



Ethernet Operation

The Cisco ONS 15454 SDH integrates Ethernet into a SDH time-division multiplexing (TDM) platform. The ONS 15454 SDH 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 SONET/SDH 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 Low-Order Tunnels” chapter of the *Cisco ONS 15454 SDH Procedure Guide*.

Chapter topics include:

- [13.1 G-Series Application, page 13-1](#)
- [13.2 E-Series Application, page 13-5](#)
- [13.3 G-Series Circuit Configurations, page 13-15](#)
- [13.4 E-Series Circuit Configurations, page 13-16](#)
- [13.5 Remote Monitoring Specification Alarm Thresholds, page 13-19](#)

13.1 G-Series Application

The G-Series cards (G1000-4/G1K-4) reliably transport Ethernet and IP data across an SDH backbone. The G-Series card maps up to four Gigabit Ethernet interfaces onto a SDH transport network and provides scalable and provisionable transport bandwidth at signal levels up to VC4-16C 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 (including hitless (< 50 ms) performance under software upgrades and all types of SDH/SONET equipment protection switches)
- Hitless reprovisioning
- Support of Gigabit Ethernet traffic at full line rate
- Full TL1-based provisioning capability

The G-Series card allows an Ethernet private line service to be provisioned and managed very much like a traditional SDH or SONET 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 STM circuit. You can independently map the four ports on the G-Series card to any combination of VC4, VC4-2c, VC4-3c, VC4-4c, VC4-8c, and VC4-16C circuit sizes, provided the sum of the circuit sizes that terminate on a card do not exceed VC4-16C.

To support a Gigabit Ethernet port at full line rate, an STM circuit with a capacity greater or equal to 1 Gbps (bidirectional 2 Gbps) is needed. A VC4-8c 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 SDH J1 Path Trace byte in the same manner as ONS 15454 SDH STM-N cards. For more information on the J1 Path Trace, refer to the “[J1 Path Trace](#)” section on page 9-14.

**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.


13.1.1 G1K-4 and G1000-4 Comparison

The G1K-4 and the G1000-4 cards comprise the ONS 15454 SDH G-Series and are hardware equivalents. Software Release 4.0 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.

13.1.2 G-Series Example

Figure 13-1 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 SDH point-to-point circuit to the Gigabit Ethernet port of another high-end router.

Figure 13-1 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 SDH payload by multiplexing the payload onto a SDH STM-N card. When the SDH 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 SDH. 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 SDH 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.

13.1.3 IEEE 802.3z Flow Control and Frame Buffering

The G-Series supports 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 SDH uses 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 13-1 on page 13-2](#) illustrates pause frames being sent and received by ONS 15454 SDHs and attached switches.

With Software R4.0 and later, the G-Series card has symmetric flow control and 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 STM 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 SDH circuit assigned to the G-Series port might be only VC4-4c (622.08 Mbps). In this example, the ONS 15454 SDH 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 (VC4-8c) is efficient because frame loss can be controlled to a large extent.

**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.

13.1.4 Ethernet Link Integrity Support

The G-Series supports end-to-end Ethernet link integrity. 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 13-2 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 13-2](#), 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.

13.1.5 Gigabit EtherChannel/IEEE 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 13-3](#) illustrates G-Series GEC support.

Figure 13-3 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 SDH network, the ONS 15454 SDH 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.

13.2 E-Series Application

The E-Series cards incorporate Layer 2 switching, whereas the G-Series card is a straight mapper card. E-Series cards in multi-card EtherSwitch Group or single-card EtherSwitch support virtual local area networks (VLANs), IEEE 802.1Q, STP, and IEEE 802.1D. The E-Series card in port-mapped mode configures the E-Series card to be a mapper card and disables the Layer 2 functions. An ONS 15454 SDH holds a maximum of ten Ethernet cards, and you can insert Ethernet cards in any multipurpose slot.

13.2.1 E-Series Modes

An E-Series card operates in one of three modes: Multi-card EtherSwitch Group, Single-card EtherSwitch, or Port-mapped. Within an ONS 15454 SDH 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.”

13.2.1.1 E-Series Multicard EtherSwitch Group

Multicard EtherSwitch Group (stitched) provisions two or more Ethernet cards to act as a single Layer 2 switch with valid circuit sizes of VC-4-2c or VC-4. You can provision up to VC-4-4c to all cards in EtherSwitch group (up to VC-4-2c in each direction). You must reserve an equal amount of bandwidth to create a stitch to another Ethernet card. [Figure 13-4 on page 13-6](#) illustrates a multicard EtherSwitch configuration.

Figure 13-4 A Multicard EtherSwitch Configuration


**Caution**

Whenever you terminate two VC4-2c multicard EtherSwitch circuits on an Ethernet card and later delete the first circuit. Also delete the remaining VC4-2c circuit, before you provision an VC4 circuit to the card. If you attempt to create an VC4 circuit after only deleting the first VC4-2c circuit, the VC4 circuit will not work, but no alarms will indicate this condition. To avoid this situation, delete the second VC4-2c before creating an VC4 circuit.

13.2.1.2 E-Series Single-Card EtherSwitch

Single-card EtherSwitch allows each Ethernet card to remain a single switching entity within the ONS 15454 SDH shelf. This option allows VC4-4c worth of bandwidth between two Ethernet circuit endpoints. Figure 13-5 illustrates a single-card EtherSwitch configuration.

Figure 13-5 A Single-card EtherSwitch Configuration

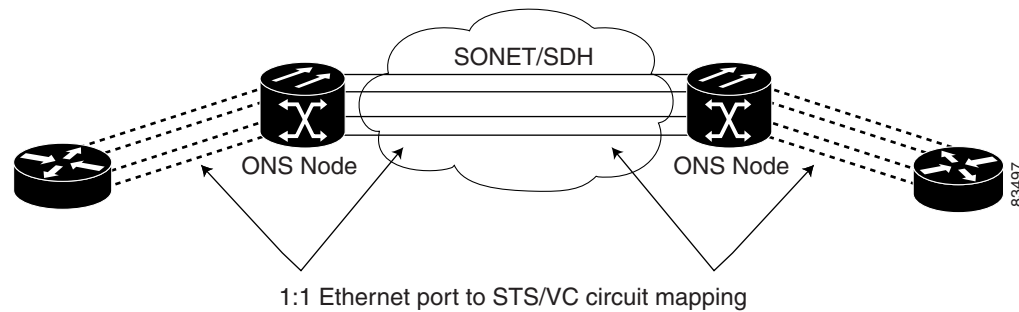


13.2.1.3 Port-Mapped (Linear Mapper)

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 STM circuits. (See Figure 13-6 on page 13-7.) Port-mapped mode ensures Layer 1 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 not line rate because the E1000-2-G has a maximum bandwidth of VC4-4c. Ethernet frame sizes up to 1522 bytes are also supported, which allows transport of 802.1Q tagged frames. The larger maximum frame size of Q-in-Q frames, 802.1Q in 802.1Q wrapped frames, are not supported.

Figure 13-6 E-Series Mapping Ethernet Ports To SDH STM circuits



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

Port-mapped mode does not support VLANs in the same manner as multcard and single-card mode. The ports of E-Series cards in multcard 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 STM-N circuits between any two E-Series cards. Port-mapped mode does not allow an E-Series cards to connect to the ML-Series or G-Series cards.

13.2.2 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 STM-N 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 SDH circuit assigned to the E-Series port in port-mapped mode is a maximum of VC4-4c (622.08 Mbps). In this scenario, the ONS 15454 SDH 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.

13.2.3 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 SDH. The definition of VLAN ports includes all Ethernet and packet-switched SDH port types. All VLAN IP address discovery, flooding, and forwarding is limited to these ports.

The ONS 15454 SDH IEEE 802.1Q-based VLAN mechanism provides logical isolation of subscriber LAN traffic over a common SDH 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.

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

E-Series cards in single-card and multicard mode support IEEE 802.1Q. IEEE 802.1Q allows the same physical port to host multiple IEEE 802.1Q VLANs. Each IEEE 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 SDH 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 SDH Ethernet port does not support IEEE 802.1Q, the ONS 15454 SDH uses Q-tags internally only. The ONS 15454 SDH associates these Q-tags with specific ports.

With Ethernet devices that do not support IEEE 802.1Q, the ONS 15454 SDH 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 IEEE 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 13-7 on page 13-9](#) illustrates Q-tag use only within an ONS network.

The ONS 15454 SDH 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 SDH 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 13-7 on page 13-9](#) illustrates the handling of packets that both enter and exit the ONS network with a Q-tag.

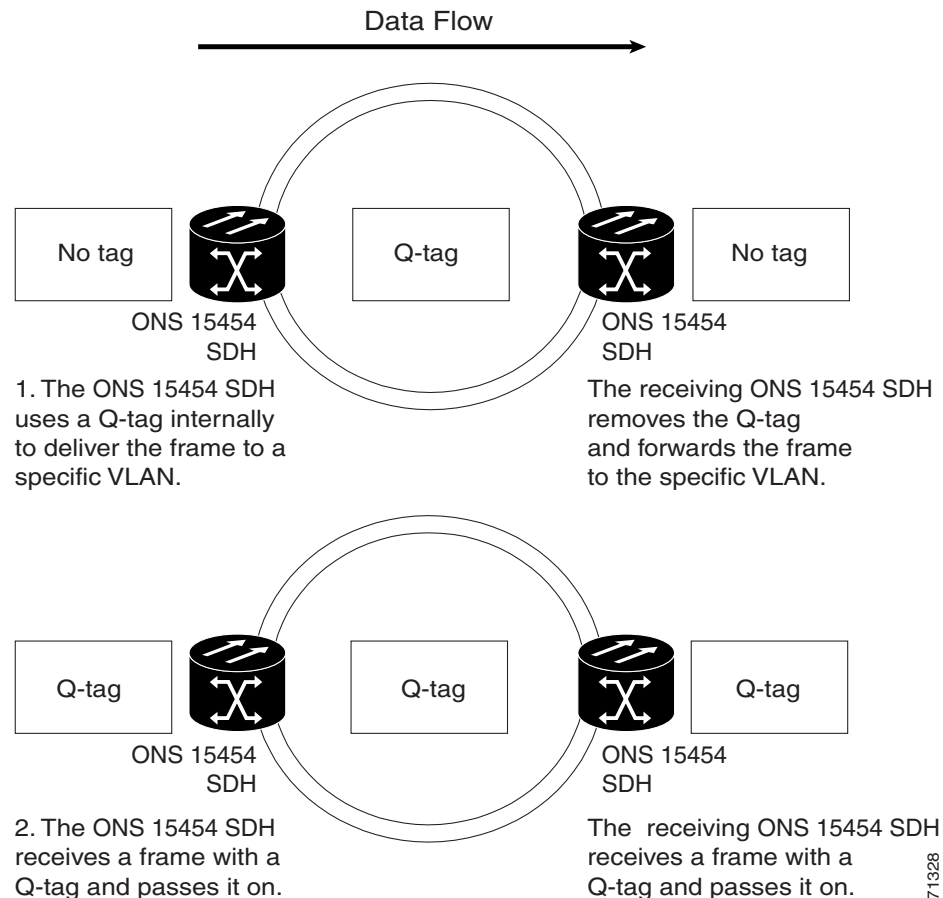
For more information about setting ports to Tagged and Untag, refer to the *Cisco ONS 15454 SDH Procedure Guide*.



Caution

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

Figure 13-7 Q-tag Moving Through VLAN



13.2.5 E-Series Priority Queuing (IEEE 802.1Q)

Networks without priority queuing handle all packets on a first-in, first-out (FIFO) basis. Priority queuing reduces the impact of network congestion by mapping Ethernet traffic to different priority levels. The ONS 15454 SDH supports priority queuing. The ONS 15454 SDH maps the eight priorities specified in IEEE 802.1Q to two queues, low priority and high priority number ([Table 13-1 on page 13-10](#)). Q-tags carry priority queuing information through the network. (See [Figure 13-8 on page 13-11](#).)

The ONS 15454 SDH 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.

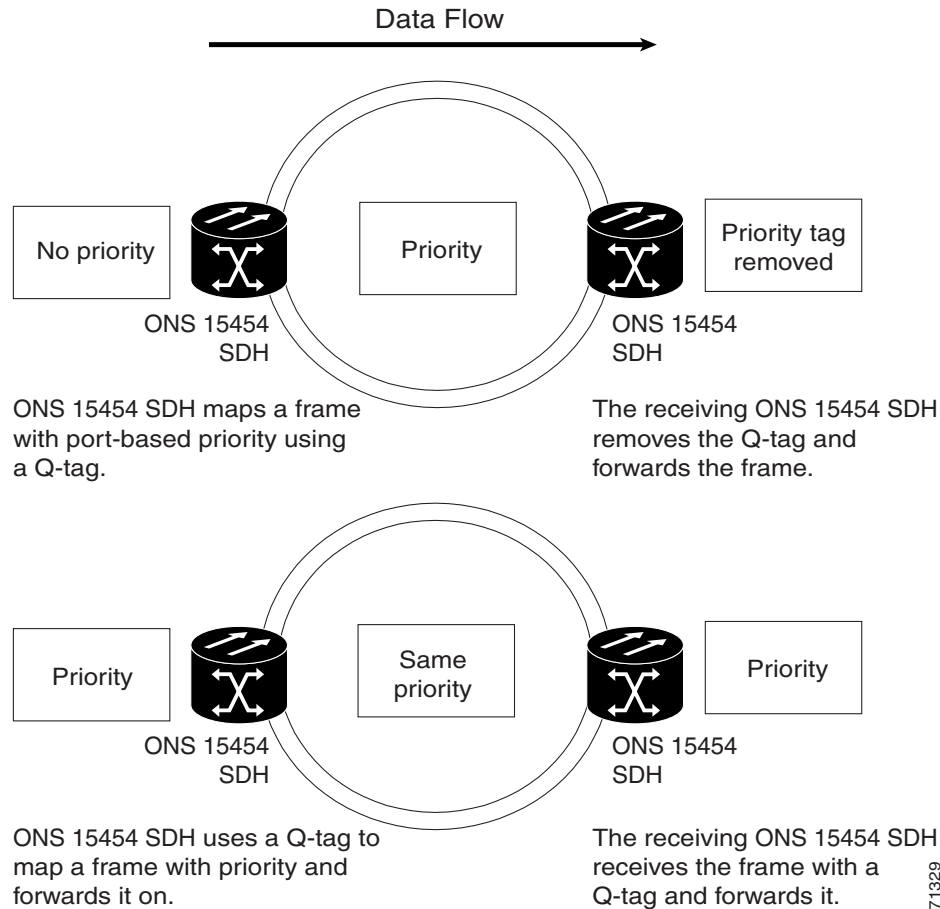


Note E-Series port-mapped mode does not support priority queuing (IEEE 802.1Q).

Table 13-1 Priority Queuing

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

Figure 13-8 Priority Queuing Process



13.2.6 E-Series Spanning Tree (IEEE 802.1D)


The Cisco ONS 15454 SDH 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 STM-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 SDH 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 STM-N interface ports, the ONS 15454 SDH 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 13-9). The single path eliminates possible bridge loops. This is crucial for shared packet rings, which naturally include a loop.

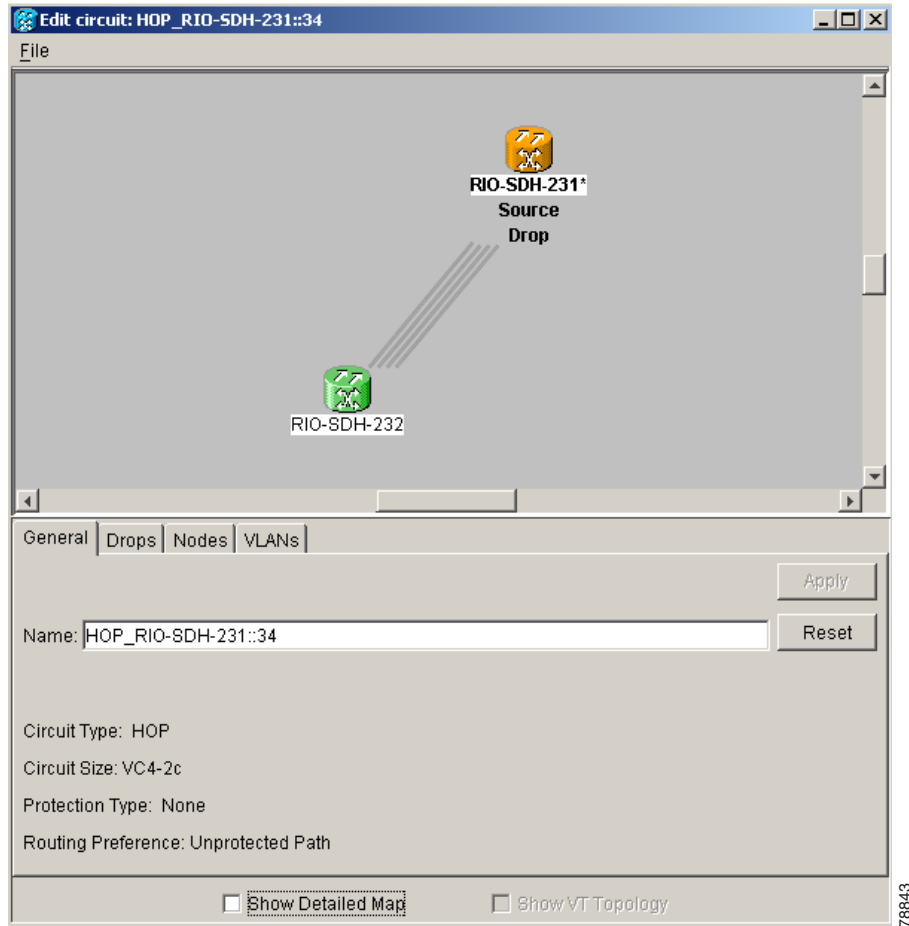
Figure 13-9 An STP Blocked Path



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 SDH supports one STP instance per circuit and a maximum of eight STP instances per ONS 15454 SDH.

The Circuit window shows forwarding spans and blocked spans on the spanning tree map. (See Figure 13-10 on page 13-13.)

Figure 13-10 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).

13.2.6.1 E-Series Multi-Instance Spanning Tree and VLANs

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

13.2.6.2 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.

13.2.6.3 E-Series Spanning Tree Parameters

Default STP parameters are appropriate for most situations. (See [Table 13-2](#).) Contact the Cisco Technical Assistance Center (TAC) at 1-877-323-7368 before you change the default STP parameters.

Table 13-2 Spanning Tree Parameters

Parameter	Description
BridgeID	ONS 15454 SDH 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 SDH 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	STP's designated root for a particular STP instance
RootCost	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

13.2.6.4 E-Series Spanning Tree Configuration

To view the spanning tree configuration, at the node view click the **Provisioning > Etherbridge > Spanning Trees** tabs. (See [Table 13-3](#).)

Table 13-3 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


13.3 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 STM-N channel on the ONS 15454 SDH optical interface and also to bridge non-ONS SDH network segments.

13.3.1 G-Series Point-to-Point Ethernet Circuits

G-Series cards support point-to-point circuit configurations. (See [Figure 13-11](#).) Provisionable circuit sizes are VC4, VC4-2c, VC4-3c, VC4-4c, VC4-8c, and VC4-16C. Each Ethernet port maps to a unique STM-N circuit of the G-Series card.

Figure 13-11 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 VC4-16.

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 SDH, and you do not need to keep track of restricted circuit combinations.

When a single VC4-8c terminates on a card, the remaining circuits on that card can be another single VC4-8c or any combination of circuits of VC4-4c size or less that add up to no more than VC4-4c (that is a total of 12 VC4s on the card).

If VC4-8c circuits are not being dropped on the card, the full VC4-16 bandwidth can be used with no restrictions (for example, using either a single VC4-16C or 4 VC4-4c circuits).



Note

The VC4-8c restriction only applies when a single VC4-8c circuit is dropped; therefore, you can easily minimize the impact of this restriction. Group the VC4-8c 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 SDH.



Note

All SDH side STM-N circuits must be adjacent to one another.



Caution

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

13.3.2 G-Series Manual Cross-Connects


ONS 15454 SDHs require end-to-end CTC visibility between nodes for normal provisioning of Ethernet circuits. When other vendors' equipment sits between ONS 15454 SDHs, OSI/TARP-based equipment does not allow tunneling of the ONS 15454 SDH TCP/IP-based DCC. To circumvent a lack of

continuous DCC, the Ethernet circuit must be manually cross connected to an STM 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. (See [Figure 13-12](#).)

**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 SDH to allow a circuit to enter and exit an ONS 15454 SDH. Circuit refers to the series of connections from a traffic source (where traffic enters the ONS 15454 SDH network) to the drop or destination (where traffic exits an ONS 15454 SDH network).

Figure 13-12 G-Series Manual Cross-connects



13.4 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 STM channel on the ONS 15454 SDH optical interface and also to bridge non-ONS SDH network segments. For step-by-step procedures to configure E-Series circuits, see the “Create Circuits and Low-Order Tunnels” chapter of the *Cisco ONS 15454 SDH Procedure Guide*.

13.4.1 Port-Mapped Mode and Single-card EtherSwitch Circuit Scenarios

Four scenarios exist for provisioning maximum single-card EtherSwitch or port-mapped mode bandwidth:

1. VC4-4c
2. VC4-2c + VC4-2c
3. VC4-2c + VC4 + VC4
4. VC4 + VC4 + VC4 + VC4

**Note**

When configuring scenario 3, the VC4-2c must be provisioned before either of the VC4 circuits. When configuring scenarios 3 and 4, the STM 6c must be provisioned before the smaller STM circuits.

13.4.2 E-Series Point-to-Point Ethernet Circuits

The ONS 15454 SDH can set up a point-to-point (straight) Ethernet circuit as single-card, port-mapped or multcard circuit. [Figure 13-13](#) shows a multcard EtherSwitch, which allows adding nodes and cards and making a shared packet ring. [Figure 13-14](#) shows a single-card EtherSwitch and port-mapped mode allows a full VC4-4c of bandwidth between two Ethernet circuit endpoints.

Figure 13-13 Multicard EtherSwitch Point-to-Point Circuit





Figure 13-14 Single-card EtherSwitch or Port-Mapped Point-to-Point Circuit




Note

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

13.4.3 E-Series Shared Packet Ring Ethernet Circuits

A shared packet ring allows additional nodes, besides the source and destination nodes, access to an Ethernet STM circuit. The E-Series card ports on the additional nodes can share the circuit's VLAN and bandwidth. [Figure 13-15](#) illustrates a shared packet ring. Your network architecture may differ from the example.


Figure 13-15 Shared Packet Ring Ethernet Circuit



13.4.4 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 13-16 on page 13-19](#) illustrates a hub-and-spoke ring. Your network architecture may differ from the example.

Figure 13-16 Hub-and-Spoke Ethernet Circuit



13.4.5 E-Series Ethernet Manual Cross-Connects

ONS 15454 SDHs require end-to-end CTC visibility between nodes for normal provisioning of Ethernet circuits. When other vendors' equipment sits between ONS 15454 SDHs, OSI/TARP-based equipment does not allow tunneling of the ONS 15454 SDH TCP/IP-based DCC. To circumvent this lack of continuous DCC, the Ethernet circuit must be manually cross connected to an STM 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 SDH to allow a circuit to enter and exit an ONS 15454 SDH. Circuit refers to the series of connections from a traffic source (where traffic enters the ONS 15454 SDH network) to the drop or destination (where traffic exits an ONS 15454 SDH network).

13.5 Remote Monitoring Specification Alarm Thresholds

The ONS 15454 SDH 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 SDH'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 13-4 on page 13-20](#) defines the variables you can provision in CTC. For example, to set the collision threshold, choose etherStatsCollisions from the Variable menu.

Table 13-4 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 sub-layer to a higher (sub-)layer, which were addressed to a broadcast address at this sub-layer (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 sub-layer, 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, i.e., 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, i.e., there is an integral number of octets, but an incorrect Frame Check Sequence (FCS)
dot3StatsSingleCollisionFrames	Number of successfully transmitted frames that had exactly one collision
dot3StatsMutlipleCollisionFrame	Number of successfully transmitted frames that had multiple collisions
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

Table 13-4 Ethernet Threshold Variables (MIBs) (continued)

Variable	Definition
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
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)

Table 13-4 Ethernet Threshold Variables (MIBs) (continued)

Variable	Definition
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)