



Manage Network Connectivity

This chapter provides an overview of ONS 15454 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 15454 nodes, and communication among networked ONS 15454 nodes. The chapter provides scenarios showing Cisco ONS 15454 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 15454 DCN communication is based on IP, ONS 15454 nodes can be networked to equipment that is based on the Open System Interconnection (OSI) protocol suites. This chapter also describes the ONS 15454 OSI implementation and provides scenarios that show how the ONS 15454 can be networked within a mixed IP and OSI environment.

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 15454 networking setup instructions, refer to the “Turn Up a Node” chapter of the *Cisco ONS 15454 DWDM Procedure Guide*.



Note

Unless otherwise specified, in this chapter “ONS 15454” refers to both ANSI and ETSI shelf assemblies.

Chapter topics include:

- [8.1 IP Networking Overview, page 8-2](#)
- [8.2 IP Addressing Scenarios, page 8-2](#)
- [8.3 Provisionable Patchcords, page 8-23](#)
- [8.4 Routing Table, page 8-24](#)
- [8.5 External Firewalls, page 8-26](#)
- [8.6 Open GNE, page 8-28](#)
- [8.7 TCP/IP and OSI Networking, page 8-30](#)
- [8.8 Link Management Protocol, page 8-34](#)



Note

To connect ONS 15454s 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 15454s 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 15454 node groups that allow you to provision non-data communication 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 15454 to serve as a gateway for ONS 15454s that are not connected to the LAN.
- Static routes can be created to enable connections among multiple Cisco Transport Controller (CTC) sessions with ONS 15454s that reside on the same subnet with multiple CTC sessions.
- ONS 15454s can be connected to Open Shortest Path First (OSPF) networks so ONS 15454 network information is automatically communicated across multiple LANs and WANs.
- The ONS 15454 proxy server can control the visibility and accessibility between CTC computers and ONS 15454 element nodes.

8.2 IP Addressing Scenarios

ONS 15454 IP addressing generally has eight 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 15454s in IP networks.


Table 8-1 General ONS 15454 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 15454s (backplane [ANSI] or MIC-C/T/P [ETSI] wire-wrap pins or RJ-45 port) and network hub/switch • Router ports and hub/switch ports
ONS 15454 hub/switch ports	If connectivity problems occur, set the hub or switch port that is connected to the ONS 15454 to 10 Mbps half-duplex.
Ping	Ping the node to test connections between computers and ONS 15454s.
IP addresses/subnet masks	Verify that ONS 15454 IP addresses and subnet masks are set up correctly.
Optical connectivity	Verify that ONS 15454 optical trunk ports are in service and that a DCC is enabled on each trunk port.

8.2.1 Scenario 1: CTC and ONS 15454s on Same Subnet

Scenario 1 shows a basic ONS 15454 LAN configuration ([Figure 8-1](#)). The ONS 15454s and CTC computer reside on the same subnet. All ONS 15454s connect to LAN A, and all ONS 15454s have DCC connections.

Figure 8-1 Scenario 1: CTC and ONS 15454s on Same Subnet (ANSI and ETSI)




8.2.2 Scenario 2: CTC and ONS 15454s 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 15454s 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). The routers each have a subnet mask of 255.255.255.0.

On the CTC computer, the default gateway is set to router interface A. If the LAN uses DHCP (Dynamic Host Configuration Protocol), the default gateway and IP address are assigned automatically. In the Figure 8-2 example, a DHCP server is not available.

Figure 8-2 Scenario 2: CTC and ONS 15454s Connected to Router (ANSI and ETSI)



8.2.3 Scenario 3: Using Proxy ARP to Enable an ONS 15454 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 15454 to respond to the ARP request for ONS 15454s not connected to the LAN. (ONS 15454 proxy ARP requires no user configuration.) For this to occur, the DCC-connected ONS 15454s must reside on the same subnet as the LAN-connected (gateway) ONS 15454. When a LAN device sends an ARP request to an ONS 15454 that is not connected to the LAN, the gateway ONS 15454 (the one connected to the LAN) returns its MAC address to the LAN

device. The LAN device then sends the datagram for the remote ONS 15454 to the MAC address of the proxy ONS 15454. The proxy ONS 15454 uses its routing table to forward the datagram to the non-LAN ONS 15454.

Scenario 3 is similar to Scenario 1, but only one ONS 15454 (Node 1) connects to the LAN ([Figure 8-3](#)). Two ONS 15454s (Node 2 and Node 3) connect to ONS 15454 Node 1 through the section DCC. Because all three ONS 15454s are on the same subnet, proxy ARP enables ONS 15454 Node 1 to serve as a gateway for ONS 15345 Node 2 and Node 3.


**Note**

This scenario assumes all CTC connections are to Node 1. If you connect a laptop to either ONS 15454 Node 2 or Node 3, network partitioning occurs; neither the laptop or the CTC computer can see all nodes. If you want laptops to connect directly to end network elements, you must create static routes (see [Scenario 5](#)) or enable the ONS 15454 proxy server (see [Scenario 7](#)).

Be aware that:

- GNE and ENE 15454 proxy ARP is disabled.
- There is exactly one proxy ARP server on any given Ethernet segment; however, there may be more than one server in an ANSI or ETSI topology.
- The proxy ARP server does not perform the proxy ARP function for any node or host that is on the same Ethernet segment.
- It is important in [Figure 8-3](#) that the CTC workstation be located within the same subnet and on the same Ethernet segment as the proxy ARP server.

Figure 8-3 Scenario 3: Using Proxy ARP (ANSI and ETSI)




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 (ANSI and ETSI)



8.2.4 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 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 (ANSI and ETSI)




8.2.5 Scenario 5: Using Static Routes to Connect to LANs

Static routes are used for two purposes:

- To connect ONS 15454s to CTC sessions on one subnet connected by a router to ONS 15454s 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 15454s 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 15454s 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 (ANSI and ETSI)




The destination and subnet mask entries control access to the ONS 15454s:

- 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 (ANSI and ETSI)



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.

ONS 15454s use the OSPF protocol in internal ONS 15454 networks for node discovery, circuit routing, and node management. You can enable OSPF on the ONS 15454s so that the ONS 15454 topology is sent to OSPF routers on a LAN. Advertising the ONS 15454 network topology to LAN routers

eliminates the need to manually enter static routes for ONS 15454 subnetworks. Figure 8-8 shows a network enabled for OSPF. 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.

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 15454 OSPF topology for advertising to an OSPF network, you must assign an OSPF area ID in decimal format to the ONS 15454 network. An area ID is a “dotted quad” value that appears similar to an IP address. Coordinate the area ID number assignment with your LAN administrator. All DCC-connected ONS 15454s should be assigned the same OSPF area ID.

**Note**

It is recommended that the number of 15454s in an OSPF area be limited, because this allows faster loading into a CTC and is less likely to incur any problems.

Figure 8-8 Scenario 6: OSPF Enabled (ANSI and ETSI)





Figure 8-9 Scenario 6: OSPF Not Enabled (ANSI and ETSI)



8.2.7 Scenario 7: Provisioning the ONS 15454 Proxy Server

The ONS 15454 proxy server is a set of functions that allows you to network ONS 15454s in environments where visibility and accessibility between ONS 15454s 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 ONS 15454s while preventing the field technicians from accessing the NOC LAN. To do this, one ONS 15454 is provisioned as a GNE and the other ONS 15454s are provisioned as end network elements (ENEs). The GNE ONS 15454 tunnels connections between CTC computers and ENE ONS 15454s, providing management capability while preventing access for non-ONS 15454 management purposes.

The ONS 15454 gateway setting 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 8-3 on page 8-16](#) and [Table 8-4 on page 8-17](#)) depend on whether the packet arrives at the ONS 15454 DCC or TCC2/TCC2P Ethernet interface.
- Processes Simple Network Time Protocol (SNTP) and Network Time Protocol (NTP) requests. ONS 15454 ENEs can derive time-of-day from an SNTP/NTP LAN server through the GNE ONS 15454.
- Processes Simple Network Management Protocol version 1 (SNMPv1) traps. The GNE ONS 15454 receives SNMPv1 traps from the ENE ONS 15454s and forwards or relays the traps to SNMPv1 trap destinations or ONS 15454 SNMP relay nodes.

The ONS 15454 proxy server is provisioned using the Enable proxy server on port check box on the Provisioning > Network > General tab. If checked, the ONS 15454 serves as a proxy for connections between CTC clients and ONS 15454s that are DCC-connected to the proxy ONS 15454. 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. In addition, you can set the proxy server as an ENE or a GNE:

- End Network Element (ENE)—If set as an ENE, the ONS 15454 neither installs nor advertises default or static routes that go through its Ethernet port. However, an ENE does install and advertise routes that go through the DCC. CTC computers can communicate with the ONS 15454 using the TCC2/TCC2P craft port, but they cannot communicate directly with any other DCC-connected ONS 15454.

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 15454 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, firewall is not enabled. CTC can communicate with any other DCC-connected ONS 15454s.



Note

If you launch CTC against a node through a Network Address Translation (NAT) or Port Address Translation (PAT) 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 8-10](#) shows an ONS 15454 proxy server implementation. A GNE ONS 15454 is connected to a central office LAN and to ENE ONS 15454s. 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 15454 ENEs. However, the craft technicians must be prevented from accessing or seeing the NOC or central office LANs.

In the example, the ONS 15454 GNE is assigned an IP address within the central office LAN and is physically connected to the LAN through its LAN port. ONS 15454 ENEs are assigned IP addresses that are outside the central office LAN and given private network IP addresses. If the ONS 15454 ENEs are collocated, the craft LAN ports could be connected to a hub. However, the hub should have no other network connections.

Figure 8-10 Scenario 7: ONS 15454 Proxy Server with GNE and ENEs on the Same Subnet (ANSI and ETSI)

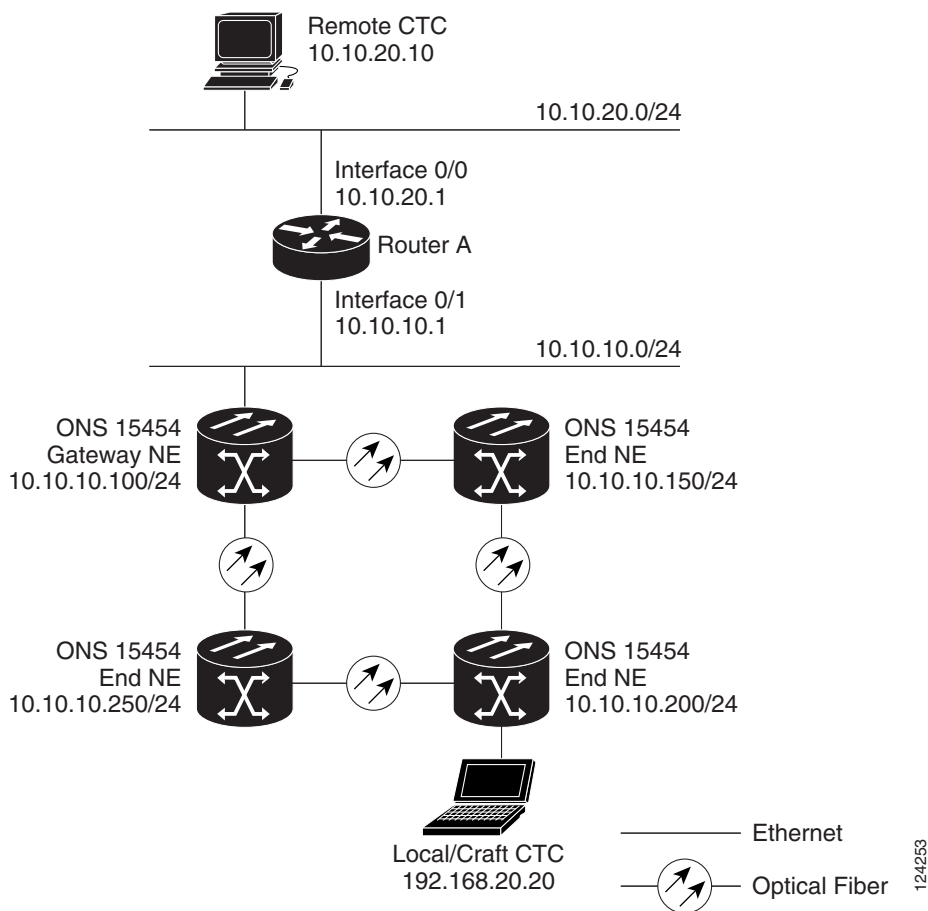


Table 8-2 shows recommended settings for ONS 15454 GNEs and ENEs in the configuration shown in Figure 8-10.

Table 8-2 ONS 15454 Gateway and End NE Settings

Setting	ONS 15454 Gateway NE	ONS 15454 End NE
OSPF	Off	Off
SNTP server (if used)	SNTP server IP address	Set to ONS 15454 GNE IP address
SNMP (if used)	SNMPv1 trap destinations	Set SNMPv1 trap destinations to ONS 15454 GNE, port 391

Figure 8-11 shows the same proxy server implementation with ONS 15454 ENEs on different subnets. The ONS 15454 GNEs and ENEs are provisioned with the settings shown in Table 8-2.

Figure 8-11 Scenario 7: ONS 15454 Proxy Server with GNE and ENEs on Different Subnets (ANSI and ETSI)

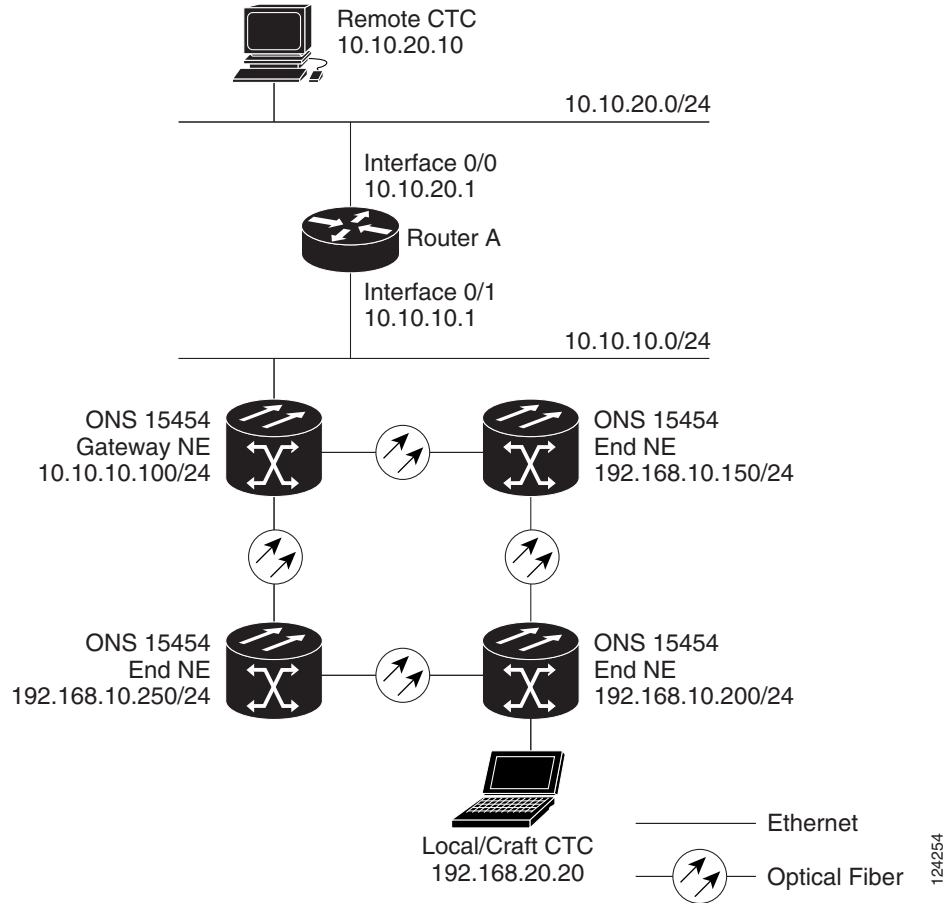


Figure 8-12 shows the same proxy server implementation with ONS 15454 ENEs in multiple rings.

Figure 8-12 Scenario 7: ONS 15454 Proxy Server With ENEs on Multiple Rings (ANSI and ETSI)

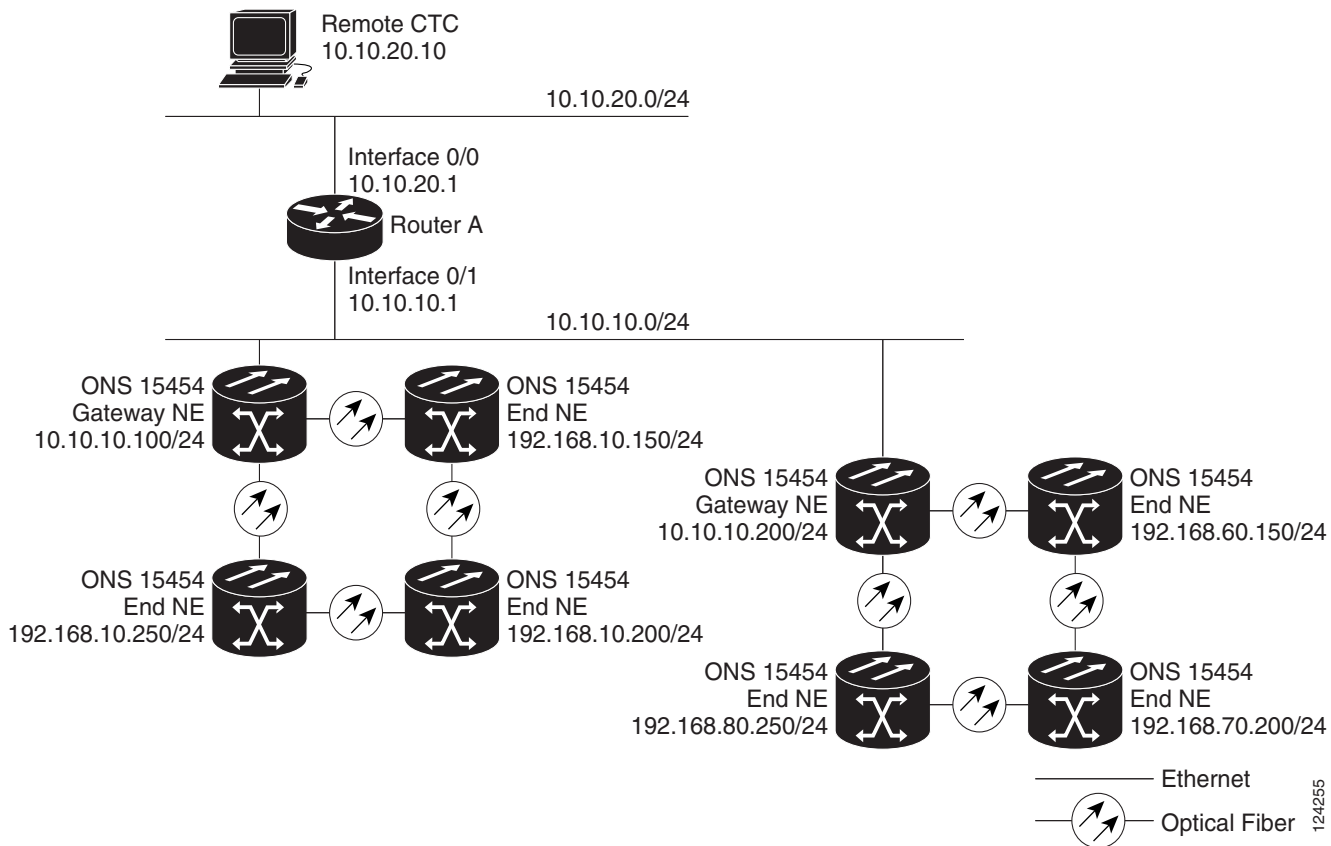


Table 8-3 shows the rules the ONS 15454 follows to filter packets for the firewall when nodes are configured as ENEs and GNEs. If the packet is addressed to the ONS 15454, additional rules (shown in Table 8-4) are applied. Rejected packets are silently discarded.

Table 8-3 Proxy Server Firewall Filtering Rules

Packets Arriving At:	Are Accepted if the Destination IP Address is:
TCC2/TCC2P Ethernet interface	<ul style="list-style-type: none"> The ONS 15454 itself The ONS 15454'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	<ul style="list-style-type: none"> The ONS 15454 itself Any destination connected through another DCC interface Within the 224.0.0.0/8 network

Table 8-4 Proxy Server Firewall Filtering Rules When Packet Addressed to ONS 15454

Packets Arriving At	Rejects
TCC2/TCC2P Ethernet interface	<ul style="list-style-type: none"> UDP (User Datagram Protocol) packets addressed to the SNMP trap relay port (391)
DCC interface	<ul style="list-style-type: none"> TCP (Transmission Control Protocol) packets addressed to the proxy server port (1080)

If you implement the proxy server, note that all DCC-connected ONS 15454s 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:

- Disconnect the craft computer from the unreachable ONS 15454. Connect to the ONS 15454 through another network ONS 15454 that has a DCC connection to the unreachable ONS 15454.
- Disconnect all DCCs to the node by disabling them on neighboring nodes. Connect a CTC computer directly to the ONS 15454 and change its provisioning.

8.2.8 Scenario 8: Dual GNEs on a Subnet

The ONS 15454 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 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, both of which enhance CTC performance.

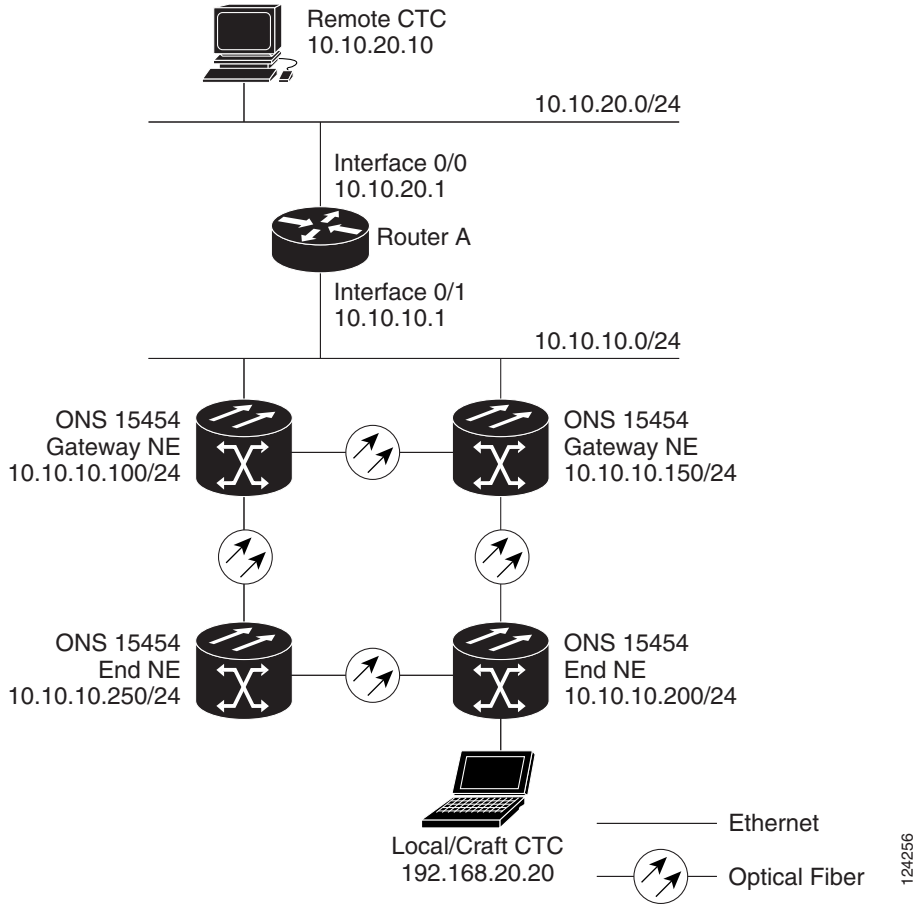


Note

Dual GNEs do not need special provisioning

Figure 8-13 shows a network with dual GNEs on the same subnet.

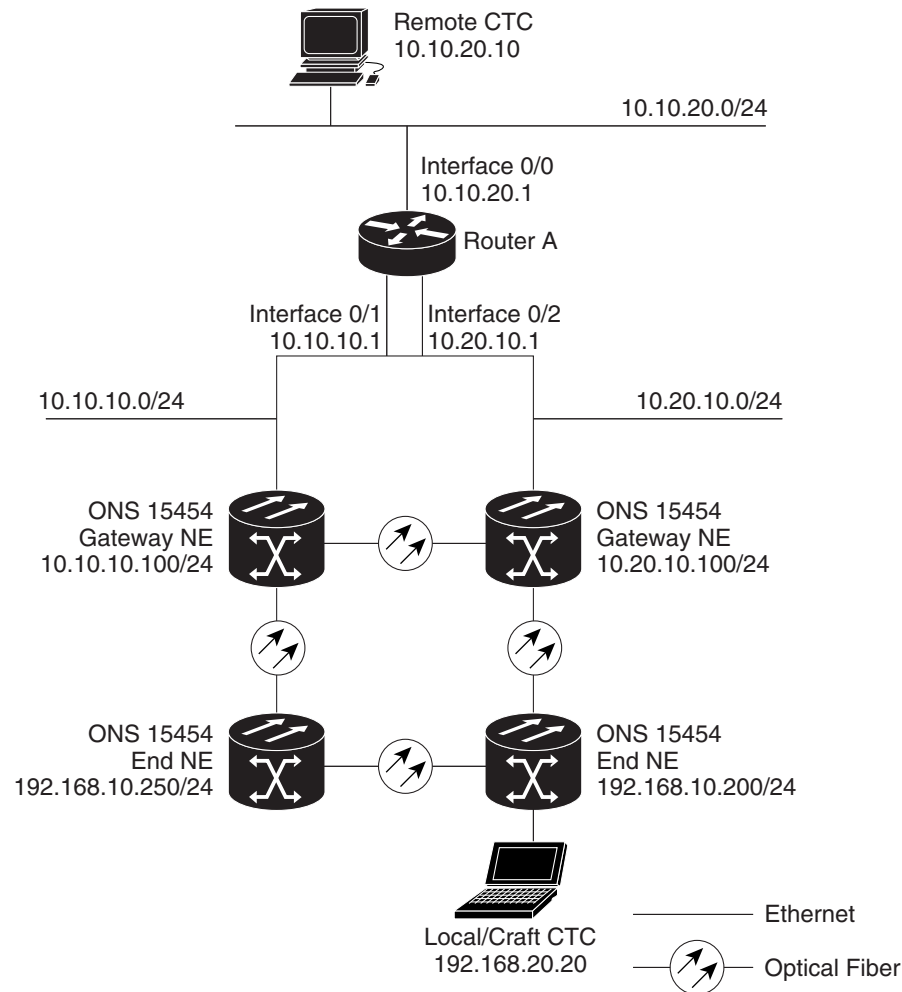
Figure 8-13 Scenario 8: Dual GNEs on the Same Subnet (ANSI and ETSI)



124/256

Figure 8-14 shows a network with dual GNEs on different subnets.

Figure 8-14 Scenario 8: Dual GNEs on Different Subnets (ANSI and ETSI)



8.2.9 Scenario 9: IP Addressing with Secure Mode Enabled

The TCC2 card and TCC2P card both default to nonsecure mode. In this mode, the front and back Ethernet (LAN) ports share a single MAC address and IP address. TCC2P cards allow you to place a node in secure mode, which prevents a front-access craft port user from accessing the LAN through the backplane port. Secure mode can be locked, which prevents the mode from being altered. To place a node in secure mode refer to the “DLP -G264 Enable Node Security Mode” task in the “Turn Up a Node” chapter of the *Cisco ONS 15454 DWDM Procedure Guide*. To lock secure node, refer to the “DLP-G265 Lock Node Security” task in the “Manage the Node” chapter of the *Cisco ONS 15454 DWDM Procedure Guide*.

8.2.9.1 Secure Mode Behavior

Changing a TCC2P node from nonsecure mode to secure mode allows you to provision two Ethernet addresses for the ONS 15454 and causes the node to assign the ports different MAC addresses. In secure mode, one IP address is provisioned for the ONS 15454 backplane LAN (Ethernet) port, and the other IP address is provisioned for the TCC2P Ethernet port. Both addresses reside on different subnets, providing an additional layer of separation between the craft access port and the ONS 15454 LAN. If secure mode is enabled, the IP addresses provisioned for both TCC2P Ethernet ports must follow general IP addressing guidelines and must reside on different subnets from each other and the default router IP address.

In secure mode, the IP address assigned to the front LAN (Ethernet) port becomes a private address, while the backplane connects the node to an Operations Support System (OSS) through a central office LAN or private enterprise network. A superuser can configure the node to hide or reveal the backplane's LAN IP address in CTC, the routing table, or autonomous message reports.

In nonsecure mode, a node can be a GNE or ENE. Placing the node into secure mode automatically turns on SOCKS proxy and defaults the node to GNE status. However, the node can be changed back to an ENE. In nonsecure mode, an ENE's SOCKS proxy can be disabled—effectively isolating the node beyond the LAN firewall—but it cannot be disabled in secure mode. To change a node's GNE or ENE status and disable the SOCKS proxy, refer to the “DLP-G56 Provision IP Settings” task in the “Turn Up a Node” chapter of the *Cisco ONS 15454 DWDM Procedure Guide*.



Caution

Enabling secure mode causes the TCC2P card to reboot; a TCC2P card reboot affects traffic.



Caution

The TCC2 card fails to boot when it is added as a standby card to a node containing an active TCC2P card configured in the secure mode.



Note

If both front and backplane access ports are disabled in an ENE and the node is isolated from DCC communication (due to user provisioning or network faults), the front and backplane ports are automatically reenabled.

Figure 8-15 shows an example of secure mode ONS 15454 nodes with front-access Ethernet port addresses that reside on the same subnet.

Figure 8-15 Scenario 9: ONS 15454 GNE and ENEs on the Same Subnet with Secure Mode Enabled

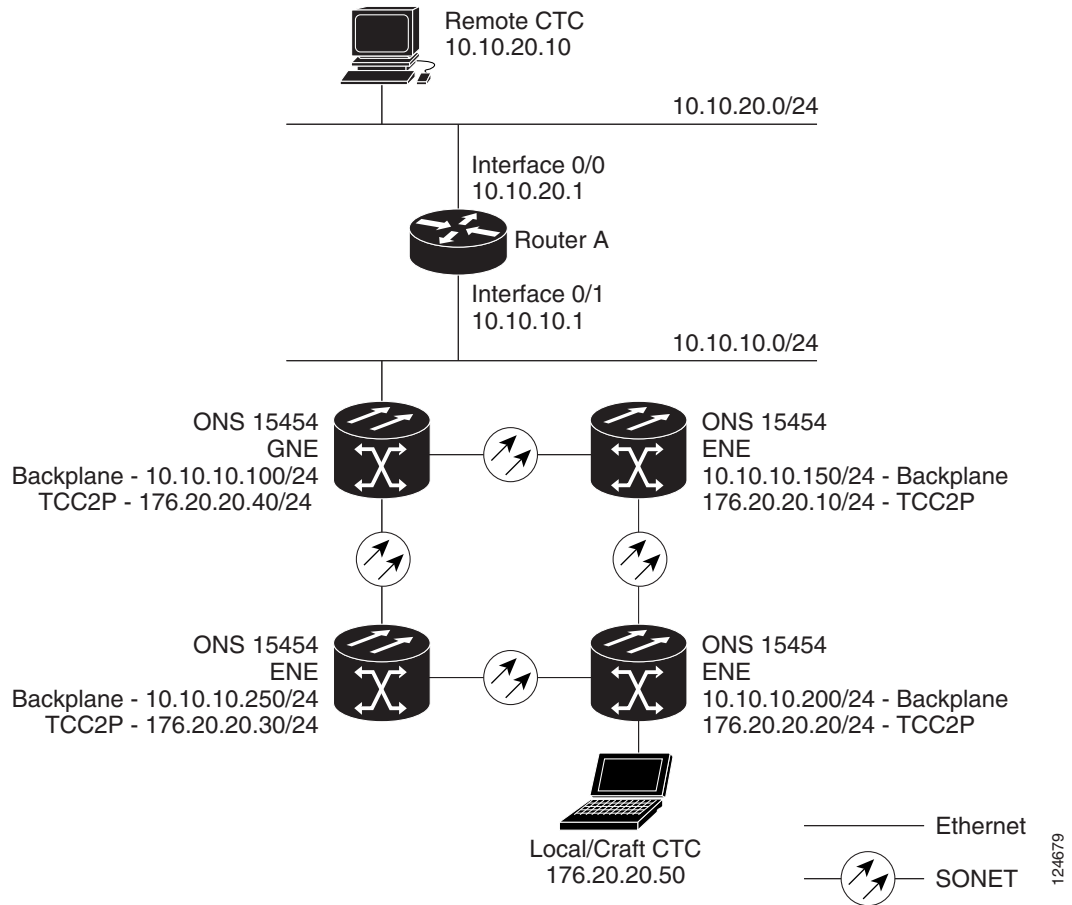



Figure 8-16 shows an example of ONS 15454 nodes connected to a router with secure mode enabled. In each example, the node's TCC2P port address (node address) resides on a different subnet from the node backplane addresses.

Figure 8-16 Scenario 9: ONS 15454 GNE and ENEs on Different Subnets with Secure Mode Enabled



8.2.9.2 Secure Node Locked and Unlocked Behavior

Secure mode can operate on a node in either locked or unlocked mode. By default, secure mode's status is unlocked; only a superuser can convert it to locked mode. Doing so permanently changes the hardware configuration on the active and standby TCC2P cards as well as the chassis.

Locked mode must be used carefully because the cards and shelf retain their locked status even if separated from each other. For example, if a node is in secure, locked mode and you perform a card pull on its standby TCC2P, then insert that as the active card into another node, the secure, locked mode is written to the new node's chassis and standby TCC2P. If you perform a card pull on a secure, locked node's active and standby TCC2Ps and insert both of them into a chassis that previously was in unlocked mode, the node becomes locked.

When it is secure and locked, a node's configuration, Ethernet port status, its secure mode, and the locked status cannot be changed by any network user—including a superuser. To have a secure node's lock removed, contact Cisco Technical Support to arrange a Return Material Authorization (RMA) for the chassis and for the TCC2Ps. Refer to the [“Obtaining Documentation and Submitting a Service Request”](#) section on page [xlvii](#) as needed.

**Caution**

It is necessary for the TCC2Ps and the chassis to be unlocked together. If only one component (such as the shelf) is unlocked, the system will return to locked mode.

8.3 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 in the following situations:

- An optical port is connected to a transponder or muxponder client port provisioned in transparent mode.
- An optical ITU port is connected to a DWDM optical channel card.
- Two transponder or muxponder trunk ports are connected to a DWDM optical channel card and the generic control channel (GCC) is carried transparently through the ring.
- Transponder or muxponder client and trunk ports are in a regenerator group, the cards are in transparent mode, and DCC/GCC termination is not available.

Provisionable patchcords are required on both ends of a physical link. The provisioning at each end includes a local patchcord ID, slot/port information, remote IP address, and remote patchcord ID. Patchcords appear as dashed lines in CTC network view.

An optical patchcord must be provisioned between an OCH filter and an OCH trunk port. A manually provisioned patchcord automatically tunes the transponder (TXP) or muxponder (MXP) trunk as an OCH filter if the TXP or MXP is set to autoprovision at the first tunable wavelength. You can automatically tune internal and external (virtual link) patchcords in CTC. In TL1, only internal patchcords can be provisioned.

Table 8-5 lists the supported card combinations for client and trunk ports in a provisionable patchcord.

Table 8-5 Cisco ONS 15454 Client/Trunk Card Combinations for Provisionable Patchcords

Trunk Cards	Client Cards						
	MXP_2.5G_10G/ TXP_MR_10G	TXP_MR_2.5G/ TXPP_MR_2.5G	MXP_2.5G_10E/ TXP_MR_10E	32MUX-0 32DMX-0	32WSS/ 32DMX	AD-xC-xx.x	4MD-xx.x
MXP_2.5G_10G/ TXP_MR_10G	—	—	—	Yes	Yes	Yes	Yes
TXP_MR_2.5G/ TXPP_MR_2.5G	—	—	—	Yes	Yes	Yes	Yes
MXP_2.5G_10E/ TXP_MR_10E	—	—	—	Yes	Yes	Yes	Yes
MXP_MR_2.5G/ MXPP_MR_2.5G	—	—	—	Yes	Yes	Yes	Yes
OC-192	Yes	—	Yes	—	—	—	—
OC-48	Yes	Yes	Yes	—	—	—	—
OC-192 ITU	—	—	—	Yes	Yes	Yes	Yes
OC-48 ITU	—	—	—	Yes	Yes	Yes	Yes

Table 8-6 lists the supported card combinations for client-to-client ports in a patchcord.

Table 8-6 Cisco ONS 15454 Client/Client Card Combinations for Provisionable Patchcords

Client Cards	MXP_2.5G_10G/ TXP_MR_10G	TXP_MR_2.5G/ TXPP_MR_2.5G	MXP_2.5G_10E/ TXP_MR_10E
MXP_2.5G_10G/TXP_MR_10G	Yes	—	Yes
TXP_MR_2.5G/TXPP_MR_2.5G	—	Yes	—
MXP_2.5G_10E/TXP_MR_10E	Yes	—	Yes

Table 8-7 lists the supported card combinations for trunk-to-trunk ports in a patchcord.

Table 8-7 Cisco ONS 15454 Trunk/Trunk Card Combinations for Provisionable Patchcords

Trunk Cards	MXP_2.5G_10G/ TXP_MR_10G	TXP_MR_2.5G/ TXPP_MR_2.5G	MXP_2.5G_10E/ TXP_MR_10E
MXP_2.5G_10G/TXP_MR_10G	Yes	—	Yes
TXP_MR_2.5G/TXPP_MR_2.5G	—	Yes	—
MXP_2.5G_10E/TXP_MR_10E	Yes	—	Yes

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

- An optical port connected to transponder/muxponder port or add/drop multiplexer or multiplexer/demultiplexer port requires section DCC/line DCC (SDCC/LDCC or RS-DCC/MS-DCC) termination.
- If the optical port is the protection port in a 1+1 group, the working port must have SDCC/LDCC or RS-DCC/MS-DCC termination provisioned.
- If the remote end of a patchcord is Y-cable protected or is an add/drop multiplexer or multiplexer/demultiplexer port, an optical port requires two patchcords.

Transponder and muxponder ports have the following requirements when used in a provisionable patchcord:

- Two patchcords are required when a transponder/muxponder port is connected to an add/drop multiplexer or multiplexer/demultiplexer port. CTC automatically prompts the user to set up the second patchcord.
- If a patchcord is on a client port in a regenerator group, the other end of the patchcord must be on the same node and on a port within the same regenerator group.
- A patchcord is allowed on a client port only if the card is in transparent mode.

DWDM cards support provisionable patchcords only on optical channel ports. Each DWDM optical channel port can have only one provisionable patchcord.

8.4 Routing Table

ONS 15454 routing information is displayed 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 15454 interface used to access the destination. Values are:
 - motfcc0—The ONS 15454 Ethernet interface, that is, the RJ-45 jack on the TCC2/TCC2P and, for ANSI shelves, the LAN 1 pins on the backplane or, for ETSI shelves, the LAN connection on the MIC-C/T/P.
 - pdcc0—An SDCC or RS-DCC interface, that is, an OC-N trunk card identified as the SDCC or RS-DCC termination.
 - lo0—A loopback interface.

Table 8-8 shows sample routing entries for an ONS 15454.

Table 8-8 Sample Routing Table Entries

Entry	Destination	Mask	Gateway	Usage	Interface
1	0.0.0.0	0.0.0.0	172.20.214.1	265103	motfcc0
2	172.20.214.0	255.255.255.0	172.20.214.92	0	motfcc0
3	172.20.214.92	255.255.255.255	127.0.0.1	54	lo0
4	172.20.214.93	255.255.255.255	0.0.0.0	16853	pdcc0
5	172.20.214.94	255.255.255.255	172.20.214.93	16853	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 are 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 (motfcc0) indicates that the ONS 15454 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 (motfcc0) indicates that the ONS 15454 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 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 a DCC interface is used to reach the gateway.

8.5 External Firewalls

This section provides information on firewall ports required for various type of connections that are established with the NE (controller card). Also, there are examples of Access Control List (ACL) for external firewall configuration that makes a connection feasible with the controller card.

8.5.1 Firewall Ports

Table 8-9 lists the ports that must be enabled to establish a communication channel with the NE (controller card).

Table 8-9 Firewall Ports for Various Sessions

Session Type	Session Description	Mode	Port Number	Firewall ACL
CORBA	CORBA listener port on the NE	Standard	57790 (default); user configurable to the standard port 683 or any other port. ¹	Inbound
		Secure	57791 (default); user configurable to the standard port 684 or any other port.	
	Standard Internet Inter-ORB Protocol (IIOP) listener port on machine running CTC	Standard	Dynamic (default); user configurable to the standard port 683 or any other port. ²	Outbound
		Secure	Dynamic (default); user configurable to the standard port 684 or any other port.	
SOCKS	CTC configured with SOCKS or GNE	—	1080	Inbound
HTTP	HTTP port on the NE	—	80	Inbound
HTTPS	HTTPS port on the NE	—	433 ³	Inbound
TL1	TL1 port on NE	Standard	3082, 3083, 2362	Inbound
		Secure	4083	

Table 8-9 Firewall Ports for Various Sessions

Session Type	Session Description	Mode	Port Number	Firewall ACL
SNMP	SNMP listener port on NE	Standard	161	Inbound
		Secure		
	SNMP trap listener port on the machine receiving the traps	Standard	162 (default); user configurable to any port between 1024 to 65535	Outbound
		Secure		

1. To configure the port, see “DLP-G61 Provision the IOP Listener Port on the ONS 15454” in the *Cisco ONS 15454 DWDM Procedure Guide*.
2. To configure the port, see “DLP-G62 Provision the IOP Listener Port on the CTC Computer” in the *Cisco ONS 15454 DWDM Procedure Guide*.
3. If this port is blocked, NE could take long time to initialize.

8.5.2 ACL Examples

The following access control list (ACL) example shows a firewall configuration when the proxy server gateway setting is not enabled. In the example, the CTC workstation's address is 192.168.10.10, and the ONS 15454 address is 10.10.10.100. The firewall is attached to the GNE; hence, inbound is from the CTC to the GNE and outbound is from the GNE to CTC. The CTC Common Object Request Broker Architecture (CORBA) standard port is 683 and the TCC CORBA default port on TCC is 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 15454 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 15454 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 15454 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 15454 (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 15454 GNE to CTC ***
```

The following ACL example shows 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 15454 address is 10.10.10.100. The firewall is attached to the GNE; hence, inbound is from the CTC to the GNE and outbound is from the GNE to CTC. The CTC Common Object Request Broker Architecture (CORBA) standard port is 683 and the TCC CORBA default port on TCC is 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 15454 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 15454 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 15454 GNE to CTC ***

```

8.6 Open GNE

The ONS 15454 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 a GCC-based network to function as an IP network for non-ONS nodes.

To configure an open GNE network, you can provision GCC 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 GCC creation. The default 0.0.0.0 IP address allows the far-end, non-ONS node to identify itself with any 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 GCC 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 GCC 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. A proxy connection is allowed if the CTC client is in a source subnet and the requested destination is in the destination subnet. Firewall tunnels allow IP traffic to route between the node Ethernet and pdcc interfaces. An inbound Ethernet packet is allowed through the firewall if its source address matches a tunnel source and its destination matches a tunnel destination. An inbound pdcc packet is allowed through the firewall if its source address matches a tunnel destination and its destination address matches a tunnel source. Tunnels only affect TCP and UDP packets.

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 8-17 shows an example of a foreign node connected to the GCC 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 8-17 Proxy and Firewall Tunnels for Foreign Terminations

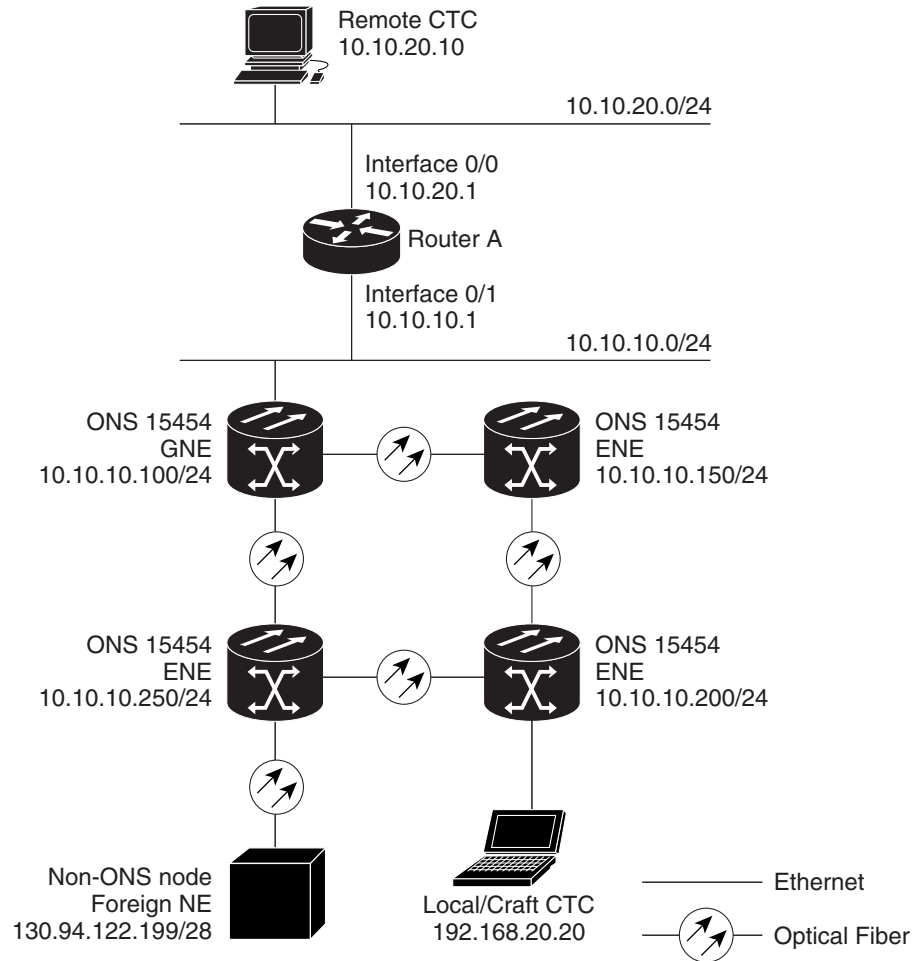
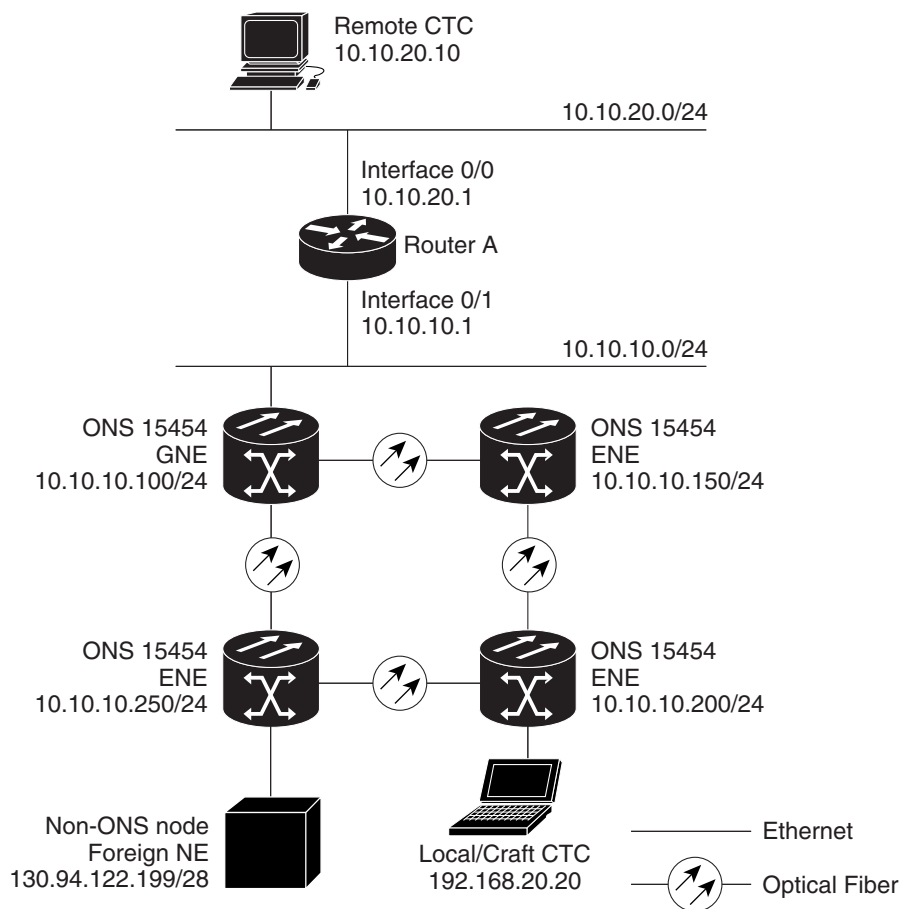


Figure 8-18 shows a remote node connected to an ENE Ethernet port. Proxy and firewall tunnels are useful in this example because the GNE would otherwise block IP access between the PC and foreign node. This configuration also requires a firewall tunnel on the ENE.

12/4261

Figure 8-18 Foreign Node Connection to an ENE Ethernet Port



8.7 TCP/IP and OSI Networking

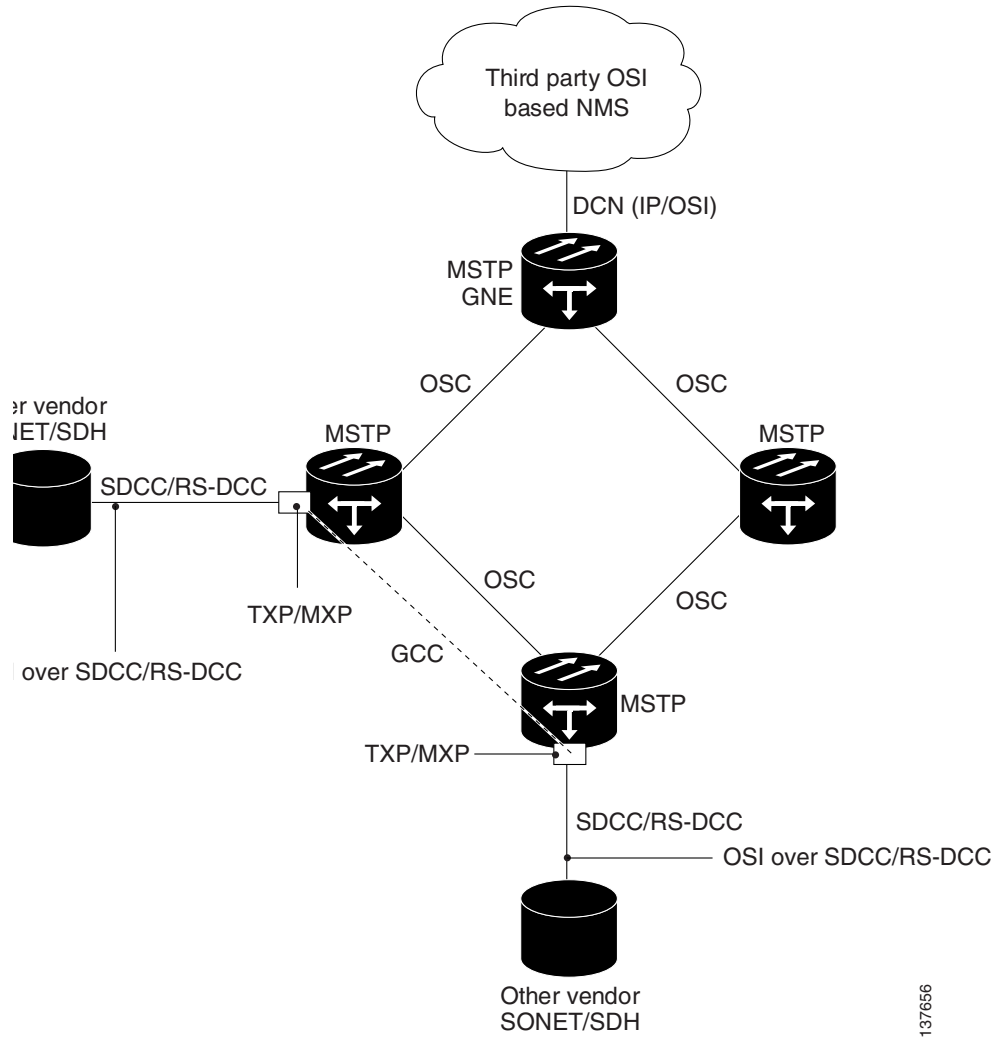
ONS 15454 DCN communication is based on the TCP/IP protocol suite. However, ONS 15454s can also be networked with equipment that uses the OSI protocol suite. While TCP/IP and OSI protocols are not directly compatible, they do have the same objectives and occupy similar layers of the OSI reference model. For detailed information about OSI protocols, processes, and scenarios, refer to the “Management Network Connectivity” chapter in the *ONS 15454 Reference Manual*. OSI/MSTP scenarios are provided in the following sections.

In OSI/MSTP Scenario 1 (Figure 8-19), an SDCC or RS-DCC carries an OC-N signal from an OSI-based third party NE to a TXP/MXP card on an ONS NE. It is carried by GCC to a TXP/MXP card on another MSTP NE and then by SDCC or RS-DCC to a second third party NE. This scenario requires TXPs/MXPs whose client interfaces can be provisioned in section or line termination mode. These include:

- TXP_MR_2.5 / TXPP_MR_2.5 (when equipped with OCn-N SFPs)
- TXP_MR_10G / TXP_MR_10E (when the client is configured as OC192)
- MXP_2.5_10G and MXP_2.5_10E

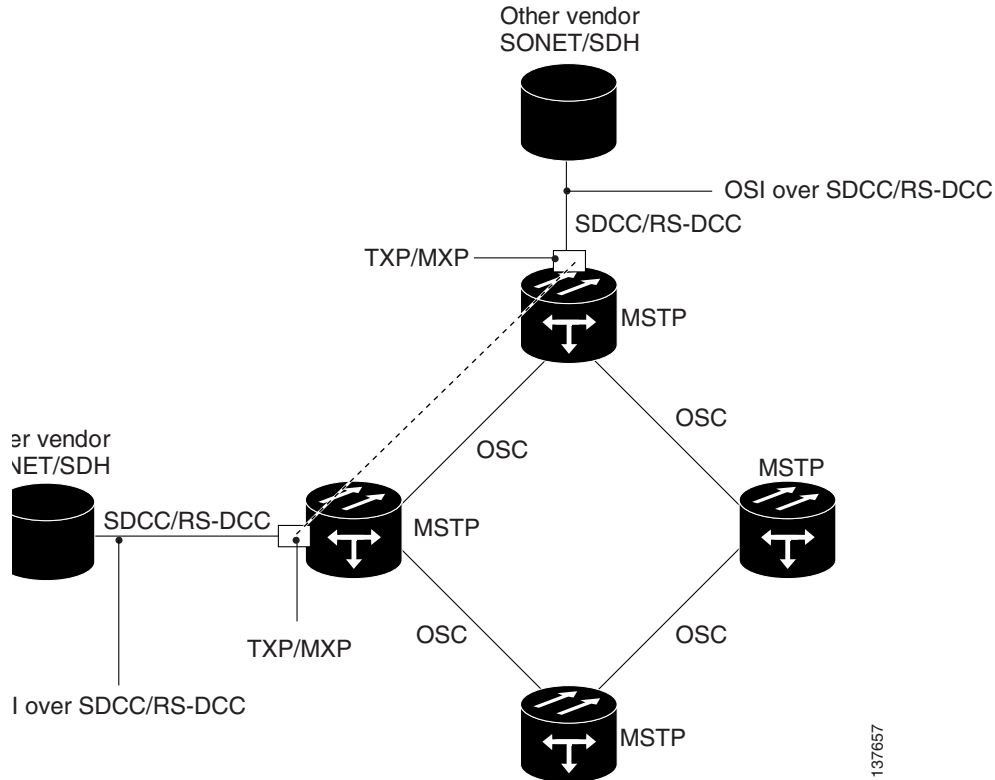
OSI has to be carried or tunnelled to the other TXP/MXP through an OSC termination, GCC termination, or both. The third party NMS has OSI connectivity to its NEs with the MSTP ONS NE serving as the GNE for third party vendor OSI-based SONET equipment.

Figure 8-19 OSI/MSTP Scenario 1



OSI/MSTP Scenario 2 (Figure 8-20) is similar to Scenario 1, except the MSTP NEs do not have connectivity to an OSI NMS.


Figure 8-20 OSI/MSTP Scenario 2



OSI/MSTP Scenario 3 (Figure 8-21) shows the following:

- OSI is carried over an SDCC or RS-DCC termination.
- OSI has to be carried or tunneled to the other peer TXP/MXP through an OSC termination, GCC termination, or both
- An OSS has IP connectivity to all the NEs
- The MSTP NE is a GNE for the third party OSI-based SONET NEs. the MSTP NEs perform all mediation functions.

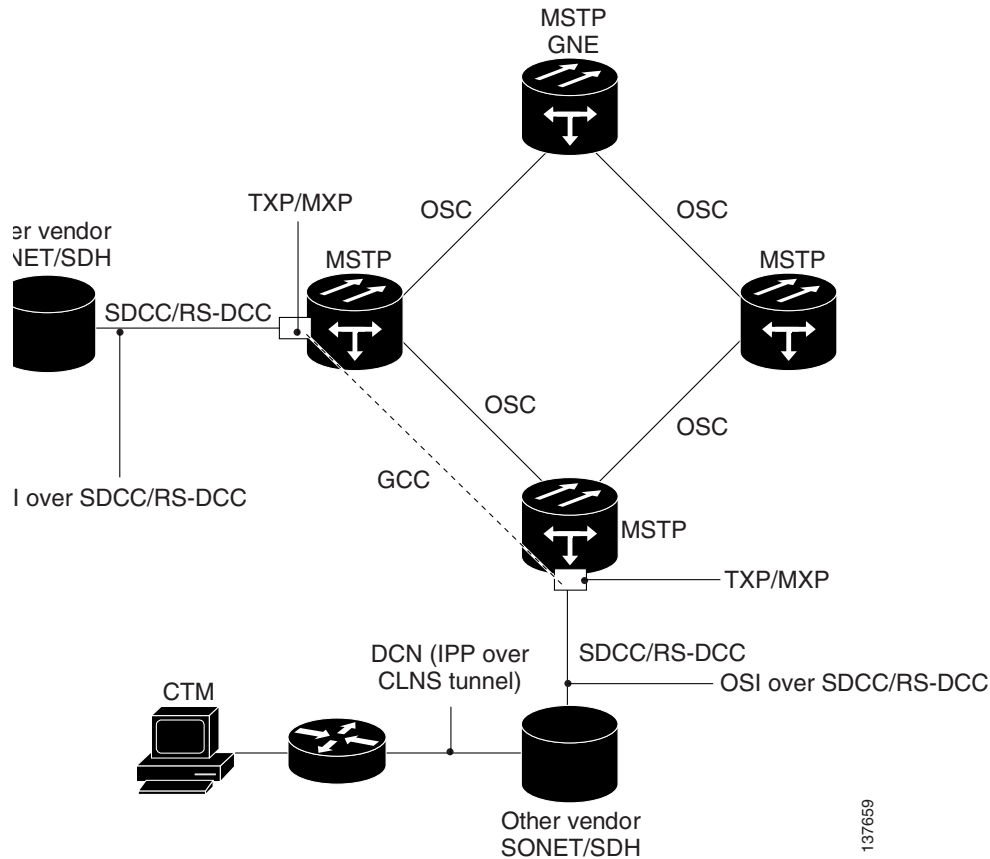
Figure 8-21 OSI/MSTP Scenario 3



OSI/MSTP Scenario 4 (Figure 8-22) shows the following:

- OSI is carried over an SDCC or RS-DCC termination.
- OSI has to be carried or tunneled to the other peer TXP/MXP through an OSC termination, GCC termination, or both
- An OSS has IP connectivity to all the NEs through third party NE network
- The MSTP NE is a GNE for the third party OSI-based SONET NEs. the MSTP NEs perform all mediation functions.
- The third party vendor NE is a GNE for the Cisco MSTP network.

Figure 8-22 OSI/IP Scenario 4



8.8 Link Management Protocol

This section describes Link Management Protocol (LMP) management and configuration. To troubleshoot specific alarms, refer to the *Cisco ONS 15454 DWDM Troubleshooting Guide*. To configure LMP, refer to the *Cisco ONS 15454 DWDM Procedure Guide*.



Note

Cisco Transport Manager (CTM) support is not required for LMP.

LMP is used to establish traffic engineering (TE) links between Cisco ONS 15454 nodes or between Cisco ONS 15454 nodes and selected non-Cisco nodes that use vendor-specific hardware.

8.8.1 Overview

LMP manages TE links between nodes through the use of control channels. TE links are designed to define the most efficient paths possible for traffic to flow over a network and through the Internet. Traffic engineering encompasses traffic management, capacity management, traffic measurement and modeling,

network modeling, and performance analysis. Traffic engineering methods include call routing, connection routing, quality of service (QoS) resource management, routing table management, and capacity management.

LMP manages TE links between peer nodes, such as two optical cross-connect (OXC) nodes. Peer nodes have equivalent signaling and routing. LMP also manages TE links between a node such as an OXC and an adjacent optical line system (OLS) node. An example of an OLS node is an ONS 15454 DWDM node.

Networks with routers, switches, OXC nodes, DWDM OLS nodes, and add/drop multiplexers (ADM) use a common control plane such as Generalized Multiprotocol Label Switching (GMPLS) to provision resources and provide network survivability using protection and restoration techniques. LMP is part of the GMPLS protocol suite.

A single TE link can be formed from several individual links. Management of TE links can be accomplished with in-band messaging, as well as with out-of-band methods. The following material describes the link management protocol between a pair of nodes that manages TE links. LMP accomplishes the following:

- Maintains control channel connectivity
- Verifies the physical connectivity of the data links
- Correlates the link property information
- Suppresses downstream alarms
- Localizes link failures for protection/restoration purposes in multiple types of networks

DWDM networks often use MPLS and GMPLS as common-control planes to control how packets are routed through the network.

LMP manages the control channel that must exist between nodes for routing, signaling, and link management. For a control channel to exist, each node must have an IP interface that is reachable from the other node. Together, the IP interfaces form a control channel. The interface for the control messages does not have to be the same interface as the one for the data.

The LMP protocol is specified by the IETF in an Internet-Draft, draft-ietf-ccamp-lmp-10.txt, which was published as a Proposed Standard, RFC 4204, (<http://www.ietf.org/rfc/rfc4204.txt>), on 2005-10-28.

8.8.1.1 MPLS

MPLS provides a mechanism for engineering network traffic patterns that is independent of routing tables and routing protocols. MPLS assigns short labels to network packets that describe how to forward the packets through the network. The traditional Layer 3 forwarding mechanism requires each hop to analyze the packet header and determine the next hop based on routing table lookup. With MPLS, the analysis of the packet header is performed just once, when a packet enters the MPLS cloud. The packet is then assigned to a stream known as a Label Switch Path (LSP), which is identified with a label. The short, fixed-length label is an index into a forwarding table, which is more efficient than the traditional routing table lookup at each hop. Using MPLS, both the control protocol (used to manage the LSPs) and user data can be carried over the same bearer interfaces.

8.8.1.2 GMPLS

GMPLS is based on MPLS, with protocol extensions to support additional technologies, including time division multiplexing (TDM) slots (such as SONET and SDH), wavelength division multiplexing (WDM) wavelengths at Layer 1, and fiber. For MPLS, the control traffic (signaling and routing) can run

over bearer interfaces. This is not the case with GMPLS, where a separate control channel is used. The GMPLS control channel is managed with LMP. With GMPLS, the control channels between two adjacent nodes do not need to use the same physical medium as the data links between those nodes.

8.8.2 Configuring LMP

Configuring LMP consists of the following four topics:

- Control channel management
- TE link management
- Link connectivity verification
- Fault management

8.8.2.1 Control Channel Management

Control channel management establishes and maintains control channels between adjacent nodes. Control channels use a Config message exchange and a fast keep-alive mechanism between the nodes. The latter is required if lower-level mechanisms are not available to detect control-channel failures. A maximum of four LMP control channels can be supported.

The nodes initially exchange configuration messages (Config, ConfigAck, and ConfigNack), which are used to exchange identifiers and negotiate parameters for the keep-alive protocol. The nodes then perform a continuous rapid exchange of Hello messages, which are used to monitor the health on the channel.



Note The identifiers are Local Node Id, Remote Node Id, Local Control Channel Id, and Remote Control Channel Id. The parameters are the HelloInterval and the HelloDeadInterval.

LMP out-of-fiber and LMP out-of-band control channels are supported and terminated on the shelf. An out-of-fiber control channel includes using the control plane network (Ethernet) for the control channel because Ethernet is separate from the fiber used for the data plane. An out-of-band control channel includes using overhead bytes, such as the SDCC and LDCC bytes, for the control channel because overhead bytes are separate from the payload. In-band means that the control messages are in the same channel as the data messages; therefore, out-of-band refers to overhead bytes in the same fiber, separate circuits dedicated to control messages in the same fiber (SONET/SDH circuits), or separate wavelengths in the same fiber (DWDM).



Note Overhead bytes are SDCC or LDCC for SONET networks, RS-DCC or MS-DCC for SDH networks, and GCC or OSC for DWDM networks.

Out-of-band implies in-fiber, but not in-band. In-fiber means that the control messages are in the same fiber as the data messages, and includes both in-band and out-of-band. Out-of-fiber means that the control messages take a path separate from the data plane. This includes separate fiber and Ethernet.

The control channel management for a peer node to OLS link is the same as that for a link between two peer nodes.

**Note**

The software supports gracefully taking a control channel down for administration purposes (refer to Section 3.2.3 of the IETF LMP document). However, there is no provision for a graceful restart (refer to Section 8 of RFC 4204).

- Graceful means that the nodes participating in the control channel agree that the link should go down. To gracefully take down a control channel, the node sets the ControlChannelDown flag in its messages to the other node until either the HelloDeadInterval expires or the other node sends a message back with the ControlChannelDown flag set. In either case, the node then stops sending messages for this control channel. Before a control channel is taken down, there should be a backup control channel in place that can be used to manage the data links.
- Non-graceful means that one of the nodes just stops sending messages. The other side would declare a failure after the HelloDeadInterval, but would continue to send Hello messages to see if the control channel will come back up.

8.8.2.2 TE Link Management

LMP ensures that links are grouped into TE links and that the properties of those links are the same at both endpoints. This is called TE link management, or link property correlation.

Link property correlation is used to synchronize the TE link properties and verify the TE link configuration. The link property correlation function of LMP aggregates one or more data links into a TE link and synchronizes the properties of the TE link with the neighbor node. The procedure starts by sending a LinkSummary message to the neighbor. The LinkSummary message includes the local and remote Link Identifier, a list of all data links that make up the TE link, and various link properties. It is mandatory that a LinkSummaryAck or LinkSummaryNack message be sent in response to the receipt of a LinkSummary message, indicating agreement or disagreement with the link properties.

**Note**

A maximum of 256 LMP TE links is supported.

8.8.2.3 Link Connectivity Verification

Link connectivity verification might be supported in future CTC software releases.

8.8.2.4 Fault Management

Fault management is particularly useful when the control channels are physically diverse from the data links. It is used for rapid notification regarding the status of one or more TE-link data channels. The use of fault management is negotiated as part of the TE link's LinkSummary exchange. Data links and TE link failures can be rapidly isolated and fault management supports both unidirectional and bidirectional LSPs. Transparent devices are useful because traditional methods for monitoring the health of allocated data links might no longer be appropriate. Instead, fault detection is delegated to the physical layer (for example, loss of light or optical monitoring of the data) instead of Layer 2 or Layer 3. Fault management uses the ChannelStatus, ChannelStatusAck, ChannelStatusRequest, and ChannelStatusResponse messages.

**Note**

The LMP Channel Activation/Deactivation Indication procedures are not supported; they are described in the IETF LMP document, Sections 6.4 and 6.5.

8.8.3 LMP WDM

LMP manages traffic engineering links between peer nodes (nodes that are peers in signaling and/or routing).




Note

LMP-WDM extensions that allow management of links between a peer node and an adjacent OLS node are described in the following IETF document: Internet-Draft, draft-ietf-ccamp-lmp-wdm-03.txt, published as a Proposed Standard, RFC 4209 (<http://www.ietf.org/rfc/rfc4209.txt>), 2005-11-1

The purpose of the LMP WDM extensions is to allow LMP to be used between an OXC node and an adjacent DWDM OLS node. [Figure 8-23](#) illustrates the relationship between LMP and LMP-WDM. OXC 1 and OXC 2 are peer nodes whose control channel is managed with LMP. LMP-WDM manages the control channel between an OXC node and an OLS node.

Figure 8-23 LMP and LMP-WDM Relationship



When the two OLS nodes can communicate their configuration and the current state of their optical link to the two peer nodes (OXC 1 and OXC 2) through LMP-WDM, network usability is improved through the reduction of manual configuration and enhanced fault detection and recovery.

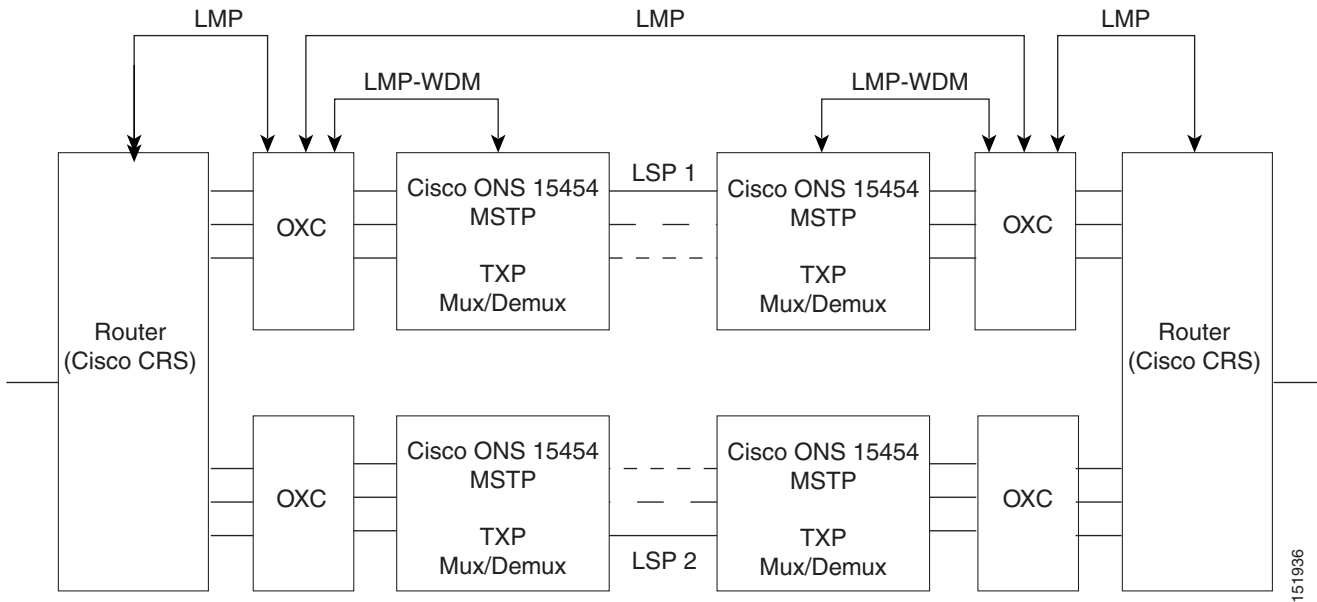
8.8.4 LMP Network Implementation

[Figure 8-24](#) shows a network-level LMP implementation. It is an IP-plus-optical network, with end-to-end routing based on MPLS and GMPLS. The primary network components are:

- Routers
 - Cisco Carrier Router System (CSR)
 - Cisco Gigabit Switch Router (GSR)
- An OXC node
- Ultra long-haul (ULH) DWDM equipment

LMP and other features allow the Cisco ONS 15454 DWDM node to fulfill the ultra long haul (ULH) DWDM role. [Figure 8-24](#) illustrates the relationship between the network components.

Figure 8-24 LMP System Implementation



151936

