



Ethernet Operation



Note

The terms "Unidirectional Path Switched Ring" and "UPSR" may appear in Cisco literature. These terms do not refer to using Cisco ONS 15xxx products in a unidirectional path switched ring configuration. Rather, these terms, as well as "Path Protected Mesh Network" and "PPMN," refer generally to Cisco's path protection feature, which may be used in any topological network configuration. Cisco does not recommend using its path protection feature in any particular topological network configuration.

The Cisco ONS 15454 integrates Ethernet into a SONET time-division multiplexing (TDM) platform. The ONS 15454 supports E-Series, G-Series, and ML-Series Ethernet cards. This chapter covers the operation of the E-Series and G-Series Ethernet cards.

For Ethernet card specifications, see Chapter 5, "Ethernet Cards." For information on the ML-Series cards, refer to the *Cisco ONS 15454 ML-Series Multilayer Ethernet Card Software Feature and Configuration Guide*. For step-by-step Ethernet card circuit configuration procedures, refer to the "Create Circuits and VT Tunnels" chapter of the *Cisco ONS 15454 Procedure Guide*.

Chapter topics include:

- G-Series Application, page 1
- E-Series Application, page 10
- G-Series Circuit Configurations, page 19
- E-Series Circuit Configurations, page 20
- Remote Monitoring Specification Alarm Thresholds, page 24

16.1 G-Series Application

The G-Series cards (G1000-4/G1K-4) reliably transport Ethernet and IP data across a SONET backbone. The G-Series card maps up to four Gigabit Ethernet interfaces onto a SONET transport network and provides scalable and provisionable transport bandwidth at signal levels up to STS-48c per card. The G-Series card provides line rate forwarding for all Ethernet frames (unicast, multicast, and broadcast) and can be configured to support Jumbo frames (defined as a maximum of 10,000 bytes). The G-Series card incorporates features optimized for carrier-class applications such as:

- High Availability (HA), including hitless (< 50 ms) performance under software upgrades and all types of SONET/SDH equipment protection switches.
- Hitless reprovisioning.
- Support of Gigabit Ethernet traffic at full line rate.

- Full TL1-based provisioning capability. Refer to the *Cisco ONS 15454 and Cisco ONS 15327 TL1 Command Guide* for G-Series TL1 provisioning commands.
- Serviceability options including enhanced port states, terminal and facility loopback, and J1 path trace.
- SONET-style alarm support.
- Ethernet performance monitoring (PM) and remote monitoring (RMON) functions.

The G-Series card allows an Ethernet private line service to be provisioned and managed very much like a traditional SONET or SDH line. G-Series card applications include providing carrier-grade transparent LAN services (TLS), 100 Mbps Ethernet private line services (when combined with an external 100 Mb Ethernet switch with Gigabit uplinks), and high-availability transport for applications such as storage over MAN/WANs.

The card maps a single Ethernet port to a single STS circuit. You can independently map the four ports on the G-Series card to any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c, and STS-48c circuit sizes, provided the sum of the circuit sizes that terminate on a card do not exceed STS-48c.

To support a Gigabit Ethernet port at full line rate, an STS circuit with a capacity greater or equal to 1 Gbps (bidirectional 2 Gbps) is needed. An STS-24c is the minimum circuit size that can support a Gigabit Ethernet port at full line rate. The G-Series card supports a maximum of two ports at full line rate.

The G-Series transmits and monitors the SONET J1 Path Trace byte in the same manner as ONS 15454 OC-N cards. For more information on the J1 Path Trace, refer to the “Path Trace” section on page 21.


Note

G-Series encapsulation is standard HDLC framing over SONET/SDH as described in RFC 1622 and RFC 2615 with the PPP protocol field set to the value specified in RFC 1841.

16.1.1 G1K-4 and G1000-4 Comparison

The G1K-4 and the G1000-4 cards comprise the ONS 15454 G-Series and are hardware equivalents.

When installed in ONS 15454s running software prior to Software Release 4.0, both cards require the XC10G card to operate. However, when installed on an ONS 15454 running Software R4.0 and later, the G1K-4 card is not limited to installation in ONS 15454s with XC10G cards but can also be installed in ONS 15454s with XC and XCVT cards. When used with XC and XCVT cards on an ONS 15454 running Release 4.0 and later, the G1K-4 is limited to the high-speed slots (Slots 5, 6, 12, and 13.)


Software R4.0 and later identifies G1K-4 cards as G1K-4s upon physical installation. Software prior to R4.0 identifies both G1000-4 and G1K-4 cards as G1000-4s upon physical installation.

These constraints do not apply to a G-Series card configured for Gigabit Ethernet Transponder Mode, see “G-Series Gigabit Ethernet Transponder Mode” section on page 5.

16.1.2 G-Series Example

Figure 16-1 on page 16-3 shows an example of a G-Series application. In this example, data traffic from the Gigabit Ethernet port of a high-end router travels across the ONS 15454 point-to-point circuit to the Gigabit Ethernet port of another high-end router.

Figure 16-1 Data Traffic On G-Series Point-To-Point Circuit



The G-Series card carries any layer three protocol that can be encapsulated and transported over Gigabit Ethernet, such as IP or IPX. The data is transmitted on the Gigabit Ethernet fiber into the standard Cisco Gigabit Interface Converter (GBIC) on a G-Series card. The G-Series card transparently maps Ethernet frames into the SONET payload by multiplexing the payload onto a SONET OC-N card. When the SONET payload reaches the destination node, the process is reversed and the data is transmitted from the standard Cisco GBIC in the destination G-Series card onto the Gigabit Ethernet fiber.

The G-Series card discards certain types of erroneous Ethernet frames rather than transport them over SONET. Erroneous Ethernet frames include corrupted frames with CRC errors and under-sized frames that do not conform to the minimum 64-byte length Ethernet standard. The G-Series card forwards valid frames unmodified over the SONET network. Information in the headers is not affected by the encapsulation and transport. For example, packets with formats that include IEEE 802.1Q information will travel through the process unaffected.

16.1.3 802.3z Flow Control and Frame Buffering

The G-Series supports IEEE 802.3z flow control and frame buffering to reduce data traffic congestion. To prevent over-subscription, 512 KB of buffer memory is available for the receive and transmit channels on each port. When the buffer memory on the Ethernet port nears capacity, the ONS 15454 uses IEEE 802.3z flow control to transmit a pause frame to the source at the opposite end of the Gigabit Ethernet connection.

The pause frame instructs the source to stop sending packets for a specific period of time. The sending station waits the requested time before sending more data. Figure 16-1 illustrates pause frames being sent and received by ONS 15454s and attached switches.

The G-Series card has symmetric flow control. It proposes symmetric flow control when auto-negotiating flow control with attached Ethernet devices. Symmetric flow control allows the G-Series to respond to pause frames sent from external devices and send pause frames to external devices. Prior to Software R4.0, flow control on the G-Series card was asymmetric, meaning the card sent pause frames and discarded received pause frames.

This flow-control mechanism matches the sending and receiving device throughput to that of the bandwidth of the STS circuit. For example, a router might transmit to the Gigabit Ethernet port on the G-Series. This particular data rate may occasionally exceed 622 Mbps, but the ONS 15454 circuit assigned to the G-Series port might be only STS-12c (622.08 Mbps). In this example, the ONS 15454 sends out a pause frame and requests that the router delay its transmission for a certain period of time. With flow control and a substantial per-port buffering capability, a private line service provisioned at less than full line rate capacity (STS-24c) is efficient because frame loss can be controlled to a large extent.

Release 4.1 introduces flow control threshold provisioning to the G-Series card, which allows a user to select one of three watermark (buffer size) settings: default, low latency, or custom. Default is the best setting for general use and was the only setting available prior to release 4.1. Low latency is good for sub-rate applications, such as VoIP over an STS-1. For attached devices with insufficient buffering, best effort traffic or long access line lengths, set the G-Series to a higher latency.

The custom setting allows you to specify an exact buffer size threshold for Flow Ctrl Lo and Flow Ctrl Hi. The flow control high setting is the watermark for sending the “Pause On” frame to the attached Ethernet device, this frame signals the device to temporarily stop transmitting. The flow control low setting is the watermark for sending the “Pause Off” frame, which signals the device to resume transmitting.

**Note**

External Ethernet devices with auto-negotiation configured to interoperate with G-Series cards running releases prior to R4.0 do not need to change auto-negotiation settings when interoperating with G-Series cards running R4.0 and later.


**Note**

With a G-Series card, you can only enable flow control on a port if auto-negotiation is enabled on the device attached to that port.

16.1.4 Ethernet Link Integrity Support

The G-Series supports end-to-end Ethernet link integrity (Figure 16-2). This capability is integral to providing an Ethernet private line service and correct operation of Layer 2 and Layer 3 protocols on the attached Ethernet devices. End-to-end Ethernet link integrity essentially means that if any part of the end-to-end path fails the entire path fails. Failure of the entire path is ensured by turning off the transmit lasers at each end of the path. The attached Ethernet devices recognize the disabled transmit laser as a loss of carrier and consequently an inactive link.

Figure 16-2 End-to-End Ethernet Link Integrity Support

**Note**

Some network devices can be configured to ignore a loss of carrier condition. If a device configured to ignore a loss of carrier condition attaches to a G-Series card at one end, alternative techniques (such as use of Layer 2 or Layer 3 keep-alive messages) are required to route traffic around failures. The response time of such alternate techniques is typically much longer than techniques that use link state as indications of an error condition.

**Note**

Enabling or disabling port-level flow control on the test set or other Ethernet device attached to the G-Series port can affect the transmit (Tx) laser and result in unidirectional traffic flow.


As shown in Figure 16-2, a failure at any point of the path causes the G-Series card at each end to disable its Tx transmit laser, which causes the devices at both ends to detect a link down. If one of the Ethernet ports is administratively disabled or set in loopback mode, the port is considered a “failure” for the purposes of end-to-end link integrity because the end-to-end Ethernet path is unavailable. The port “failure” also disables both ends of the path.

16.1.5 Gigabit EtherChannel/802.3ad Link Aggregation

The end-to-end Ethernet link integrity feature can be used in combination with Gigabit EtherChannel capability on attached devices. The combination provides an Ethernet traffic restoration scheme that has a faster response time than alternate techniques such as spanning tree rerouting, yet is more bandwidth efficient because spare bandwidth does not need to be reserved.

The G-Series supports all forms of link aggregation technologies including Gigabit EtherChannel (GEC), which is a Cisco proprietary standard, and the IEEE 802.3ad standard. The end-to-end link integrity feature of the G-Series allows a circuit to emulate an Ethernet link. This allows all flavors of Layer 2 and Layer 3 rerouting to work correctly with the G-Series. Figure 16-3 illustrates G-Series GEC support.

Figure 16-3 G-Series Gigabit EtherChannel (GEC) Support



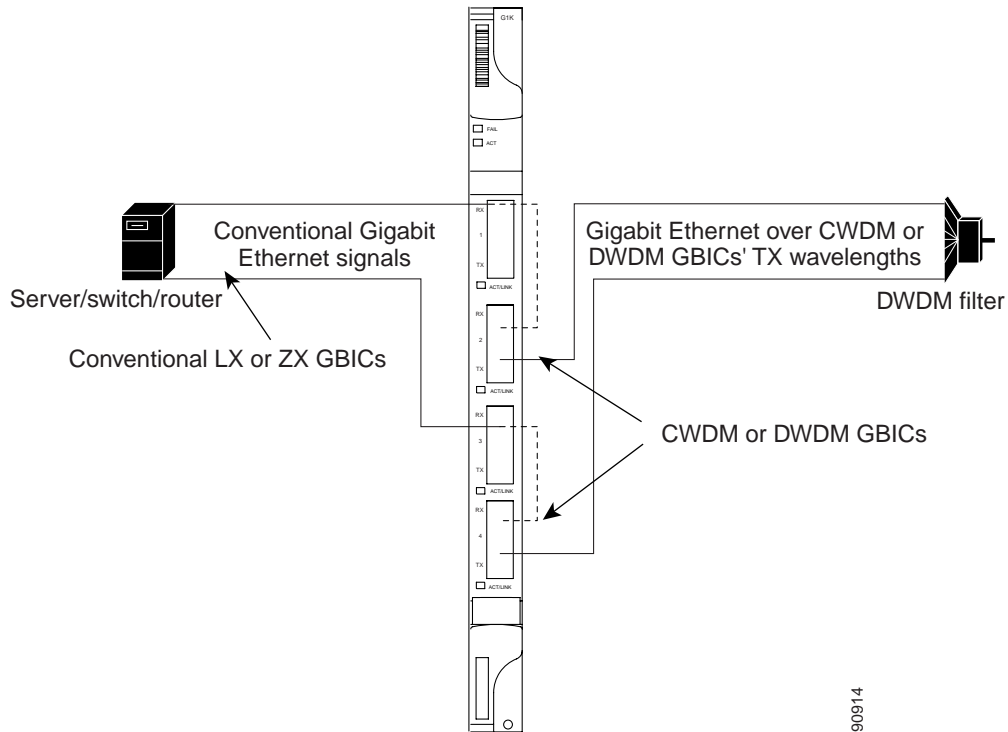
Although the G-Series card does not actively run GEC, it supports the end-to-end GEC functionality of attached Ethernet devices. If two Ethernet devices running GEC connect through G-Series cards to an ONS 15454 network, the ONS 15454 SONET side network is transparent to the EtherChannel devices. The EtherChannel devices operate as if they are directly connected to each other. Any combination of G-Series parallel circuit sizes can be used to support GEC throughput.

GEC provides line-level active redundancy and protection (1:1) for attached Ethernet equipment. It can also bundle parallel G-Series data links together to provide more aggregated bandwidth. STP operates as if the bundled links are one link and permits GEC to utilize these multiple parallel paths. Without GEC, STP permits only a single non-blocked path. GEC can also provide G-Series card-level protection or redundancy because it can support a group of ports on different cards (or different nodes) so that if one port or card has a failure, traffic is rerouted over the other port or card.

16.2 G-Series Gigabit Ethernet Transponder Mode

Starting with Software Release 4.1, the G-Series card can be configured as a transponder. Transponder mode can be used with any G-Series supported GBIC (SX, LX, Zx, CWDM, or DWDM). [Figure 16-4 on page 16-6](#) shows a card level overview of a transponder mode application.

Figure 16-4 Card Level Overview of G-Series One-Port Transponder Mode Application



A G-Series card configured as a transponder operates quite differently than a G-Series card configured for SONET. In SONET configuration, the G-Series card receives and transmits Gigabit Ethernet traffic out the Ethernet ports and GBICs on the front of the card. This Ethernet traffic is multiplexed on and off the SONET network through the cross-connect card and the OC-N card (Figure 16-5). In transponding mode, the G-Series Ethernet traffic never comes into contact with the cross-connect card or the SONET network, but stays internal to the G-Series card and is routed back to a GBIC on that card (Figure 16-6).

Figure 16-5 G-Series in Default SONET Mode

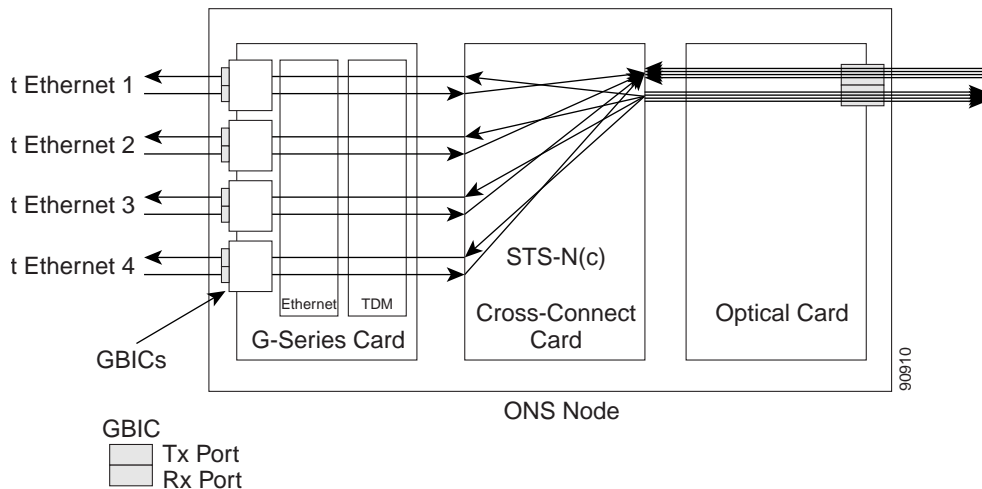
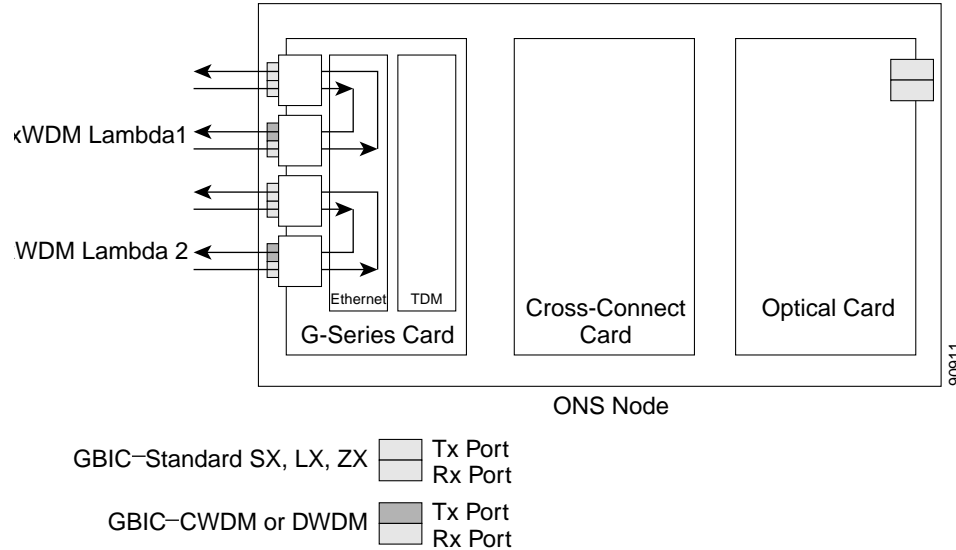


Figure 16-6 G-Series Card in Transponder Mode (Two-Port Bidirectional)



A G-Series card can either be configured for transponding mode or as the SONET default. Once any port is provisioned in transponding mode, the card is in transponding mode and no SONET circuits can be configured until every port on the card goes back to SONET mode. Refer to the *Cisco ONS 15454 Procedure Guide* for detailed instructions on how to provision G-Series ports for transponder mode.

All SONET circuits must be deleted before a G-Series card can be configured in transponding mode. An ONS 15454 can host the card in any or all of the twelve traffic slots on the ONS 15454 and supports a maximum of 24 bidirectional or 48 unidirectional lambdas.

A G-Series card configured as a transponder can be in one of three modes:

- Two-port bidirectional transponding mode
- One-port bidirectional transponding mode
- Two-port unidirectional transponding mode

16.2.1 Two-Port Bidirectional Transponder

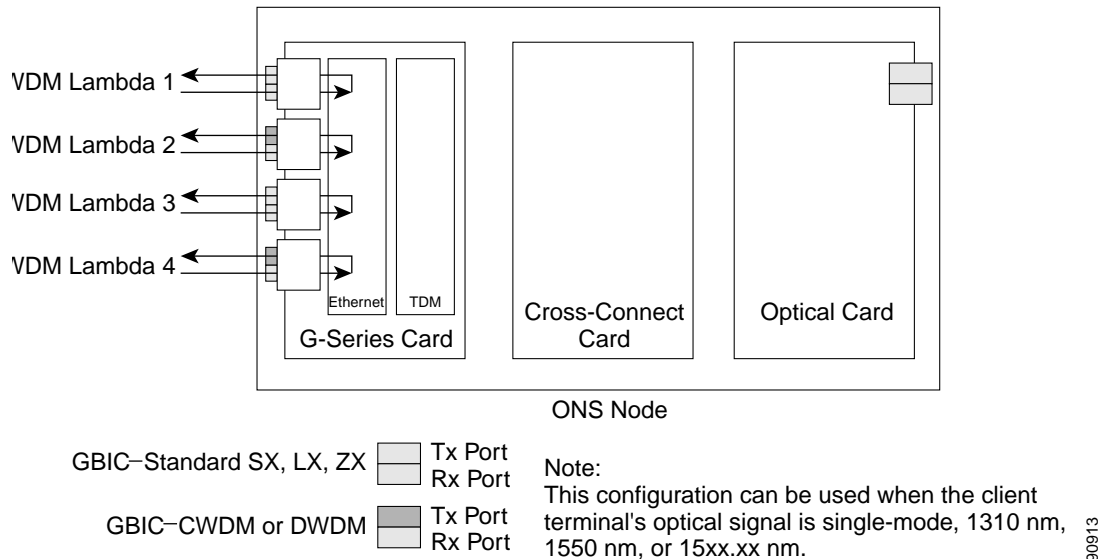
Two-port bidirectional transponder mode maps the transmitted and received Ethernet frames of one G-Series card port into the transmit and receive of another port (Figure 16-6). Transponder bidirectional port mapping can be any port to any port on the same card.

16.2.2 One-Port Bidirectional Transponder

One-port bidirectional transponder mode (Figure 16-7 on page 16-8) maps the Ethernet frames received at a port out the transmitter of the same port. This mode is similar to two-port bidirectional transponder mode except that a port is mapped only to itself instead of to another port. Although the data path of the one-port bidirectional transponder mode is identical to that of a facility loopback. The transponding mode is not a maintenance mode and does not suppress non-SONET alarms, like loss of carrier (CARLOSS).

This mode can be used for intermediate DWDM signal regeneration and to take advantage of the wide band capability of the CWDM and DWDM GBICs, which allows you to receive on multiple wavelengths but transmit on a fixed wavelength.

Figure 16-7 One-Port Bidirectional Transponding Mode

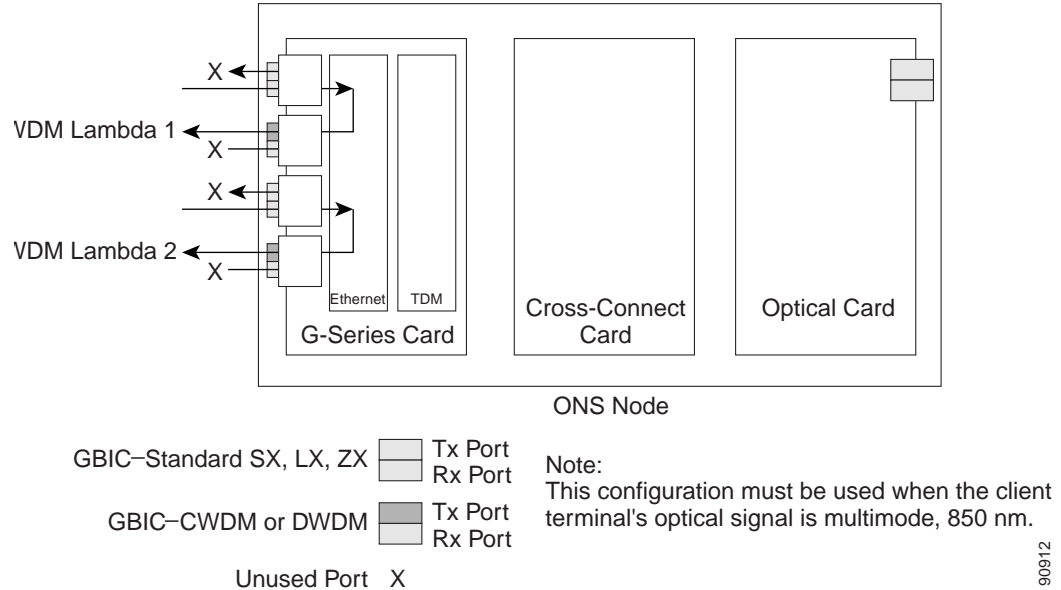


16.2.3 Two-Port Unidirectional Transponder

Ethernet frames received at one port's receiver will be transmitted out the transmitter of another port. This mode is similar to two-port bidirectional transponder mode except only one direction is used (Figure 16-8). One port has to be provisioned as unidirectional transmit only and the other port as unidirectional receive. The port configured as unidirectional transmit ignores any lack of signal on the receive port, so the receive port fiber does not need to be connected. The port configured as unidirectional receive does not turn on the transmit laser, and so the transmit port fiber does not need to be connected.

This mode can be used when only one direction needs to be transmitted over CWDM/DWDM, for example certain VoD applications.

Figure 16-8 Two-Port Unidirectional Transponder



16.2.4 G-Series Transponder Mode Characteristics

The operation of a G-Series card in transponder mode differs from a G-Series card in SONET mode in several ways:

- A G-Series card set to transponder mode will not show up in the CTC list of provisionable cards when the user is provisioning a SONET circuit.
- G-Series cards set to transponder mode do not require cross-connect cards (XC, XCVT or XC10G), but do require timing communications and control cards (TCC+/TCC2).
- G-Series ports configured as transponders do not respond to flow control pause frames and pass the pause frames transparently through the card. In SONET mode, ports can respond to pause frames and do not pass the pause frames through the card.
- There is no TL1 provisioning support for configuring transponding mode. However, transponding mode and port information can be retrieved in the results for the TL1 command RTRV-G1000.
- All SONET related alarms are suppressed when a card is set in transponding mode.
- There are no slot number or cross-connect restrictions for G1000-4 or G1K-4 cards in transponder mode.
- Facility and terminal loopbacks are not fully supported in unidirectional transponding mode, but are supported in both bidirectional transponding modes.
- Ethernet autonegotiation is not supported and cannot be provisioned in unidirectional transponding mode. Autonegotiation is supported in both bidirectional transponding modes.
- No end-to-end link integrity function is available in transponding mode.

**Note**

In normal SONET mode the G-Series cards supports an end-to-end link integrity function. This function causes an Ethernet or SONET failure to disable and turn the transmitting laser off the corresponding mapped Ethernet port. In transponder mode, the loss of signal on an Ethernet port has no impact on the transmit signal of the corresponding mapped port.

The operation of a G-Series card in transponder mode is also similar to the operation of a G-Series card in SONET mode:

- G-Series Ethernet statistics are available for ports in both modes.
- Ethernet port level alarms and conditions are available for ports in both modes.
- Jumbo frame and non-jumbo frame operation is the same in both modes.
- Collection, reporting, and threshold crossing conditions for all existing counters and PM parameters are the same in both modes.
- SNMP and RMON support is the same in both modes.

16.3 E-Series Application

The ONS 15454 E-Series cards include the E100T-12/E100T-G and E1000-2/E1000-2-G. An ONS 15454 supports a maximum of ten E-Series cards, and you can insert Ethernet cards in any multipurpose slot.

The E100T-G is the functional equivalent of the E100T-12. The E1000-2 is the functional equivalent of the E1000-2-G. An ONS 15454 using XC10G cards requires the G versions of the E-Series Ethernet cards.

16.3.1 E-Series Modes

An E-Series card operates in one of three modes: multicard EtherSwitch group, single-card EtherSwitch, or port-mapped. E-Series cards in multicard EtherSwitch Group or single-card EtherSwitch mode support Layer 2 features, including virtual local area networks (VLANs), IEEE 802.1Q, STP, and IEEE 802.1D. Port-mapped mode configures the E-Series to operate as a straight mapper card and does not support these Layer 2 features. Within an ONS 15454 containing multiple E-Series cards, each E-Series card can operate in any of the three separate modes. At the Ethernet card view in CTC, click the **Provisioning > Ether Card** tabs to reveal the card modes.


**Note**

Port-mapped mode eliminates issues inherent in other E-Series modes and detailed in the field notice, “E-Series Ethernet Line Card Packet Forwarding Limitations.”

16.3.1.1 E-Series Multicard EtherSwitch Group

Multicard EtherSwitch Group provisions two or more Ethernet cards to act as a single Layer 2 switch. It supports one STS-6c circuits, two STS-3c circuits, or six STS-1 circuits. Each multicard switch may connect up to a total of STS-6c in SONET circuits. When provisioned as an add or drop node of a shared packet ring circuit, the effective bandwidth doubles, supporting STS-6c in each direction of the ring. [Figure 16-9 on page 16-11](#) illustrates a multicard EtherSwitch configuration.

Figure 16-9 Multicard EtherSwitch Configuration


**Caution**

Whenever you terminate two STS-3c multicard EtherSwitch circuits on an Ethernet card and later delete the first circuit. Also delete the remaining STS-3c circuit, before you provision an STS-1 circuit to the card. If you attempt to create an STS-1 circuit after only deleting the first STS-3c circuit, the STS-1 circuit will not work, but no alarms will indicate this condition. To avoid this situation, delete the second STS-3c before creating an STS-1 circuit.

16.3.1.2 E-Series Single-Card EtherSwitch

Single-card EtherSwitch allows each Ethernet card to remain a single switching entity within the ONS 15454 shelf. This option allows STS-12c worth of bandwidth between two Ethernet circuit endpoints. Figure 16-10 illustrates a single-card EtherSwitch configuration.

Figure 16-10 Single-card EtherSwitch Configuration

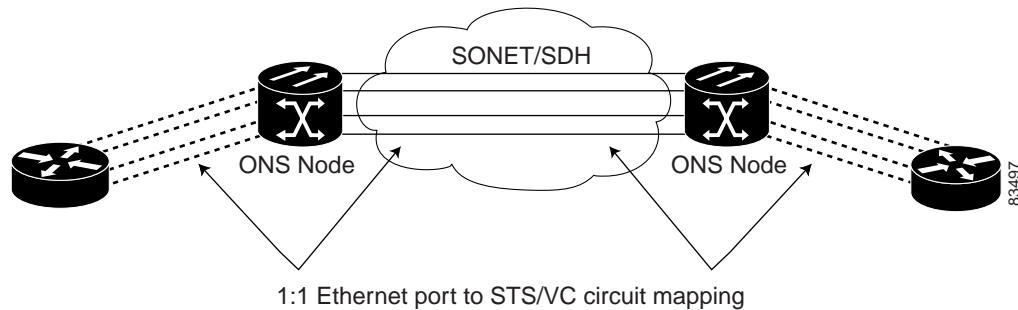


16.3.1.3 Port-Mapped (Linear Mapper)

Port-mapped mode, also referred to as linear mapper, configures the E-Series card to map a specific E-Series Ethernet port to one of the card's specific STS circuits. (See [Figure 16-11 on page 16-12.](#)) Port-mapped mode ensures Layer one transport has low latency for unicast, multicast, and mixed traffic.

Ethernet and Fast Ethernet on the E100T-G card operate at line-rate speed. Gigabit Ethernet transport is limited to a maximum of 600 Mbps because the E1000-2-G has a maximum bandwidth of STS-12c. Ethernet frame sizes up to 1522 bytes are also supported, which allows transport of IEEE 802.1Q tagged frames. The larger maximum frame size of Q-in-Q frames, IEEE 802.1Q in IEEE 802.1Q wrapped frames, are not supported.

Figure 16-11 E-Series Mapping Ethernet Ports to SONET STS Circuits



Port-mapped mode disables Layer 2 functions supported by the E-Series in single-card and multicard mode, including STP, VLANs, and MAC address learning. It significantly reduces the service-affecting time for cross-connect and TCC+/TCC2 card switches.

Port-mapped mode does not support VLANs in the same manner as multicard and single-card mode. The ports of E-Series cards in multicard and single-card mode can join specific VLANs. E-Series cards in port-mapped mode do not have this Layer 2 capability and only transparently transport external VLANs over the mapped connection between ports. An E-Series card in port-mapped mode does not inspect the tag of the transported VLAN, so a VLAN range of 1 through 4096 can be transported in port-mapped mode.

Port-mapped mode does not perform any inspection or validation of the Ethernet frame header. The Ethernet Cyclic Redundancy Check (CRC) is validated, and any frame with an invalid Ethernet CRC is discarded.

Port-mapped mode also allows the creation of STS circuits between any two E-Series cards, including the E100T-G, E1000-G, or the E10/100-4 (the ONS 15327 E-Series card). Port-mapped mode does not allow an E-Series cards to connect to the ML-Series or G-Series cards.

16.3.2 E-Series IEEE 802.3z Flow Control

The E100T-G in any mode and the E1000-G in port-mapped mode support IEEE 802.3z symmetrical flow control and propose symmetric flow control when auto-negotiating with attached Ethernet devices. For flow control to operate, both the E-Series port and the attached Ethernet device must be set to auto-negotiation (AUTO) mode. The attached Ethernet device may also need to have flow control enabled. The flow-control mechanism allows the E-Series to respond to pause frames sent from external devices and send pause frames to external devices.

Flow control matches the sending and receiving device throughput to that of the bandwidth of the STS circuit. For example, a router might transmit to the Gigabit Ethernet port on the E-Series in port mapped mode. The data rate transmitted by the router may occasionally exceed 622 Mbps, but the ONS 15454 circuit assigned to the E-Series port in port-mapped mode is a maximum of STS-12c (622.08 Mbps). In this scenario, the ONS 15454 sends out a pause frame and requests that the router delay its transmission for a certain period of time.

**Note**

To enable flow control between an E-Series in port mapped mode and a SmartBits test set, manually set bit 5 of the MII register to 0 on the SmartBits test set. To enable flow control between an E-Series in port mapped mode and an Ixia test set, select Enable the flow control in the properties menu of the attached Ixia port.

16.3.3 E-Series VLAN Support

Users can provision up to 509 VLANs per network with the CTC software. Specific sets of ports define the broadcast domain for the ONS 15454. The definition of VLAN ports includes all Ethernet and packet-switched SONET port types. All VLAN IP address discovery, flooding, and forwarding is limited to these ports.


The ONS 15454 IEEE 802.1Q-based VLAN mechanism provides logical isolation of subscriber LAN traffic over a common SONET transport infrastructure. Each subscriber has an Ethernet port at each site, and each subscriber is assigned to a VLAN. Although the subscriber's VLAN data flows over shared circuits, the service appears to the subscriber as a private data transport.

**Note**

Port-mapped mode does not support VLANs.

The number of VLANs used by circuits and the total number of VLANs available for use appears in CTC on the VLAN counter (Figure 16-12).

Figure 16-12 Edit Circuit Dialog Featuring Available VLANs



16.3.4 E-Series Q-Tagging (IEEE 802.1Q)

E-Series cards in single-card and multiscard mode support IEEE 802.1Q. IEEE 802.1Q allows the same physical port to host multiple 802.1Q VLANs. Each 802.1Q VLAN represents a different logical network. E-Series cards in port-mapped mode transport IEEE 802.1Q tags (Q-tags), but do not remove or add these tags.

The ONS 15454 works with Ethernet devices that support IEEE 802.1Q and those that do not support IEEE 802.1Q. If a device attached to an ONS 15454 Ethernet port does not support IEEE 802.1Q, the ONS 15454 uses Q-tags internally only. The ONS 15454 associates these Q-tags with specific ports.

With Ethernet devices that do not support IEEE 802.1Q, the ONS 15454 takes non-tagged Ethernet frames that enter the ONS network and uses a Q-tag to assign the packet to the VLAN associated with the ONS network's ingress port. The receiving ONS node removes the Q-tag when the frame leaves the ONS network (to prevent older Ethernet equipment from incorrectly identifying the 802.1Q packet as an illegal frame). The ingress and egress ports on the ONS network must be set to Untag for the removal to occur. Untag is the default setting for ONS ports. Example 1 in Figure 16-13 illustrates Q-tag use only within an ONS network.

The ONS 15454 uses the Q-tag attached by the external Ethernet devices that support IEEE 802.1Q. Packets enter the ONS network with an existing Q-tag; the ONS 15454 uses this same Q-tag to forward the packet within the ONS network and leaves the Q-tag attached when the packet leaves the ONS network. The entry and egress ports on the ONS network must be set to Tagged for this process to occur. Example 2 in Figure 16-13 on page 16-15 illustrates the handling of packets that both enter and exit the ONS network with a Q-tag.

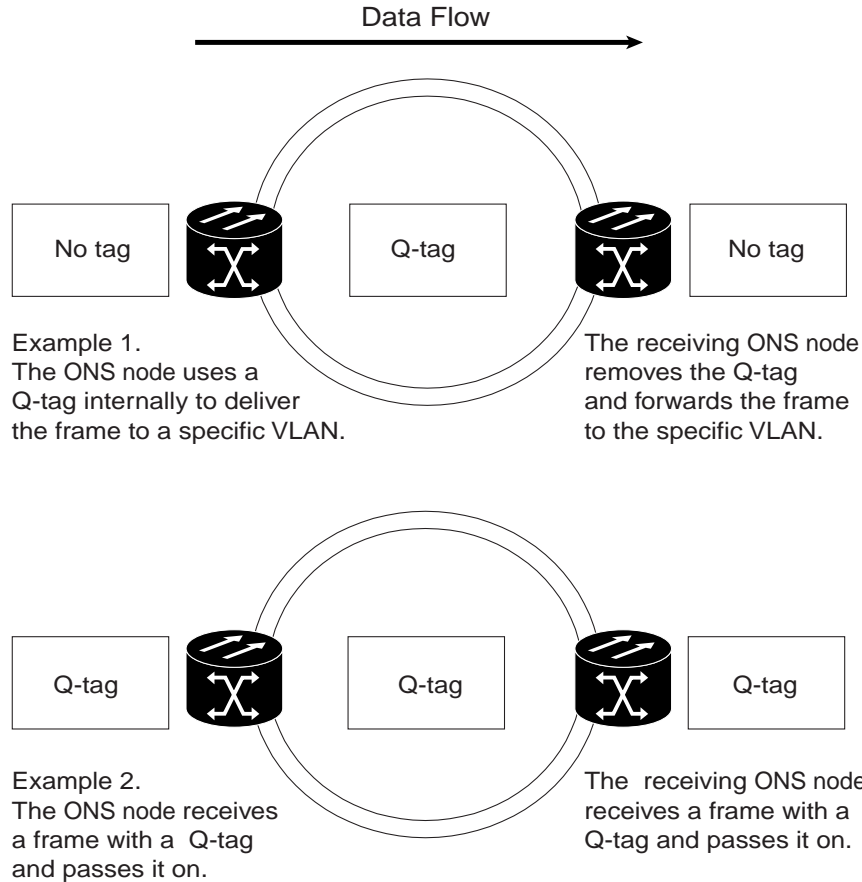
For more information about setting ports to Tagged and Untag, refer to the *Cisco ONS 15454 Procedure Guide*.



Caution

ONS 15454s propagate VLANs whenever a node appears on the network view of another node, regardless of whether the nodes are in the same SONET network or connect through DCC. For example, if two ONS 15454s without DCC connectivity belong to the same login node group, VLANs propagate between the two ONS 15454s. VLAN propagation happens even though the ONS 15454s do not belong to the same SONET ring.

Figure 16-13 Q-tag Moving Through VLAN



16.3.5 E-Series Priority Queuing (IEEE 802.1Q)

Networks without priority queuing handle all packets on a first-in-first-out basis. Priority queuing reduces the impact of network congestion by mapping Ethernet traffic to different priority levels. The ONS 15454 supports priority queuing. The ONS 15454 maps the eight priorities specified in IEEE 802.1Q to two queues, low priority and high priority (Table 16-1 on page 16-16). Q-tags carry priority queuing information through the network. (See [Figure 16-14 on page 16-16](#).)

The ONS 15454 uses a “leaky bucket” algorithm to establish a weighted priority (not a strict priority). A weighted priority gives high-priority packets greater access to bandwidth, but does not totally preempt low-priority packets. During periods of network congestion, roughly 70 percent of bandwidth goes to the high-priority queue and the remaining 30 percent goes to the low-priority queue. A network that is too congested will drop packets.



Note

IEEE 802.1Q was formerly IEEE 802.1P.



Note

E-Series cards in port-mapped mode and G-Series cards do not support priority queuing.



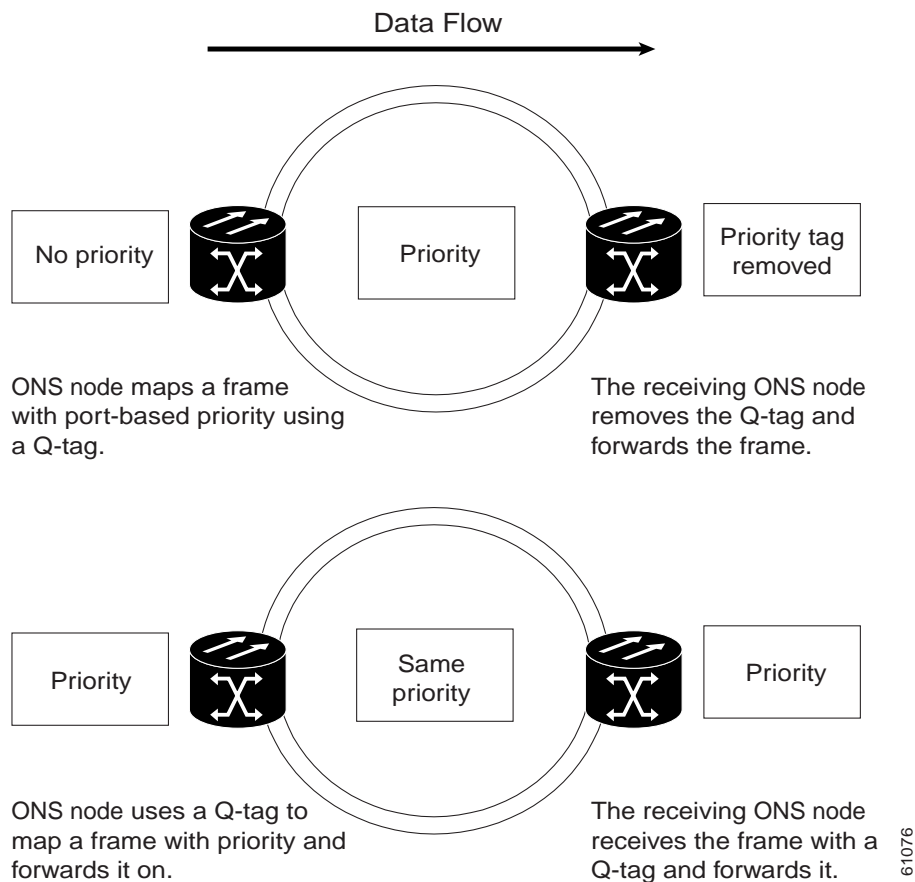
Note

E-Series port-mapped mode does not support priority queuing (IEEE 802.1Q).

Table 16-1 Priority Queuing

User Priority	Queue	Allocated Bandwidth
0,1,2,3	Low	30%
4,5,6,7	High	70%

Figure 16-14 Priority Queuing Process



16.3.6 E-Series Spanning Tree (IEEE 802.1D)

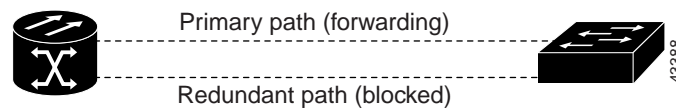
The ONS 15454 operates Spanning Tree Protocol (STP) according to IEEE 802.1D, when an Ethernet card is installed. The E-Series Card supports common STPs on a per circuit basis up to a total of eight STP instances. It does not support per-VLAN STP. In single-card mode, STP can be disabled or enabled on a per circuit basis during circuit creation. Disabling STP will preserve the number of available STP instances.

STP operates over all packet-switched ports including Ethernet and OC-N ports. On Ethernet ports, STP is enabled by default but may be disabled. A user can also disable or enable STP on a circuit-by-circuit basis on unstitched Ethernet cards in a point-to-point configuration. However, turning off STP protection on a circuit-by-circuit basis means that the ONS 15454 system is not protecting the Ethernet traffic on this circuit, and the Ethernet traffic must be protected by another mechanism in the Ethernet network. On OC-N interface ports, the ONS 15454 activates STP by default, and STP cannot be disabled.

The Ethernet card can enable STP on the Ethernet ports to create redundant paths to the attached Ethernet equipment. STP connects cards so that both equipment and facilities are protected against failure.

STP detects and eliminates network loops. When STP detects multiple paths between any two network hosts, STP blocks ports until only one path exists between any two network hosts (Figure 16-15). The single path eliminates possible bridge loops. This is crucial for shared packet rings, which naturally include a loop.

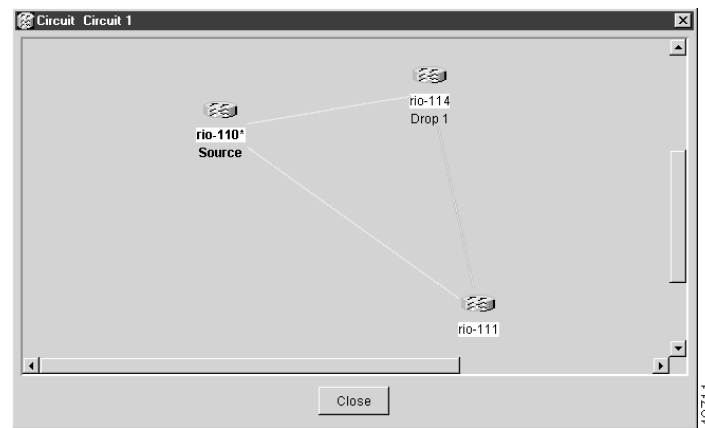
Figure 16-15 STP Blocked Path



To remove loops, STP defines a tree that spans all the switches in an extended network. STP forces certain redundant data paths into a standby (blocked) state. If one network segment in the STP becomes unreachable, the STP algorithm reconfigures the STP topology and reactivates the blocked path to reestablish the link. STP operation is transparent to end stations, which do not discriminate between connections to a single LAN segment or to a switched LAN with multiple segments. The ONS 15454 supports one STP instance per circuit and a maximum of eight STP instances per ONS 15454.

The Circuit window shows forwarding spans and blocked spans on the spanning tree map (Figure 16-16).

Figure 16-16 Spanning Tree Map on Circuit Window



Note

Green represents forwarding spans and purple represents blocked (protect) spans. If you have a packet ring configuration, at least one span should be purple.



Caution

Multiple circuits with STP protection enabled will incur blocking, if the circuits traverse a common card and uses the same VLAN.

**Note**

E-Series port-mapped mode does not support STP (IEEE 8021.D).

16.3.6.1 E-Series Multi-Instance Spanning Tree and VLANs

The ONS 15454 can operate multiple instances of STP to support VLANs in a looped topology. You can dedicate separate circuits across the SONET ring for different VLAN groups. Each circuit runs its own STP to maintain VLAN connectivity in a multiring environment.

16.3.6.2 Spanning Tree on a Circuit-by-Circuit Basis

You can also disable or enable STP on a circuit-by-circuit basis on single-card EtherSwitch E-Series cards in a point-to-point configuration. This feature allows customers to mix spanning tree protected circuits with unprotected circuits on the same card. It also allows two single-card EtherSwitch E-Series cards on the same node to form an intranode circuit.

16.3.6.3 E-Series Spanning Tree Parameters

Default STP parameters are appropriate for most situations (Table 16-2). Contact the Cisco Technical Assistance Center (Cisco TAC) before you change the default STP parameters. Contact information for Cisco TAC is listed in the Preface of this manual.

Table 16-2 Spanning Tree Parameters

Parameter	Description
BridgeID	ONS 15454 unique identifier that transmits the configuration bridge protocol data unit (BPDU); the bridge ID is a combination of the bridge priority and the ONS 15454 MAC address
TopoAge	Amount of time in seconds since the last topology change
TopoChanges	Number of times the STP topology has been changed since the node booted up
DesignatedRoot	Identifies the STP's designated root for a particular STP instance
RootCost	Identifies the total path cost to the designated root
RootPort	Port used to reach the root
MaxAge	Maximum time that received-protocol information is retained before it is discarded
HelloTime	Time interval, in seconds, between the transmission of configuration BPDUs by a bridge that is the spanning tree root or is attempting to become the spanning tree root
HoldTime	Minimum time period, in seconds, that elapses during the transmission of configuration information on a given port
ForwardDelay	Time spent by a port in the listening state and the learning state

16.3.6.4 E-Series Spanning Tree Configuration

To view the spanning tree configuration, at the node view click the **Provisioning > Etherbridge > Spanning Trees** tabs (Table 16-3).

Table 16-3 Spanning Tree Configuration

Column	Default Value	Value Range
Priority	32768	0–65535
Bridge max age	20 seconds	6–40 seconds
Bridge Hello Time	2 seconds	1–10 seconds
Bridge Forward Delay	15 seconds	4–30 seconds

16.4 G-Series Circuit Configurations

This section explains G-Series point-to-point circuits and manual cross-connects. Ethernet manual cross-connects allow you to cross connect individual Ethernet circuits to an STS channel on the ONS 15454 optical interface and also to bridge non-ONS SONET network segments.

16.4.1 G-Series Point-to-Point Ethernet Circuits

G-Series cards support point-to-point circuit configurations (Figure 16-17). Provisionable circuit sizes are STS 1, STS 3c, STS 6c, STS 9c, STS 12c, STS 24c, and STS 48c. Each Ethernet port maps to a unique STS circuit of the G-Series card.

Figure 16-17 G-Series Point-to-Point Circuit



The G-Series supports any combination of up to four circuits from the list of valid circuit sizes; however, the circuit sizes can add up to no more than 48 STSs.

Due to hardware constraints, the card imposes an additional restriction on the combinations of circuits that can be dropped onto a G-Series card. These restrictions are transparently enforced by the ONS 15454, and you do not need to keep track of restricted circuit combinations.

When a single STS-24c terminates on a card, the remaining circuits on that card can be another single STS-24c or any combination of circuits of STS-12c size or less that add up to no more than 12 STSs (that is a total of 36 STSs on the card).

If STS-24c circuits are not being dropped on the card, the full 48 STSs bandwidth can be used with no restrictions (for example using either a single STS-48c or 4 STS-12c circuits).



Note

The STS-24c restriction only applies when a single STS-24c circuit is dropped; therefore, you can easily minimize the impact of this restriction. Group the STS-24c circuits together on a card separate from circuits of other sizes. The grouped circuits can be dropped on other G-Series cards on the ONS 15454.



Note

The G-Series uses STS cross-connects only. No VT level cross-connects are used.

**Caution**

G-Series cards do not connect with E-Series cards.


16.4.2 G-Series Manual Cross-Connects

ONS 15454s require end-to-end CTC visibility between nodes for normal provisioning of Ethernet circuits. When other vendors' equipment sits between ONS 15454s, OSI/TARP-based equipment does not allow tunneling of the ONS 15454 TCP/IP-based DCC. To circumvent a lack of continuous DCC, the Ethernet circuit must be manually cross connected to an STS channel using the non-ONS network. Manual cross-connects allows an Ethernet circuit to run from ONS node to ONS node while utilizing the non-ONS network (Figure 16-18).

**Note**

In this chapter, “cross-connect” and “circuit” have the following meanings: Cross-connect refers to the connections that occur within a single ONS 15454 to allow a circuit to enter and exit an ONS 15454. Circuit refers to the series of connections from a traffic source (where traffic enters the ONS 15454 network) to the drop or destination (where traffic exits an ONS 15454 network).

Figure 16-18 G-Series Manual Cross-connects



16.5 E-Series Circuit Configurations

Ethernet circuits can link ONS nodes through point-to-point (straight), shared packet ring, or hub and spoke configurations. Two nodes usually connect with a point-to-point configuration. More than two nodes usually connect with a shared packet ring configuration or a hub-and-spoke configuration. Ethernet manual cross-connects allow you to cross connect individual Ethernet circuits to an STS channel on the ONS 15454 optical interface and also to bridge non-ONS SONET network segments. For step-by-step procedures to configure E-Series circuits, see the “Create Circuits and VT Tunnels” chapter of the *Cisco ONS 15454 Procedure Guide*.

16.5.1 E-Series Circuit Protection

Different combinations of E-Series circuit configurations and SONET network topologies offer different levels of E-Series circuit protection. Table 16-4 details the available protection.

Table 16-4 Protection for E-Series Circuit Configurations

Configuration	Path Protection	BLSR	1 + 1
Point-to-Point Multicard Etherswitch	None	SONET	SONET
Point-to-Point Single-Card Etherswitch	SONET	SONET	SONET
Point-to-Point Port-mapped Mode	SONET	SONET	SONET
Shared Packet Ring Multicard Etherswitch	STP	SONET	SONET
Common Control Card Switch	STP	STP	STP

**Note**

Before making Ethernet connections, choose an STS-1, STS-3c, STS-6c, or STS-12c circuit size.

**Note**

To make an STS-12c Ethernet circuit, Ethernet cards must be configured in Single-card EtherSwitch or port-mapped mode. Multicard mode does not support STS-12c Ethernet circuits.

16.5.2 Port-mapped Mode and Single-card EtherSwitch Circuit Scenarios

Seven scenarios exist for provisioning circuits on an E-Series card in single-card EtherSwitch or port-mapped mode:

1. STS 12c
2. STS 6c + STS 6c
3. STS 6c + STS 3c + STS 3c
4. STS 6c + 6 STS-1s
5. STS 3c + STS 3c + STS 3c + STS 3c
6. STS 3c + STS 3c + 6 STS-1s
7. 12 STS-1s

**Note**

When configuring scenarios 3 and 4, the STS 6c must be provisioned before the smaller STS circuits.

16.5.3 ONS 15454 E-Series and ONS 15327 EtherSwitch Circuit Combinations

Table 16-5 shows the Ethernet circuit combinations available in ONS 15454 E-Series cards and ONS 15327 E-Series cards.

Table 16-5 ONS 15454 and ONS 15327 E-Series Ethernet Circuit Combinations

15327 E-Series Port-mapped and Single-card EtherSwitch	15327 E-Series Multicard EtherSwitch	15454 E-Series Port-mapped and Single-card EtherSwitch	15454 E-Series Multicard EtherSwitch
Six STS-1s	Three STS-1s	One STS 12c	Six STS-1s
Two STS 3cs	One STS 3c	Two STS 6cs	Two STS 3cs
One STS 6c	—	One STS 6c and two STS 3cs	One STS 6c
One STS 12c	—	One STS 6c and six STS-1s	—
—	—	Four STS 3cs	—
—	—	Two STS 3cs and six STS-1s	—
—	—	Twelve STS-1s	—

16.5.4 E-Series Point-to-Point Ethernet Circuits

The ONS 15454 can set up a point-to-point (straight) Ethernet circuit as single-card, port-mapped, or multicard circuit. Multicard EtherSwitch limits bandwidth to STS-6c of bandwidth between two Ethernet circuit points, but allows adding nodes and cards and making a shared packet ring (Figure 16-19 on page 16-22). Single-card EtherSwitch and port-mapped mode allows a full STS-12c of bandwidth between two Ethernet circuit endpoints (Figure 16-20 on page 16-23).

Figure 16-19 Multicard EtherSwitch Point-to-point Circuit





Figure 16-20 Single-card EtherSwitch or Port-mapped Point-to-Point Circuit




Note A port-mapped point-to-point circuit does not contain a VLAN.

16.5.5 E-Series Shared Packet Ring Ethernet Circuits

A shared packet ring allows additional nodes, besides the source and destination nodes, access to an Ethernet STS circuit. The E-Series card ports on the additional nodes can share the circuit's VLAN and bandwidth. Figure 16-21 on page 16-23 illustrates a shared packet ring. Your network architecture might differ from the example.


Figure 16-21 Shared Packet Ring Ethernet Circuit



16.5.6 E-Series Hub and Spoke Ethernet Circuit Provisioning

The hub and spoke configuration connects point-to-point circuits (the spokes) to an aggregation point (the hub). In many cases, the hub links to a high-speed connection and the spokes are Ethernet cards. Figure 16-22 illustrates a hub and spoke ring. Your network architecture may differ from the example.

Figure 16-22 Hub and Spoke Ethernet Circuit



16.5.7 E-Series Ethernet Manual Cross-Connects

ONS 15454s require end-to-end CTC visibility between nodes for normal provisioning of Ethernet circuits. When other vendors' equipment sits between ONS 15454s, OSI/TARP-based equipment does not allow tunneling of the ONS 15454 TCP/IP-based DCC. To circumvent this lack of continuous DCC, the Ethernet circuit must be manually cross connected to an STS channel using the non-ONS network. The manual cross-connect allows an Ethernet circuit to run from ONS node to ONS node utilizing the non-ONS network.



Note

In this chapter, “cross-connect” and “circuit” have the following meanings: cross-connect refers to the connections that occur within a single ONS 15454 to allow a circuit to enter and exit an ONS 15454. Circuit refers to the series of connections from a traffic source (where traffic enters the ONS 15454 network) to the drop or destination (where traffic exits an ONS 15454 network).

16.6 Remote Monitoring Specification Alarm Thresholds

The ONS 15454 features remote monitoring (RMON) that allows network operators to monitor the health of the network with a network management system (NMS).

One of the ONS 15454's RMON MIBs is the Alarm group, which consists of the alarmTable. An NMS uses the alarmTable to find the alarm-causing thresholds for network performance. The thresholds apply to the current 15-minute interval and the current 24-hour interval. RMON monitors several variables, such as Ethernet collisions, and triggers an event when the variable crosses a threshold during that time

interval. For example, if a threshold is set at 1000 collisions and 1001 collisions occur during the 15-minute interval, an event triggers. CTC allows you to provision these thresholds for Ethernet statistics.



Note

Table 16-6 on page 16-25 define the variables you can provision in CTC. For example, to set the collision threshold, choose etherStatsCollisions from the Variable menu.

Table 16-6 Ethernet Threshold Variables (MIBs)

Variable	Definition
ifInOctets	Total number of octets received on the interface, including framing octets
ifInUcastPkts	Total number of unicast packets delivered to an appropriate protocol
ifInMulticastPkts	Number of multicast frames received error free (not supported by E-Series)
ifInBroadcastPkts	Number of packets, delivered by this sublayer to a higher (sub)layer, which were addressed to a broadcast address at this sublayer (not supported by E-Series)
ifInDiscards	Number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol (not supported by E-Series)
ifInErrors	Number of inbound packets discarded because they contain errors
ifOutOctets	Total number of transmitted octets, including framing packets
ifOutUcastPkts	Total number of unicast packets requested to transmit to a single address
ifOutMulticastPkts	Number of multicast frames transmitted error free (not supported by E-Series)
ifOutBroadcastPkts	Total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sublayer, including those that were discarded or not sent (not supported by E-Series)
ifOutDiscards	Number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted (not supported by E-Series)
dot3statsAlignmentErrors	Number of frames with an alignment error, that is, the length is not an integral number of octets and the frame cannot pass the Frame Check Sequence (FCS) test
dot3StatsFCSErrors	Number of frames with framecheck errors, that is, there is an integral number of octets, but an incorrect FCS
dot3StatsSingleCollisionFrames	Number of successfully transmitted frames that had exactly one collision
dot3StatsMutlipleCollisionFrame	Number of successfully transmitted frames that had multiple collisions

Table 16-6 Ethernet Threshold Variables (MIBs) (continued)

Variable	Definition
dot3StatsDeferredTransmissions	Number of times the first transmission was delayed because the medium was busy
dot3StatsExcessiveCollision	Number of frames where transmissions failed because of excessive collisions
dot3StatsLateCollision	Number of times that a collision was detected later than 64 octets into the transmission (also added into collision count)
dot3StatsFrameTooLong	Number of received frames that were larger than the maximum size permitted
dot3StatsCarrierSenseErrors	Number of transmission errors on a particular interface that are not otherwise counted (not supported by E-Series)
dot3StatsSQETestErrors	Number of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface (not supported by E-Series)
etherStatsJabbers	Total number of Octets of data (including bad packets) received on the network
etherStatsUndersizePkts	Number of packets received with a length less than 64 octets
etherStatsFragments	Total number of packets that are not an integral number of octets or have a bad FCS, and that are less than 64 octets long
etherStatsOversizePkts	Total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed
etherStatsOctets	Total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets)
etherStatsPkts64Octets	Total number of packets received (including error packets) that were 64 octets in length
etherStatsPkts65to127Octets	Total number of packets received (including error packets) that were 65–172 octets in length
etherStatsPkts128to255Octets	Total number of packets received (including error packets) that were 128–255 octets in length
etherStatsPkts256to511Octets	Total number of packets received (including error packets) that were 256–511 octets in length
etherStatsPkts512to1023Octets	Total number of packets received (including error packets) that were 512–1023 octets in length
etherStatsPkts1024to1518Octets	Total number of packets received (including error packets) that were 1024–1518 octets in length
etherStatsJabbers	Total number of packets longer than 1518 octets that were not an integral number of octets or had a bad FCS
etherStatsCollisions	Best estimate of the total number of collisions on this segment
etherStatsCollisionFrames	Best estimate of the total number of frame collisions on this segment

Table 16-6 Ethernet Threshold Variables (MIBs) (continued)

Variable	Definition
etherStatsCRCAlignErrors	Total number of packets with a length between 64 and 1518 octets, inclusive, that had a bad FCS or were not an integral number of octets in length
receivePauseFrames	Number of received 802.x pause frames (not supported by E-Series)
transmitPauseFrames	Number of transmitted 802.x pause frames (not supported by E-Series)
receivePktsDroppedInternalCongestion	Number of received frames dropped because of frame buffer overflow and other reasons (not supported by E-Series)
transmitPktsDroppedInternalCongestion	Number of frames dropped in the transmit direction because of frame buffer overflow and other reasons (not supported by E-Series)
txTotalPkts	Total number of transmit packets (not supported by E-Series)
rxTotalPkts	Total number of receive packets (not supported by E-Series)

