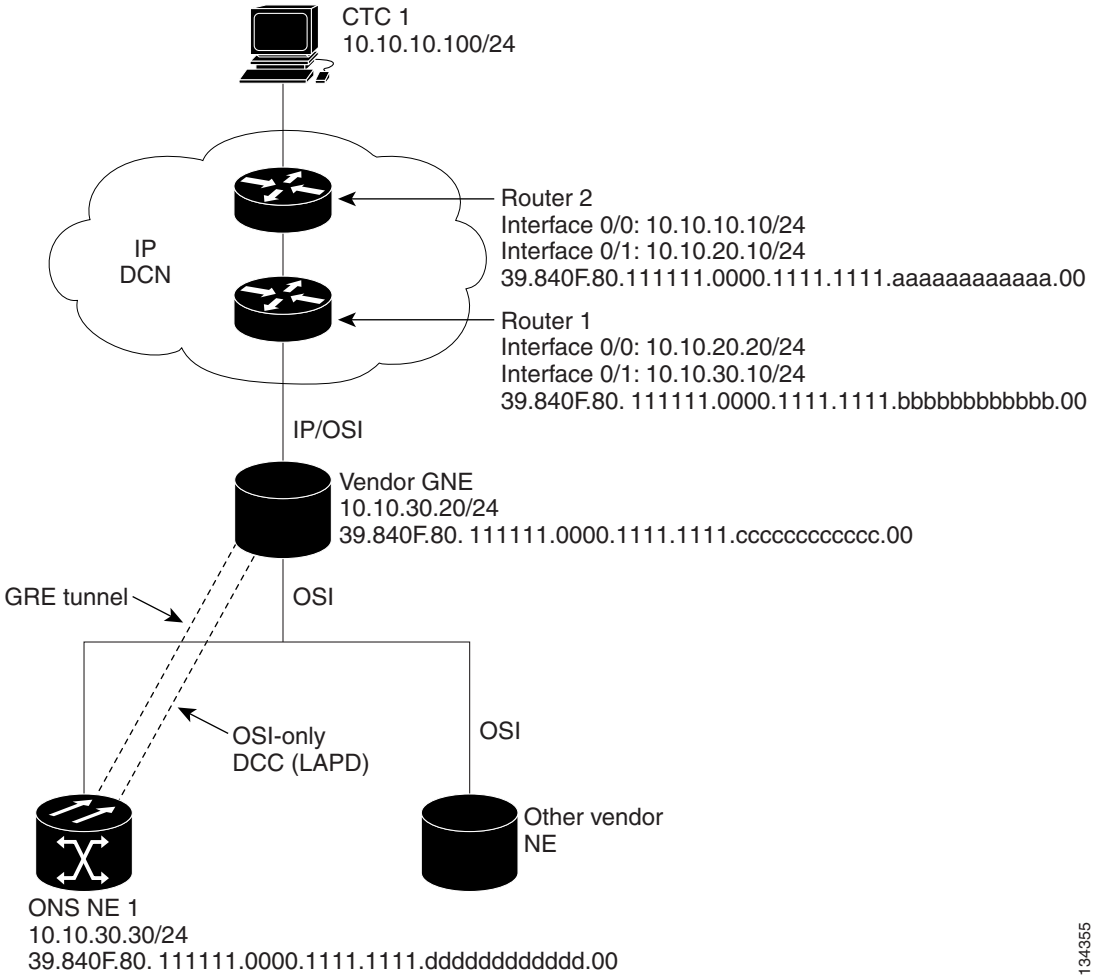

8.6.7.3 IP-Over-CLNS Tunnel Scenario 2: ONS Node to Router

Figure 8-20 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) RS-DCC is created between the ONS NE 1 and the other vender 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 provisioning 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) provisioning 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.dddddddddddd.00
interface Ethernet0/1
    clns router isis
router isis
    net 39.840F.80.1111.0000.1111.1111.bbbbbbbbbbbb.00
```

Figure 8-20 IP-Over-CLNS Tunnel Scenario 2: ONS Node to Router

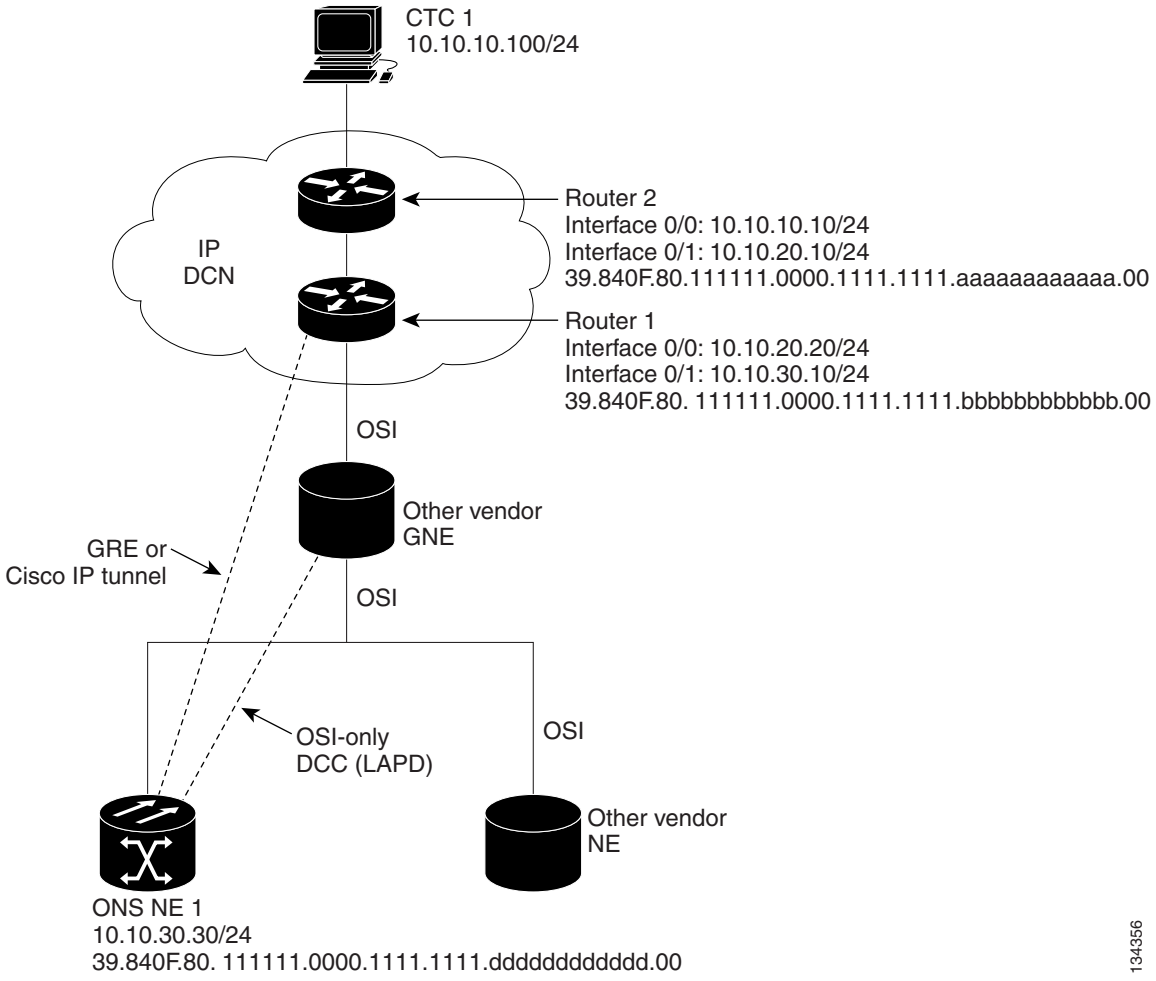

8.6.7.4 IP-Over-CLNS Tunnel Scenario 3: ONS Node to Router Across an OSI DCN

Figure 8-21 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) RS-DCC is created between the ONS NE 1 and the other vender 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 provisioning 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 vender GNE NSAP address
- Metric: 110
- Tunnel Type: Cisco IP

IP-over-OSI tunnel provisioning 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.dddddddddddd.00
interface Ethernet0/1
    clns router isis
router isis
    net 39.840F.80.1111.0000.1111.1111.aaaaaaaaaaaa.00
```

Figure 8-21 IP-Over-CLNS Tunnel Scenario 3: ONS Node to Router Across an OSI DCN

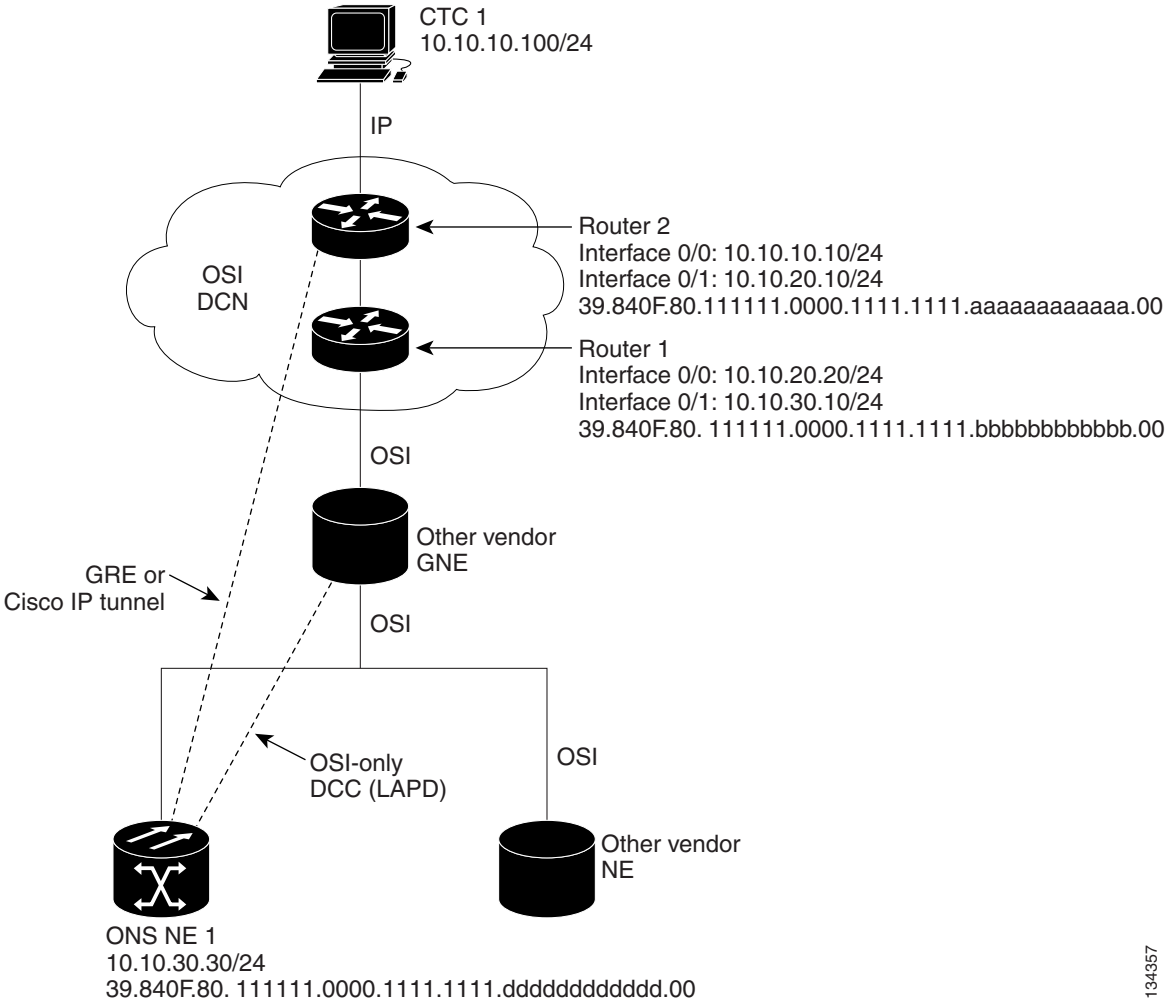

8.6.8 Provisioning OSI in CTC

Table 8-14 shows the OSI actions that can be performed in CTC using the node view Provisioning tab. Refer to the *Cisco ONS 15310-MA SDH Procedure Guide* for OSI procedures and tasks.

Table 8-14 OSI Actions from the CTC Node View Provisioning Tab

Tab	Actions
OSI > Main Setup	• View and edit Primary Area Address. • Change OSI routing mode. • Change LSP buffers.
OSI > TARP > Config	Configure the TARP parameters: • PDU L1/L2 propagation and origination. • TARP data cache and loop detection buffer. • LAN storm suppression. • Type 4 PDU on startup. • TARP timers: LDB, E1, T2, DS3/E3, 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	• Enable and disable routers. • Add, delete, and edit manual area addresses.
OSI > Routers > Subnets	Edit RS-DCC, MS-DCC, and LAN subnets that are provisioned for OSI.
OSI > Tunnels	Add, delete, and edit Cisco and IP-over-CLNS tunnels.
Comm Channels > RS-DCC	• Add OSI configuration to an RS-DCC. • Choose the data link layer protocol, PPP or LAP-D.
Comm Channels > MS-DCC	• Add OSI configuration to an RS-DCC.

Table 8-15 shows the OSI actions that can be performed in CTC using the node view Maintenance tab.

Table 8-15 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	• View the TARP data cache and identify static and dynamic entries. • Perform TID to NSAP resolutions. • Flush the TDC.

8.7 IPv6 Network Compatibility

IPv6 simplifies IP configuration and administration and has a larger address space than IPv4 to support the future growth of the Internet and Internet related technologies. It uses 128-bit addresses as against the 32-bit used in IPv4 addresses. Also, IPv6 gives more flexibility in designing newer addressing architectures.

Cisco ONS 15310 MA SDH can function in an IPv6 network when an Internet router that supports Network Address Translation-Protocol Translation (NAT-PT) is positioned between the GNE, such as an ONS 15310 MA SDH, and the client workstation. NAT-PT is a migration tool that helps users transition from IPv4 networks to IPv6 networks. NAT-PT is defined in RFC-2766. IPv4 and IPv6 nodes communicate with each other using NAT-PT by allowing both IPv6 and IPv4 stacks to interface between the IPv6 DCN and the IPv4 DCC networks.

8.8 IPv6 Native Support

Cisco ONS 15310 MA SDH Software R9.0 supports native IPv6. ONS 15310 MA SDH can be managed over IPv6 DCN networks by enabling the IPv6 feature. After you enable IPv6 in addition to IPv4, you can use CTC, TL1, and SNMP over an IPv6 DCN to manage ONS 15310 MA SDH. Each NE can be assigned an IPv6 address in addition to the IPv4 address. You can access the NE by entering the IPv4 address, an IPv6 address or the DNS name of the device. The IPv6 address is assigned only on the LAN interface of the NE. DCC/GCC interfaces use the IPv4 address.

By default, when IPv6 is enabled, the node processes both IPv4 and IPv6 packets on the LAN interface. If you want to process only IPv6 packets, you need to disable IPv4 on the node. Before you disable IPv4, ensure that IPv6 is enabled and the node is not in multishelf mode.

Figure 8-22 shows how an IPv6 DCN interacts with and IPv4 DCC.

Figure 8-22 IPv6-IPv4 Interaction

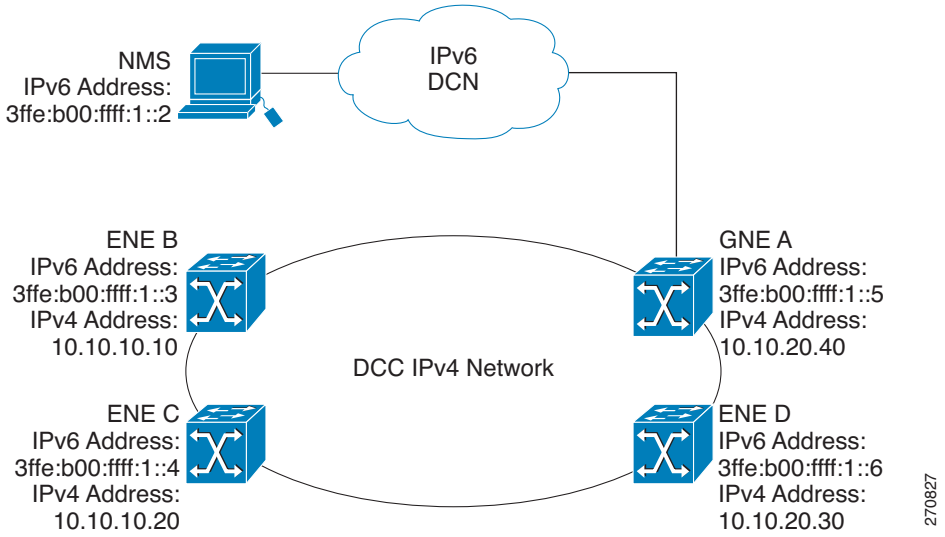

You can manage MSTP multishelf nodes over IPv6 DCN. RADIUS, FTP, SNTP, and other network applications support IPv6 DCN. To enable IPv6 addresses, you need to make the necessary configuration changes from the CTC or TL1 management interface. After you enable IPv6, you can start a CTC or TL1 session using the provisioned IPv6 address. The ports used for all IPv6 connections to the node are the same as the ports used for IPv4.

An NE can either be in IPv6 mode or IPv4 mode. In IPv4 mode, the LAN interface does not have an IPv6 address assigned to it. An NE, whether it is IPv4 or IPv6, has an IPv4 address and subnet mask. TCC2/TCC2P cards do not reboot automatically when you provision an IPv6 address, but a change in IPv4 address initiates a TCC2/TCC2P card reset. Table 8-16 describes the differences between an IPv4 node and an IPv6 node.

Table 8-16 Differences Between an IPv6 Node and an IPv4 Node

IPv6 Node	IPv4 Node
Has both IPv6 address and IPv4 address assigned to its craft Ethernet interface.	Does not have an IPv6 address assigned to its craft Ethernet interface.
The default router has an IPv6 address for IPv6 connectivity, and an IPv4 address for IPv4 connectivity.	The default router has an IPv4 address.
Cannot enable OSPF on LAN. Cannot change IPv4 NE to IPv6 NE if OSPF is enabled on the LAN.	Can enable OSPF on the LAN.
Cannot enable RIP on the LAN. Cannot change IPv4 NE to IPv6 NE if RIP is enabled on the LAN.	Can enable static routes/RIP on the LAN.
Not supported on static routes, proxy tunnels, and firewall tunnels.	Supported on static routes, proxy tunnels, and firewall tunnels.
Routing decisions are based on the default IPv6 router provisioned.	

8.8.1 IPv6 Enabled Mode

The default IP address configured on the node is IPv4. You can use either CTC or the TL1 management interface to enable IPv6. For more information about enabling IPv6 from the CTC interface, see the *Cisco ONS 15310-MA SDH Procedure Guide*. For more information about enabling IPv6 using TL1 commands, see the *Cisco ONS 15454 SDH, Cisco ONS 15600 SDH, and Cisco ONS 15310 MA SDH TL1 Command Guide*.

8.8.2 IPv6 Disabled Mode

You can disable IPv6 either from the CTC or from the TL1 management interface. For more information about disabling IPv6 from the CTC interface, see the *Cisco ONS 15310-MA SDH Procedure Guide*. For more information about disabling IPv6 using TL1 commands, see the *Cisco ONS 15454 SDH, Cisco ONS 15600 SDH, and Cisco ONS 15310 MA SDH TL1 Command Guide*.

8.8.3 IPv6 in Non-secure Mode

In non-secure mode, IPv6 is supported on the front and the rear Ethernet interfaces. You can start a CTC or TL1 session using the IPv6 address provisioned on the on the front and rear ports of the NE.

8.8.4 IPv6 in Secure Mode

In secure mode, IPv6 is only supported on the rear Ethernet interface. The front port only supports IPv4 even if it is disabled on the rear Ethernet interface. For more information about provisioning IPv6 addresses in secure mode, see the *Cisco ONS 15310-MA SDH Procedure Guide*.

8.8.5 IPv6 Limitations

IPv6 has the following configuration restrictions:

- You can provision an NE as IPv6 enabled only if the node is a SOCKS-enabled or firewall-enabled GNE/ENE.
- IPSec is not supported.
- OSPF/RIP cannot be enabled on the LAN interface if NE is provisioned as an IPv6 node.
- Static route/firewall/proxy tunnel provisioning is applicable only to IPv4 addresses even if the IPv6 is enabled.
- In secure mode, IPv6 is supported only on the rear Ethernet interface. IPv6 is not supported on the front port.
- ONS platforms use NAT-PT internally for providing IPv6 native support. NAT-PT uses the IPv4 address range 128.x.x.x for packet translation. Do not use the 128.x.x.x address range when you enable IPv6 feature.

8.9 FTP Support for ENE Database Backup

The Cisco ONS 15310-MA SDH provides FTP database backup and restore download to ENEs when proxy/firewall is enabled. This feature allows you to provision a list of legal FTP hosts in CTC, that can be used with TL1 commands to perform database backup/restore or software download. The FTP hosts can be provisioned to elapse after a specified time interval with the enable FTP relay function.

Once FTP host are provisioned, and FTP Releay is enabled, TL1 users can then use the COPY-RFILE command to perform database backup/restore or software download to and from this list of legal FTP hosts that are provisioned to ENEs. Also, TL1 supports TID to IP address translation for the GNE TID that is specified in the FTP URL of COPY-RFILE and COPY-IOSCFG commands.

Using the FTP Host provisioning feature in CTC and TL1 you can configure up to 12 valid FTP hosts.

ENEs are allowed access through the firewall according to the time configured in the FTP Relay Timer in CTC or TL1. The time interval is 1 to 60 minutes, and once the timer elapses, all FTP access to the FTP host is blocked again. A time of 0 disallows ENE access to FTP commands through the firewall.

When the firewall is not enabled (Proxy only), all FTP operations to the ENE will be allowed – software download, database backup/restore and IOS config file backup/restore. All FTP operations to the ENEs will be blocked when firewall is enabled.