CHAPTER 17

CE-Series Ethernet Cards

This chapter describes the operation of the CE-100T-8 (Carrier Ethernet) card supported on the Cisco ONS 15310-CL and ONS 15310-MA (15310-CE-100T-8) and the CE-MR-6 card supported on the Cisco ONS 15310-MA (15310-CE-MR-6). A CE-100T-8 card is also supported on the ONS 15454 (15454-CE-100T-8). Provisioning is done through Cisco Transport Controller (CTC) or Transaction Language One (TL1). Cisco IOS is not supported on the CE-100T-8 card.

For Ethernet card specifications, refer to the *Cisco ONS 15454 Reference Manual*. For step-by-step Ethernet card circuit configuration procedures and hard-reset and soft-reset procedures, refer to the *Cisco ONS 15454 Procedure Guide*. For TL1 provisioning commands, refer to the *Cisco ONS SONET TL1 Command Guide*. For specific details on ONS 15310-CL and ONS 15310-MA Ethernet card interoperability with other ONS platforms, refer to the "POS on ONS Ethernet Cards" chapter of the *Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide*.

Chapter topics include:

CE-100T-8 Ethernet Card

This section describes the operation of the CE-100T-8 (Carrier Ethernet) card supported on the ONS 15310-CL and ONS 15310-MA.

Provisioning is done through Cisco Transport Controller (CTC) or Transaction Language One (TL1). Cisco IOS is not supported on the CE-100T-8 card.

For Ethernet card specifications, refer to the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Reference Manual*. For step-by-step Ethernet card circuit configuration procedures, refer to the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Procedure Guide*. For TL1 provisioning commands, refer to the *Cisco ONS SONET TL1 Command Guide* .

Section topics include:

CE-100T-8 Overview

The CE-100T-8 is a Layer 1 mapper card with eight 10/100 Ethernet ports. It maps each port to a unique SONET circuit in a point-to-point configuration. Figure 17-1 illustrates a sample CE-100T-8 application. In this example, data traffic from the Fast Ethernet port of a switch travels across the point-to-point circuit to the Fast Ethernet port of another switch.

Figure 17-1 CE-100T-8 Point-to-Point Circuit

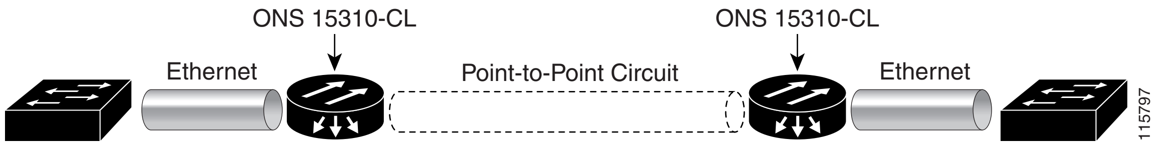

The CE-100T-8 cards allow you to provision and manage an Ethernet private line service like a traditional SONET line. CE-100T-8 card applications include providing Ethernet private line services and high-availability transport. It supports ITU-T G.707 and Telcordia GR-253 based standards for SONET.

The CE-100T-8 offers full TL1-based provisioning capability. Refer to the *Cisco ONS SONET TL1 Command Guide* for CE-100T-8 TL1 provisioning commands.

CE-100T-8 Ethernet Features

The CE-100T-8 card has eight front-end Ethernet ports which use standard RJ-45 connectors for 10BASE-T Ethernet/100BASE-TX Ethernet media. Ethernet Ports 1 through 8 each map to a POS port with a corresponding number. The console port on the CE-100T-8 card is not functional.

The CE-100T-8 cards forward valid Ethernet frames unmodified over the SONET network. Information in the headers is not affected by the encapsulation and transport. For example, included IEEE 802.1Q information will travel through the process unaffected.

The ONS 15454 CE-100T-8 and the ONS 15310 CE-100T-8 support maximum Ethernet frame sizes of 1600 bytes including the CRC. The MTU size is not configureable and is set at a 1500 byte maximum (standard Ethernet MTU). Baby giant frames in which the standard Ethernet frame is augmented by IEEE 802.1 Q tags or MPLS tags are also supported. Full Jumbo frames (9000 byte maximum) are not supported.

The CE-100T-8 cards discard certain types of erroneous Ethernet frames rather than transport them over SONET. Erroneous Ethernet frames include corrupted frames with cyclic redundancy check (CRC) errors and undersized frames that do not conform to the minimum 64-byte length Ethernet standard.

Note Many Ethernet attributes are also available through the network element default feature. For more information on NE defaults, refer to the "Network Element Defaults" appendix in the *Cisco ONS 15454 Reference Manual*.

Autonegotiation, Flow Control, and Frame Buffering

On the CE-100T-8 card, Ethernet link autonegotiation is on by default when the speed or duplex of the port is set to auto. The user can also set the link speed, duplex, selective autonegotiation, and flow control manually under the card-level Provisioning tab of CTC.

The CE-100T-8 card supports selective autonegotiation on the Ethernet ports. If selective autonegotiation is enabled, the port attempts to autonegotiate only to a specific speed and duplex. The link will come up if both the speed and duplex of the attached autonegotiating device matches that of the port. You cannot enable selective autonegotiation if either the speed or duplex of the port is set to auto.

The CE-100T-8 card supports IEEE 802.3x flow control and frame buffering to reduce data traffic congestion. Flow control is on by default.

To prevent over-subscription, buffer memory is available for each port. When the buffer memory on the Ethernet port nears capacity, the CE-100T-8 card uses IEEE 802.3x flow control to transmit a pause frame to the attached Ethernet device. Flow control and autonegotiation frames are local to the Fast Ethernet interfaces and the attached Ethernet devices. These frames do not continue through the POS ports.

The CE-100T-8 card has symmetric flow control and proposes symmetric flow control when autonegotiating flow control with attached Ethernet devices. Symmetric flow control allows the CE-100T-8 cards to respond to pause frames sent from external devices and to send pause frames to external devices.

The pause frame instructs the source to stop sending packets for a specific period of time. The sending station waits the requested amount of time before sending more data. Figure 17-2 illustrates pause frames being sent and received by CE-100T-8 cards and attached switches.

Figure 17-2 Flow Control

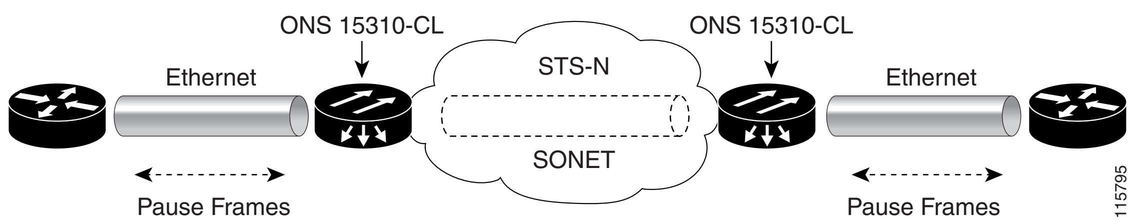

This flow-control mechanism matches the sending and receiving device throughput to that of the bandwidth of the STS circuit. For example, a router might transmit to the Ethernet port on the CE-100T-8 card. This particular data rate might occasionally exceed 51.84 Mbps, but the SONET circuit assigned to the CE-100T-8 port might be only STS-1 (51.84 Mbps). In this example, the CE-100T-8 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-1) is efficient because frame loss can be controlled to a large extent.

Ethernet Link Integrity Support

The CE-100T-8 supports end-to-end Ethernet link integrity (Figure 17-3). 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 means that if any part of the end-to-end path fails, the entire path fails. It disables the Ethernet port on the CE-100T-8 card if the remote Ethernet port is unable to transmit over the SONET network or if the remote Ethernet port is disabled.

Failure of the entire path is ensured by turning off the transmit pair at each end of the path. The attached Ethernet devices recognize the disabled transmit pair as a loss of carrier and consequently an inactive link or link fail.

Figure 17-3 End-to-End Ethernet Link Integrity Support

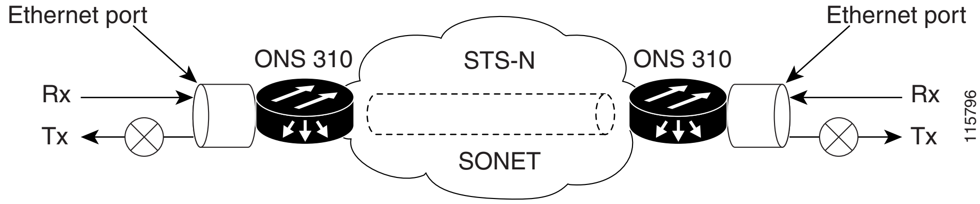

Note Some network devices can be configured to ignore a loss of carrier condition. If a device configured to ignore a loss of carrier condition attaches to a CE-100T-8 card at one end, alternative techniques (such as use of Layer 2 or Layer 3 keep-alive messages) are required to route traffic around failures. The response time of such alternate techniques is typically much longer than techniques that use link state as indications of an error condition.

Enhanced State Model for Ethernet and SONET Ports

The CE-100T-8 supports the Enhanced State Model (ESM) for the Ethernet ports, as well as for the SONET circuit. For more information about the ESM, refer to the "Administrative and Service States" appendix in the *Cisco ONS 15454 Reference Manual*.

The Ethernet ports can be set to the ESM service states including the In-Service, Automatic In-Service (IS,AINS) administrative state. IS,AINS initially puts the port in the Out-of-Service and Autonomous, Automatic In-Service (OOS-AU,AINS) state. In this service state, alarm reporting is suppressed, but traffic is carried. After the soak period passes, the port changes to In-Service and Normal (IS-NR). Raised fault conditions, whether their alarms are reported or not, can be retrieved on the CTC Conditions tab or by using the TL1 RTRV-COND command.

Two Ethernet port alarms/conditions, CARLOSS and TPTFAIL, can prevent the port from going into service. This occurs even though alarms are suppressed when a CE-100T-8 circuit is provisioned with the Ethernet ports set to the IS,AINS state, because the CE-100T-8 link integrity function is active and ensures that the links at both ends are not enabled until all SONET and Ethernet errors along the path are cleared. As long as the link integrity function keeps the end-to-end path down, both ports will have at least one of the two conditions needed to suppress the AINS-to-IS transition. Therefore, the ports will remain in the AINS state with alarms suppressed.

ESM also applies to the SONET circuits of the CE-100T-8 card. If the SONET circuit is set up in IS,AINS state and the Ethernet error occurs before the circuit transitions to IS, then link integrity will also prevent the circuit transition to the IS state until the Ethernet port errors are cleared at both ends. The service state will be OOS-AU,AINS as long as the administrative state is IS,AINS. When there are no Ethernet or SONET errors, link integrity enables the Ethernet port at each end. Simultaneously, the AINS countdown begins as normal. If no additional conditions occur during the time period, each port transitions to the IS-NR state. During the AINS countdown, the soak time remaining is available in CTC and TL1. The AINS soaking logic restarts from the beginning if a condition appears again during the soak period.

A SONET circuit provisioned in the IS,AINS state remains in the initial Out-of-Service (OOS) state until the Ethernet ports on each end of the circuit transition to the IS-NR state. The SONET circuit transports Ethernet traffic and counts statistics when link integrity turns on the Ethernet port, regardless of whether this AINS-to-IS transition is complete.

IEEE 802.1Q CoS and IP ToS Queuing

The CE-100T-8 references IEEE 802.1Q class of service (CoS) thresholds and IP type of service (ToS) (IP Differentiated Services Code Point [DSCP]) thresholds for priority queueing. CoS and ToS thresholds for the CE-100T-8 are provisioned on a per port level. This allows the user to provide priority treatment based on open standard quality of service (QoS) schemes already existing in the data network attached to the CE-100T-8. The QoS treatment is applied to both Ethernet and POS ports.

Any packet or frame with a priority greater than the set threshold is treated as priority traffic. This priority traffic is sent to the priority queue instead of the normal queue. When buffering occurs, packets on the priority queue preempt packets on the normal queue. This results in lower latency for the priority traffic, which is often latency-sensitive traffic, such as VoIP.

Because these priorities are placed on separate queues, the priority queuing feature should not be used to separate rate-based CIR/EIR marked traffic (sometimes done at a Metro Ethernet service provider edge). This could result in out-of-order packet delivery for packets of the same application, which would cause performance issues with some applications.

For an IP ToS-tagged packet, the CE-100T-8 can map any of the 256 priorities specified in IP ToS to priority or best effort. The user can configure a different ToS on CTC at the card-level view under the Provisioning > Ether Ports tabs. Any ToS class higher than the class specified in CTC is mapped to the priority queue, which is the queue geared towards low latency. By default, the ToS is set to 255, which is the highest ToS value. This results in all traffic being treated with equal priority by default.

Table 17-3 shows which values are mapped to the priority queue for sample IP ToS settings. (ToS settings span the full 0 to 255 range, but only selected settings are shown.)

Table 17-1 IP ToS Priority Queue Mappings

ToS Setting in CTC	ToS Values Sent to Priority Queue
255 (default)	None
250	251–255
150	151–255
100	101–255
50	51–255
0	1–255

For a CoS-tagged frame, the CE-100T-8 can map the eight priorities specified in CoS to priority or best effort. The user can configure a different CoS on CTC at the card-level view under the **Provisioning > Ether Ports** tabs. Any CoS class higher than the class specified in CTC is mapped to the priority queue, which is the queue geared towards low latency. By default, the CoS is set to 7, which is the highest CoS value. This results in all traffic being treated with equal priority by default.

Table 17-2 shows which values are mapped to the priority queue for CoS settings.

Table 17-2 CoS Priority Queue Mappings

CoS Setting in CTC	CoS Values Sent to Priority Queue
7 (default)	none
6	7
5	6, 7

Table 17-2 CoS Priority Queue Mappings (continued)

CoS Setting in CTC	CoS Values Sent to Priority Queue
4	5, 6, 7
3	4, 5, 6, 7
2	3, 4, 5, 6, 7
1	2, 3, 4, 5, 6, 7
0	1, 2, 3, 4, 5, 6, 7

Ethernet frames without VLAN tagging use ToS-based priority queueing if both ToS and CoS priority queueing is active on the card. The CE-100T-8 card's ToS setting must be lower than 255 (default) and the CoS setting lower than 7 (default) for CoS and ToS priority queueing to be active. A ToS setting of 255 (default) disables ToS priority queueing, so in this case the CoS setting would be used.

Ethernet frames with VLAN tagging use CoS-based priority queueing if both ToS and CoS are active on the card. The ToS setting is ignored. CoS based priority queueing is disabled if the CoS setting is the 7 (default), so in this case the ToS setting would be used.

If the CE-100T-8 card's ToS setting is 255 (default) and the CoS setting is 7 (default), priority queueing is not active on the card, and data gets sent to the default normal traffic queue. Also if data is not tagged with a ToS value or a CoS value before it enters the CE-100T-8 card, it gets sent to the default normal traffic queue.

Note Priority queuing has no effect when flow control is enabled (default) on the CE-100T-8. Under flow control a 6 kilobyte single-priority first in first out (FIFO) buffer fills, then a PAUSE frame is sent. This results in the packet ordering priority becoming the responsibility of the external device, which is buffering as a result of receiving the PAUSE flow-control frames.

Note Priority queuing has no effect when the CE-100T-8 is provisioned with STS-3C circuits. The STS-3c circuit has more data capacity than Fast Ethernet, so CE-100T-8 buffering is not needed. Priority queuing only takes effect when buffering occurs.

RMON and SNMP Support

The CE-100T-8 card features remote monitoring (RMON) that allows network operators to monitor the health of the network with a network management system (NMS). The CE-100T-8 uses the ONG RMON. The ONG RMON contains the statistics, history, alarms, and events MIB groups from the standard RMON MIB, as well as Simple Network Management Protocol (SNMP). A user can access RMON threshold provisioning through TL1 or CTC. For RMON threshold provisioning with CTC, refer to the *Cisco ONS 15454 Procedure Guide* and the *Cisco ONS 15454 Troubleshooting Guide*. For TL1 information, refer to the *Cisco ONS SONET TL1 Command Guide*.

Statistics and Counters

The CE-100T-8 has a full range of Ethernet and POS statistics under **Performance > Ether Ports** or **Performance > POS Ports**. These are detailed in the "Performance Monitoring" chapter of the *Cisco ONS 15454 Reference Manual.*

CE-100T-8 SONET Circuits and Features

The CE-100T-8 has eight POS ports, numbered one through eight, which are exposed to management with CTC or TL1. Each POS port is statically mapped to a matching Ethernet port. By clicking the card-level Provisioning tab > POS Ports tab, the user can configure the Administrative State, Framing Type, and Encapsulation Type. By clicking the card-level Performance tab > POS Ports tab, the user can view the statistics, utilization, and history for the POS ports.

Available Circuit Sizes and Combinations

Each POS port terminates an independent contiguous SONET concatenation (CCAT) or virtual SONET concatenation (VCAT). The SONET circuit is created for these ports through CTC or TL1 in the same manner as a SONET circuit for a non-Ethernet line card. Table 17-3 shows the circuit sizes available for the CE-100T-8 on the ONS 15310-CL and ONS 15310-MA.

Table 17-3 CE-100T-8 Supported Circuit Sizes

CCAT High Order	VCAT High Order	VCAT Low Order
STS-1	STS-1-1v	VT1.5-*n*V (*n*= 1 to 64)
STS-3c	STS-1-2v	
	STS-1-3v	

A single circuit provides a maximum of 100 Mbps of throughput, even when an STS-3c circuit, which has a bandwidth equivalent of 155 Mbps, is provisioned. This is due to the hardware restriction of the Fast Ethernet port. A VCAT circuit is also restricted in this manner. Table 17-3 shows the minimum SONET circuit sizes required for 10 Mbps and 100 Mbps wire speed service.

Table 17-4 SONET Circuit Size Required for Ethernet Wire Speeds

Ethernet Wire Speed	CCAT High Order	VCAT High Order	VCAT Low Order
Line Rate 100BaseT	STS-3c	STS-1-3v, STS-1-2v*	Not applicable
Sub Rate 100BaseT	STS-1	STS-1-1v	VT1.5-*xV* (x=1-64)
Line Rate 10BaseT	STS-1	Not applicable	VT1.5-7V
Sub Rate 10BaseT	Not applicable	Not applicable	VT1.5-xV (x=1-6)

*STS-1-2v provides a total transport capacity of 98 Mbps.

The number of available circuits and total combined bandwidth for the CE-100T-8 depends on the combination of circuit sizes configured. Table 17-5 shows the circuit size combinations available for CE-100T-8 CCAT high-order circuits on the ONS 15310-CL and ONS 15310-MA. Table 17-6 shows the circuit size combinations available for CE-100T-8 VCAT high-order circuits on the ONS 15310-CL and ONS 15310-MA.

Table 17-5 CCAT High Order Circuit Size Combinations

Number of STS-3c Circuits	Maximum Number of STS-1 Circuits
None	6

Table 17-5 CCAT High Order Circuit Size Combinations

Number of STS-3c Circuits	Maximum Number of STS-1 Circuits
1	3
2	None

Table 17-6 VCAT High Order Circuit Size Combinations

Number of STS-1-3v Circuits	Maximum Number of STS-1-2v Circuits
None	2
1	1
2	None

The CE-100T-8 supports up to eight low order VCAT circuits. The available circuit sizes are VT1.5-*n*v, where *n* ranges from 1 to 64. The total number of VT members cannot exceed 168 VT1.5s with each of the two pools on the card supporting 84 VT1.5s. The user can create a maximum of two circuits at the largest low order VCAT circuit size, VT1.5-64v.

A user can combine CCAT high order, VCAT high order, and VCAT low order circuits in any way as long as there is a maximum of eight circuits and the mapper chip bandwidth restrictions are observed. The following table details the maximum density service combinations.

Table 17-7 CE-100T-8 Maximum Service Densities

Service Combination	STS-3c or STS-1-3v	STS-1-2v	STS-1	VT1.5-xV (x=1-7)	Number of Active Service
1	2	0	0	0	2
2	1	1	1	0	3
3	1	0	3	0	4
4	1	0	0	7(x=1-12)*	8*
5	0	2	2	0	4
6	0	1	1	6(x=1-14)	8
7	0	1	0	7(x=1-12)*	8*
8	0	0	6	0	6
9	0	0	3	5(x=1-16)	8
10	0	0	0	8 (x=1-21)	8

* This LO-VCAT Circuit combination is achievable if the first circuit created on the card is an LO VCAT circuit. If the first circuit created on the card is HO-VCAT or CCAT STS circuits, then a maximum of six LO-VCAT circuits can be added on the card.

CE-100T-8 STS/VT Allocation Tab

The CE-100T-8 has two pools, each with a maximum capacity of three STSs. At the CTC card-level view under the Maintenance tab, the STS/VT Allocation tab displays how the provisioned circuits populate the two pools. This information can be useful in freeing up the bandwidth required for provisioning a circuit, if there is not enough existing capacity on any one pool for provisioning the desired circuit. The user can look at the distribution of the existing circuits among the two pools and decide which circuits to delete in order to free up space for the desired circuit.

Figure 17-4 CE-100T-8 STS/VT Allocation Tab

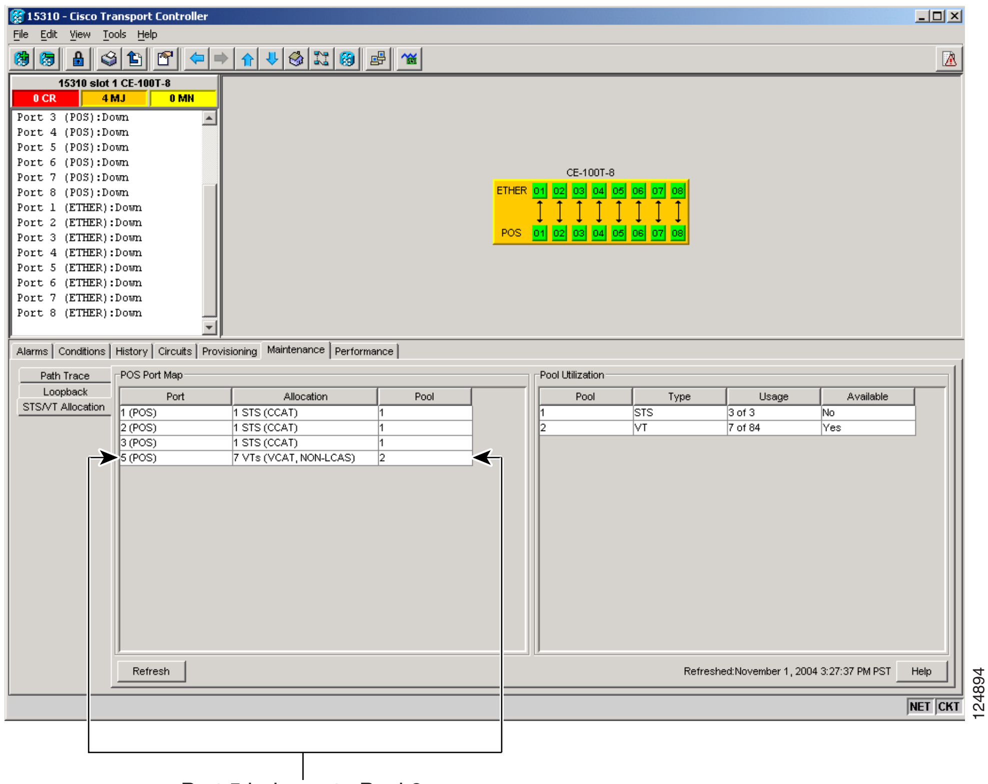

For example if a user needs to provision an STS-3c or STS-1-3v on the CE-100T-8 card shown in Figure 17-4, an STS-3c or STS-1-3v worth of bandwidth is not available from either of the two pools. The user needs to delete circuits from the same pool to free up bandwidth. If the bandwidth is available but scattered among the pools, the circuit cannot be provisioned.

Looking at the POS Port Map table, the user can determine which circuits belong to which pools. The Pool and Port columns in Figure 17-4 show that the circuit on port 5 is drawn from Pool 2, and no other circuits are drawn from Pool 2. Deleting this one circuit will free up an STS-3c or STS-1-3v worth of bandwidth from a single pool.

The POS Port table has a row for each port with three columns (Figure 17-4). They show the port number, the circuit size and type, and the pool it is drawn from. The Pool Utilization table has two columns and shows the pool number, the type of circuits on that pool, how much of the pool's capacity is being used, and whether additional capacity is available.

CE-100T-8 VCAT Characteristics

The ML-100T-8 card and the CE-100T-8 card (both the version for the ONS 15310-CL and ONS 15310-MA and the version for the ONS 15454 SONET/SDH) have hardware-based support for the ITU-T G.7042 standard link capacity adjustment scheme (LCAS). This allows the user to dynamically resize a high order or low order VCAT circuit through CTC or TL1 without affecting other members of the VCG (errorless). ML-100T-8 LCAS support is high order only and is limited to a two member VCG.

To enable end-to-end connectivity in a VCAT circuit that traverses through a third-party network, you must create a server trail between the ports. For more details, refer to the "Create Circuits and VT Tunnels" chapter in the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Procedure Guide*.

The ONS 15454 SONET/SDH ML-Series card has a software-based LCAS (SW-LCAS) scheme. This scheme is also supported by both the ML-100T-8 card and both versions of the CE-100T-8, but only for circuits with the other end terminating on an ONS 15454 SONET/SDH ML-Series card.

The SW-LCAS is not supported on CE-100T-8 cards for interoperation with the CE-MR-10, CE-MR-6, and ML-MR-10 cards.

The CE-100T-8 card allows independent routing and protection preferences for each member of a VCAT circuit. The user can also control the amount of VCAT circuit capacity that is fully protected, unprotected or if the circuit is on a bidirectional line switched ring (BLSR), uses protection channel access (PCA). Alarms are supported on a per-member as well as per virtual concatenation group (VCG) basis.

Note The maximum tolerable VCAT differential delay for the CE-100T-8 is 48 milliseconds. The VCAT differential delay is the relative arrival time measurement between members of a virtual concatenation group (VCG).

On ML-100T-8 and CE-100T-8 cards, members of a HW-LCAS circuit must be moved to the OOS,OOG (locked, outOfGroup) state before you delete them.

A traffic hit is seen under the following conditions:

- A hard reset of the card containing the trunk port.
- Trunk port moved to OOS,DSBLD(locked,disabled) state.
- Trunk fiber pull.
- Deletion of members of the HW-LCAS circuit in IG (In Group) state.

CE-100T-8 POS Encapsulation, Framing, and CRC

The CE-100T-8 uses Cisco EoS LEX (LEX). LEX is the primary encapsulation of ONS Ethernet cards. In this encapsulation the protocol field is set to the values specified in Internet Engineering Task Force (IETF) Request For Comments (RFC) 1841. The user can provision GPF-F framing (default) or high-level data link control (HDLC) framing. With GFP-F framing, the user can also configure a 32-bit CRC (the default) or no CRC (none). aming. With GFP-F framing, the user can also configure a 32-bit CRC (the default) or no CRC (none). On CTC go to CE card view and click the Provisioning >pos ports tab, to see the various parameters that can be configured on the POS ports, see Displaying ML-Series Ethernet Statistics in CTC, page 2-2. Various parameters like, admin state, service state, framing type, CRC , MTU and soak time for a port can be configured hereWhen LEX is used over GFP-F it is standard Mapped Ethernet over GFP-F according to ITU-T G.7041. HDLC framing provides a set 32-bit CRC.

Figure 17-5 illustrates CE-100T-8 framing and encapsulation.

Figure 17-5 ONS CE-100T-8 Encapsulation and Framing Options

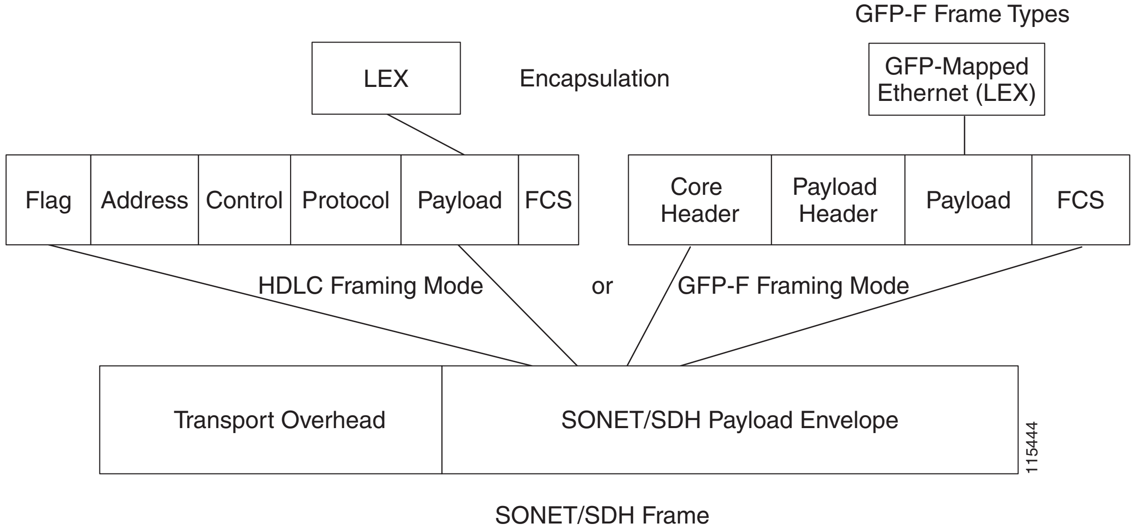

The CE-100T-8 card supports GFP-F null mode. GFP-F CMFs are counted and discarded.

The CE-100T-8 card is interoperable with the ML-100T-8 card and several other ONS Ethernet cards. For specific details on ONS Ethernet card interoperability, refer to the "POS on ONS Ethernet Cards" chapter of the *Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide.*

CE-100T-8 Loopback, J1 Path Trace, and SONET Alarms

The CE-100T-8 card supports terminal and facility loopbacks when in the Out of Service, Maintenance state (OOS, MT). It also reports SONET alarms and transmits and monitors the J1 Path Trace byte in the same manner as OC-N cards. Support for path termination functions includes:

- H1 and H2 concatenation indication
- C2 signal label
- Bit interleaved parity 3 (BIP-3) generation
- G1 path status indication
- C2 path signal label read/write
- Path level alarms and conditions, including loss of pointer, unequipped, payload mismatch, alarm indication signal (AIS) detection, and remote defect indication (RDI)
- J1 path trace for high order paths
- J2 path trace for low order paths
- J2 path trace for low order VCAT circuits at the member level
- Extended signal label for the low order paths

CE-MR-6 Ethernet Card

This section describes the operation of the CE-MR-6 card supported on the ONS 15310-MA. The CE-MR-6 card installed in ONS 15310-MA is restricted to SONET operations.

Provisioning is done through Cisco Transport Controller (CTC) or Transaction Language One (TL1). Configurations through Cisco IOS terminal are not supported on the CE-MR-6 card.

For Ethernet card specifications, refer to the *Cisco ONS 15310 Reference Manual*. For step-by-step Ethernet card circuit configuration procedures, refer to the *Cisco ONS 15310 Procedure Guide*. Refer to the *Cisco ONS SONET TL1 Command Guide* for TL1 provisioning commands.

This section include the following topics:

- CE-MR-6 Overview, page 17-12
- CE-MR-6 Ethernet Features, page 17-13
- CE-MR-6 Circuits and Features, page 17-19

CE-MR-6 Overview

CE-MR-6 card is a 5 Gbps data module for use in the Cisco ONS 15310-MA. It provides support for L1 packet mapping functions (Ethernet to SONET). The 10/100/1000 Mbps Ethernet-encapsulated traffic is mapped to SONET circuits. Each circuit has three main attributes:

- Low order or high order
- Contiguous concatenation (CCAT) or virtual concatenation (VCAT)
- Generic framing procedure (GFP), LEX, high-level data link control (HDLC), or PPP (point-to-point protocol) based framing.

The CE-MR-6 cards support LCAS that allows hitless dynamic adjustment of SONET link bandwidth.

The CE-MR-6 is a Layer 1 (Ethernet Private Line) and Layer 1+ (Virtual Private Wire Services) mapper card with six IEEE 802 compliant 10/100/1000 Mbps Ethernet ports that provide 1:1 mapping of Ethernet ports to circuits. It maps each port to a unique SONET circuit in a point-to-point configuration. Figure 16-6 illustrates a sample CE-MR-6 application. In this example, data traffic from the Fast Ethernet port of a switch travels across the point-to-point circuit to the Fast Ethernet port of another switch.

Figure 16-6 CE-MR-6 Point-to-Point Circuit

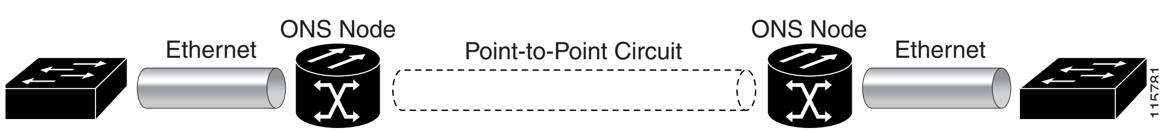

The CE-MR-6 card allows you to provision and manage an Ethernet private line service like a traditional SONET line. CE-MR-6 card applications include providing carrier-grade Ethernet private line services and high-availability transport.

The CE-MR-6 card carries any Layer 3 protocol that can be encapsulated and transported over Ethernet, such as IP or IPX. The Ethernet frame from the data network is transmitted on the Ethernet cable into the 10/100/1000 Mbps Ethernet ports on a CE-MR-6 card. The CE-MR-6 card transparently maps Ethernet frames into the SONET payload using packet-over-SONET/SDH (POS) encapsulation. The POS circuit, with its encapsulated Ethernet inside, is then multiplexed onto an optical card like any other

SONET synchronous transport signal (STS). When the payload reaches the destination node, the process is reversed and the data is transmitted from the 10/100/1000 Mbps Ethernet ports in the destination CE-MR-6 card onto the Ethernet cable and data network. The POS process is covered in detail in Chapter 6, "Configuring POS on the ML-Series Card."

The CE-MR-6 card supports ITU-T G.707-based standards. It allows a soft reset, which is errorless in most cases. During the soft reset, if there is a provisioning change, or if the firmware is replaced during a software upgrade process, the reset is equivalent to a hard reset. For more information on a soft reset of a CE-MR-6 card using Cisco Transport Controller (CTC), refer to the *Cisco ONS 15310 Procedure Guide*.

CE-MR-6 Ethernet Features

The Ethernet interface of the CE-MR-6 card comprises six front-end Small Form-factor Pluggable (SFP) slots. For each slot, the interface speed and media type is determined by the installed SFP module. The SFP slots support 10 Mbps, 100 Mbps, and 1000 Mbps (Gigabit Ethernet) operation. The SFP modules supporting the intended rate can be copper (10/100/1000 Mbps) or optical (100/1000 Mbps). SFP modules are offered as separate orderable products for flexibility. For SFP details, refer to the *Cisco ONS 15310 Reference Manual or* Installing the GBIC, SFP, SFP+, XFP, CXP, and CFP Optical Modules in Cisco ONS Platforms. Ethernet Ports 1 through 6 each map to a POS port with a corresponding number. The console port on the CE-MR-6 card is not functional.

The CE-MR-6 card forwards valid Ethernet frames without modifying it over the SONET network. Information in the headers is not affected by encapsulation and transport. IEEE 802.1Q information travels through the process unaffected.

The CE-MR-6 supports jumbo frames with MTU sizes of 64 to 9600 bytes.

The CE-MR-6 card discards certain types of erroneous Ethernet frames rather than transport them over SONET. Erroneous Ethernet frames include corrupted frames with CRC errors and undersized frames that do not conform to the minimum 64-byte length, or oversized frames greater than 9600 bytes Ethernet standard.

Note Many Ethernet attributes are also available through the network element (NE) defaults feature. For more information on NE defaults, refer to the "Network Element Defaults" appendix in the *Cisco ONS 15310 Reference Manual*.

Autonegotiation, Flow Control, and Frame Buffering

On the CE-MR-6 card, Ethernet link autonegotiation is on by default when the duplex or speed of the port is set to auto. The user can also set the link speed, duplex, selective autonegotiation, and flow control manually under the card-level Provisioning tab in CTC.

The CE-MR-6 card supports selective autonegotiation on the Ethernet ports. If selective autonegotiation is enabled, the port attempts to autonegotiate only to a specific speed and duplex. The link will come up if both the speed and duplex of the attached autonegotiating device matches that of the port. You cannot enable selective autonegotiation if either the speed or duplex of the port is set to auto.

The CE-MR-6 card supports IEEE 802.3x flow control and frame buffering to reduce data traffic congestion. Flow control is on by default.

To prevent over-subscription, buffer memory is available for each port. When the buffer memory on the Ethernet port nears capacity, the CE-MR-6 card uses IEEE 802.3x flow control to transmit a pause frame to the attached Ethernet device. Flow control and autonegotiation frames are local to the Ethernet (10 Mbps), Fast Ethernet (100 Mbps), Gigabit Ethernet (1000 Mbps) interfaces and attached Ethernet devices. These frames do not continue through the POS ports.

The CE-MR-6 card has asymmetric flow control and proposes asymmetric flow control when autonegotiating flow control with attached Ethernet devices.

The pause frame instructs the source to stop sending packets for a specific period of time. The sending station waits the requested amount of time before sending more data. Figure 16-7 illustrates pause frames being sent and received by CE-MR-6 cards and attached switches.

Figure 16-7 Flow Control

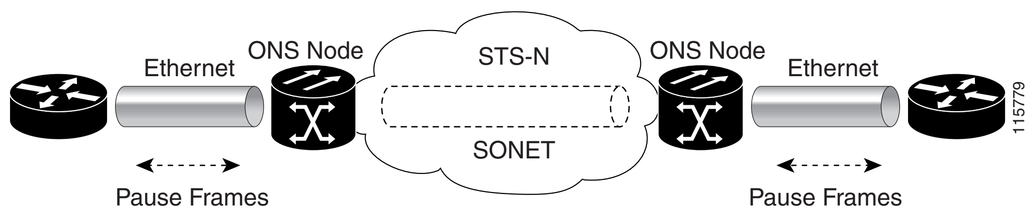

This flow-control mechanism matches the sending and receiving device throughput the bandwidth of the STS circuit. For example, a router might transmit to the Ethernet port on the CE-MR-6 card. This particular data rate might occasionally exceed 51.84 Mbps, but the SONET circuit assigned to the CE-MR-6 port might be only STS-1 (51.84 Mbps). Under this condition, the CE-MR-6 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-1) is efficient because frame loss is controlled to a large extent.

Ethernet Link Integrity Support

The CE-MR-6 supports end-to-end Ethernet link integrity (Figure 16-8). 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. Link integrity is implemented so that the Ethernet over SONET connection behaves more like an Ethernet cable from the viewpoint of the attached Ethernet devices.

End-to-end Ethernet link integrity means that if any part of the end-to-end path fails, the entire path fails. It disables the Ethernet port transmitter on the CE-MR-6 card when the remote Ethernet port does not have a receive signal or when the SONET near end or far-end failure is detected. The failure of the entire path is ensured by turning off the transmit pair at each end of the path. If any part of the end-to-end path fails, the CE-MR-6 card soaks the defect for a fixed duaration of 200 ms. The attached Ethernet devices recognize the disabled transmit pair as a loss of carrier and consequently an inactive link or link failure. The transport fail alarm is raised when the port transmitter is disabled. Link integrity supports a double fault, that is, when both the Ethernet ports do not receive a signal.

Note Only bidirectional link integrity is supported on the CE-MR-6 card.

Figure 16-8 End-to-End Ethernet Link Integrity Support

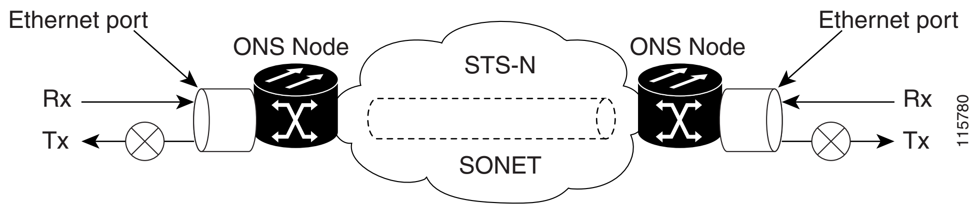

Note Some network devices can be configured to ignore a loss of carrier condition. If a device configured to ignore a loss of carrier condition attaches to a CE-MR-6 card at one end, alternative techniques (such as use of Layer 2 or Layer 3 keep-alive messages) are required to route traffic around failures. The response time of such alternate techniques is typically much longer than techniques that use link state as indications of an error condition.

In certain network configurations, the restoration time, for example, after a protection switch can be more than 200 ms. Such disruptions necessitate the link integrity to be initiated at an interval greater than 200 ms. To allow link integrity to be initiated at an interval greater than 200 ms, set the link integrity timer in the range between 200 and 5000 ms, in multiples of 100 ms.

Note The accuracy of the Link Integrity timer is less on CE-MR-6 card compared to the G1000 or CE-1000 cards. The accuracy of Link Integrity timer is within 200 ms for the CE-MR-6 card.

Administrative and Service States with Soak Time for Ethernet and SONET Ports

The CE-MR-6 card can be managed by TL1, SNMP, CTC or CTM. The card supports the administrative and service states for the Ethernet ports and the SONET circuit. For more information about card and circuit service states, refer to the "Administrative and Service States" appendix in the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Reference Manual*.

The Ethernet ports can be set to the Enhanced State Model (ESM) service states including the In-Service, Automatic In-Service (IS, AINS) administrative state. IS, AINS initially puts the port in the Out-of-Service and Autonomous, Automatic In-Service (OOS-AU, AINS) state. In this service state, alarm reporting is suppressed, but traffic is carried and loopbacks are allowed. After the soak period passes, the port changes to In-Service and Normal (IS-NR). Raised fault conditions, whether their alarms are reported or not, can be retrieved from the CTC Conditions tab or by using the TL1 RTRV-COND command.

Two Ethernet port alarms/conditions, CARLOSS and TPTFAIL, can prevent the port from going into service. This occurs even though alarms are suppressed when a CE-MR-6 circuit is provisioned with the Ethernet ports set to the IS,AINS state, because the CE-MR-6 link integrity function is active and ensures that the links at both ends are not enabled until all SONET and Ethernet errors along the path are cleared. As long as the link integrity function keeps the end-to-end path down, both ports will have at least one of the two conditions needed to suppress the AINS-to-IS transition. Therefore, the ports will remain in the AINS state with alarms suppressed.

ESM also applies to the SONET circuits of the CE-MR-6 card. If the SONET circuit is set up in IS,AINS state and the Ethernet error occurs before the circuit transitions to IS, then link integrity will also prevent the circuit transition to the IS state until the Ethernet port errors are cleared at both ends. The service state will be OOS-AU,AINS as long as the administrative state is IS,AINS. When there are no Ethernet or SONET errors, link integrity enables the Ethernet port at each end. Simultaneously, the AINS countdown begins as normal. If no additional conditions occur during the time period, each port transitions to the IS-NR state. During the AINS countdown, the soak time remaining is available in CTC and TL1. The AINS soaking logic restarts from the beginning if a condition appears again during the soak period.

A SONET circuit provisioned in the IS,AINS state remains in the initial Out-of-Service (OOS) state until the Ethernet ports on each end of the circuit transition to the IS-NR state. The SONET circuit transports Ethernet traffic and counts statistics when link integrity turns on the Ethernet port, regardless of whether this AINS-to-IS transition is complete.

IEEE 802.1Q CoS and IP ToS Queuing

The CE-MR-6 references IEEE 802.1Q class of service (CoS) thresholds and IP type of service (ToS) (IP Differentiated Services Code Point [DSCP]) thresholds for priority queueing. CoS and ToS thresholds for the CE-MR-6 are provisioned on a per port level. This allows the user to provide priority treatment based on open standard quality of service (QOS) schemes already existing in the data network attached to the CE-MR-6. The QOS treatment is applied to both Ethernet and POS ports.

Any packet or frame with a priority greater than the set threshold is treated as priority traffic. This priority traffic is sent to the priority queue instead of the normal queue. When buffering occurs, packets on the priority queue preempt packets on the normal queue. This results in lower latency for the priority traffic, which is often latency-sensitive traffic such as voice-over-IP (VoIP).

Because these priorities are placed on separate queues, the priority queuing feature should not be used to separate rate-based CIR/EIR marked traffic (sometimes done at a Metro Ethernet service provider end). This could result in out-of-order packet delivery for packets of the same application, which would cause performance issues with some applications.

For an IP ToS-tagged packet, the CE-MR-6 can map any of the 256 priorities specified in IP ToS to priority or best effort. The user can configure a different ToS in CTC at the card-level view under the Provisioning > Ether Ports tabs. Any ToS class higher than the class specified in CTC is mapped to the priority queue, which is the queue geared towards low latency. By default, the ToS is set to 255, which is the highest ToS value. This results in all traffic being treated with equal priority by default.

Table 16-8 shows which values are mapped to the priority queue for sample IP ToS settings. (ToS settings span the full 0 to 255 range, but only selected settings are shown in Table 16-8.)

Table 16-8 IP ToS Priority Queue Mappings

ToS Setting in CTC	ToS Values Sent to Priority Queue
255 (default)	None
250	251–255
150	151–255
100	101–255
50	51–255
0	1–255

For a CoS-tagged frame, the CE-MR-6 can map the eight priorities specified in CoS to priority or best effort. The user can configure a different CoS in CTC at the card-level view under the Provisioning > Ether Ports tabs. Any CoS class higher than the class specified in CTC is mapped to the priority queue, which is the queue geared towards low latency. By default, the CoS is set to 7, which is the highest CoS value. This results in all traffic being treated with equal priority by default.

Table 16-9 shows values that are mapped to the priority queue for CoS settings.

Table 16-9 CoS Priority Queue Mappings

CoS Setting in CTC	CoS Values Sent to Priority Queue
7 (default)	None
6	7
5	6, 7
4	5, 6, 7
3	4, 5, 6, 7
2	3, 4, 5, 6, 7
1	2, 3, 4, 5, 6, 7
0	1, 2, 3, 4, 5, 6, 7

Ethernet frames without VLAN tagging use ToS-based priority queueing if both ToS and CoS priority queueing is active on the card. The CE-MR-6 card's ToS setting must be lower than 255 (default) and the CoS setting lower than 7 (default) for CoS and ToS priority queueing to be active. A ToS setting of 255 (default) disables ToS priority queueing, so in this case the CoS setting would be used.

Ethernet frames with VLAN tagging use CoS-based priority queueing if both ToS and CoS are active on the card. The ToS setting is ignored. CoS based priority queueing is disabled if the CoS setting is 7 (default), so in this case the ToS setting would be used.

If the CE-MR-6 card's ToS setting is 255 (default) and the CoS setting is 7 (default), priority queueing is not active on the card, and data gets sent to the default normal traffic queue. If data is not tagged with a ToS value or a CoS value before it enters the CE-MR-6 card, it also gets sent to the default normal traffic queue.

Note Priority queuing has no effect when flow control is enabled (default) on the CE-MR-6. When flow control is enabled, a 6KB, single-priority, first-in first-out (FIFO) buffer fills, then a PAUSE frame is sent. This results in the packet ordering priority becoming the responsibility of the external device, which is buffering as a result of receiving the PAUSE flow-control frames.

Note Priority queuing takes effect only when there is congestion at the egress POS. For example, priority queuing has no effect when the CE-MR-6 card is provisioned with STS-3C circuits and the front-end is 100 Mbps. The STS-3c circuit has more data capacity than Fast Ethernet, so CE-MR-6 buffering is not needed. Priority queuing only takes effect during buffering.

RMON and SNMP Support

The CE-MR-6 card features remote monitoring (RMON) that allows network operators to monitor the health of the network with a network management system (NMS). The CE-MR-6 card uses ONG RMON. ONG RMON contains statistics, history, alarms, and events MIB groups from the standard RMON MIB as well as Simple Network Management Protocol (SNMP). A user can access RMON threshold provisioning through TL1 or CTC. For RMON threshold provisioning with CTC, see the *Cisco ONS 15310 Procedure Guide* and the *Cisco ONS 15310 Troubleshooting Guide*.

SNMP MIBs Supported

The following SNMP MIBs are supported by the CE-MR-6 card:

- RFC2819 -MIB
  - etherStatsOversizePkts
  - etherStatsUndersizePkts
  - etherStatsJabbers
  - etherStatsCollisions
  - etherStatsDropEvents
  - etherStatsOctets.
  - etherStatsBroadcastPkts.
  - etherStatsMulticastPkts.
  - etherStatsCRCAlignErrors.
  - etherStatsFragments
  - etherStatsPkts64Octets.
  - etherStatsPkts65to127Octets
  - etherStatsPkts128to255Octets
  - etherStatsPkts256to511Octets
  - etherStatsPkts512to1023Octets.
  - etherStatsPkts1024to1518Octets.
  - Rx Utilization
  - Tx Utilization
- RFC2233-MIB
  - ifInUcastPkts
  - ifOutUcastPkts
  - ifInMulticastPkts
  - ifInBroadcastPkts
  - ifInDiscards
  - ifInOctets
  - ifOutOctets
  - ifInErrors
- RFC2358-MIB
  - dot3StatsFCSErrors
  - dot3StatsSingleCollisionFrames
  - dot3StatsFrameTooLong

Statistics and Counters

The CE-MR-6 has a full range of Ethernet and POS statistics information under Performance > Ether Ports or Performance > POS Ports.

Supported Cross-connects

There is no restriction on the number of CE-MR-6 cards that could be added in one chassis or the slot where the CE-MR-6 cards can be placed. CE-MR-6 card is supported with XC cards.

CE-MR-6 Circuits and Features

The CE-MR-6 card has 6 POS ports, numbered 1 through 6, which can be managed via CTC or TL1. Each POS port is statically mapped to a matching Ethernet port. By clicking the card-level Provisioning > POS Ports tab, the user can configure the Administrative State, and Encapsulation Type. By clicking the card-level Performance > POS Ports tab, the user can view the statistics, utilization, and history of the POS ports.

Available Circuit Sizes and Combinations

Each POS port terminates an independent CCAT or VCAT circuit. The circuit is created for these ports through CTC or TL1 in the same manner as a circuit for a non-Ethernet line card.

Note If a CE-MR-6 card with a 1 Gbps SFP is installed and cross connected to another card that supports only 10 or 100 Mbps, (for example, CE-100T-8 cards in the ONS 15310-MA) packet loss may occur if the SONET circuit between the two cards is more than 100 Mbps. In the CE-100T-8 example, a STS-1-2v circuit is errorless; however, if a STS-1-3v is used there will be a packet loss in the CE-MR-6 to CE-100T-8 direction.

Table 16-10 show the circuit sizes available for CE-MR-6 on ONS 15310.

Table 16-10 Supported SONET Circuit Sizes of CE-MR-6 on ONS 15310

CCAT	VCAT High Order	VCAT Low Order
STS-1	STS-1-*n*v (*n*=1 to 21)	VT1.5-*n*v (*n*=1 to 64)
STS-3c	STS-3C-*n*v (*n*=1 to 7)	
STS-6c		
STS-9c		
STS-12c		
STS-24c		
STS-48c		

Table 16-11 shows the minimum SONET circuit sizes required for wire speed service delivery.

Table 16-11 Minimum SONET Circuit Sizes for Ethernet Speeds

Ethernet Wire Speed	CCAT High Order	VCAT High Order STS-1-nv(n=1 to 21)	VCAT Low Order
Line Rate 1000 Mbps	STS-48c or STS-24c	STS-1-21v or STS-3-7v	Not applicable
Sub Rate 1000 Mbps	STS-12c, STS-9c, STS-6c, STS-3c, and STS-1	STS-1-1v to STS-1-20v	VT1.5-*xv* (x=1-64)
Line Rate 100 Mbps	STS-3c	STS-1-3v or STS-1-2v[1]	VT1.5-*xv* (x=56-64)
Sub Rate 100 Mbps	STS-1	STS-1-1v	VT1.5-*xv* (x=1-55)
Line Rate 10 Mbps	STS-1	Not applicable	VT1.5-7v
Sub Rate 10 Mbps	Not applicable	Not applicable	VT1.5-xv (x=1-6)

1. STS-1-2v provides a total transport capacity of 98 Mbps.

Table 16-12 shows VCAT high-order circuit size combinations available for CE-MR-6 on ONS 15310 for slots 1, 2, 5, and 6.

Table 16-12 VCAT High-Order Circuit Combinations for STS on ONS 15310 (Slots 1 , 2, 5, and 6)

STS Circuit Combinations	VT Circuits
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c, STS-48c, or STS-nv circuits up to a maximum of 10 circuits or maximum of: • CCAT—48 STSs • STS-1 VCAT—47 STSs • STS-3c VCAT—45 STSs	No VTs
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c circuits up to a maximum of 9 circuits or maximum of: • CCAT—46 STSs • STS-1 VCAT—45 STSs • STS-3c VCAT—39 STSs	1 VT1.5-48v circuit
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c circuits up to a maximum of 9 circuits or maximum of: • CCAT—44 STSs • STS-1 VCAT—43 STSs • STS-3c VCAT—33 STSs	1 VT1.5-64v circuit
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c circuits up to a maximum of 8 circuits or maximum of: • CCAT—44 STSs • STS-1 VCAT—43 STSs • STS-3c VCAT—33 STSs	2 VT1.5-48v circuits
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c circuits up to a maximum of 8 circuits or maximum of: • CCAT—42 STSs • STS-1 VCAT—41 STSs • STS-3c VCAT—27 STSs	2 VT1.5-64v circuits
Any combination of STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c circuits up to a maximum of 7 circuits or maximum of: • CCAT—42 STSs • STS-1 VCAT—41 STSs • STS-3c VCAT—27 STSs	3 VT1.5-48v circuits

Note You can replace STSs for VTs at the following rates:
Add 48 VT 1.5s at the cost of two STS-1s. If all the high order (HO) circuits are VCAT, one additional STS-1 is lost. Alternatively, you can add 48 VT 1.5s at the cost of two STS-3cs. If all the HO circuits are VCAT, one additional STS-3c is lost.
In some cases, circuits can be added by reducing the circuits for other concatenation rates.

CE-MR-6 Pool Allocation

> **Note** CE-MR Pool allocation can be set for ONS 15310-MA only and can be provisioned only when the card is in automatic provisioning mode.

CE-MR-6 has the following characteristics:

One pool can support only one circuit (from any front port), and must be a STS-48c circuit only. Two pools of 24 STSs each can be independently allocated.

- Ports 1-4 assigned to pool 1; ports 5-6 to pool 2.
- Each pool can support only one circuit type at a time. The circuit types are:
  - 1 STS-24c
  - Up to 2 STS-12cs
  - Up to 2 STS-9cs
  - Up to 4 STS-6cs
  - Up to 4 CCATs/ VCATs made up of STS-1s
  - Up to 4 CCATs/VCATs made up of STS-3cs
  - Number of members available for VCAT is subject to the total limit of 24 STSs in pool 1 and 23 STSs (or 21 STSs in case of STS-3C/VC4) in pool 2 (per VCG limit 21 STSs)
  - In addition, Pool1 can also support low order circuits (but not pool2)
  - VT1.5 based VCATs are subject to a total limit of 144 VT1.5s (per VCG limit is 64)
- The 1 pool versus 2 pool operation is decided dynamically by the first circuit provisioned.
- If two pool exists, then the mode of each pool is decided dynamically by the circuit types provisioned.

CE-MR-6 VCAT Characteristics

The CE-MR-6 card has hardware-based support for the ITU-T G.7042 standard LCAS. This allows the user to dynamically resize a high order or low order VCAT circuit through CTC or TL1 without affecting other members of the virtual concatenation group (VCG) (errorless).

The ONS 15310 ML100X-8 card has a software-based LCAS (SW-LCAS) scheme. Software LCAS is supported on CE-MR-6 cards for interoperation with these cards.

The SW-LCAS is not supported on CE-MR-6 cards for interoperation with the CE-100T-8 and ML-MR-10 cards.

> **Note** The CE-MR-10 card does not support interoperation between the LCAS and non-LCAS circuits.

The CE-MR-6 card allows independent routing and protection preferences for each member of a VCAT circuit. Alarms are supported on a per-member as well as per-VCG basis.

> **Note** A differential delay of 135 ms is supported for high order circuits and low order circuits.

On the CE-MR-6 card, members of a HW-LCAS circuit must be moved to the OOS,OOG (locked, outOfGroup) state before:

- Creating or deleting HW-LCAS circuits.
- Adding or deleting HW-LCAS circuit members.
- Changing the state to OOS,DSBLD.
- Changing the state from OOS,DSBLD to any other state.

A traffic hit is seen under the following conditions:

- A hard reset of the card containing the trunk port.
- Trunk port moved to OOS,DSBLD(locked,disabled) state.
- Trunk fiber pull.
- Deletion of members of the HW-LCAS circuit in IG (In Group) state.

Note CE-MR-6 cards display symmetric bandwidth behavior when an AIS, UNEQ, LOP, SF, SD, PLM, ENCAP, OOF, or PDI alarm is raised at the near-end member of the HW-LCAS circuit. The LCAS-SINK-DNU alarm and the RDI condition are raised at the far-end member of the circuit. The LCAS-SINK-DNU alarm changes the member state to outOf Group (OOG) and hence, the traffic goes down in both directions. For more information about alarms, refer to the "Alarm Troubleshooting" chapter in the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Troubleshooting Guide* or the *Cisco ONS 15310-MA SDH Troubleshooting Guide*.

Caution Packet losses might occur when an optical fiber is reinserted or when a defect is cleared on members of the HW-LCAS split fiber routed circuits.

CE-MR-6 POS Encapsulation, Framing, and CRC

The CE-MR-6 card uses Cisco EoS LEX (LEX). LEX is the primary encapsulation of ONS Ethernet cards. In this encapsulation, the protocol field is set to the values specified in Internet Engineering Task Force (IETF) Request For Comments (RFC) 1841. The user can provision frame-mapped generic framing procedure (GFP-F) framing (default) or HDLC framing. With GFP-F framing, the user can also configure a 32-bit CRC (the default) or no CRC (none). When LEX is used over GFP-F it is standard Mapped Ethernet over GFP-F according to ITU-T G.7041. HDLC framing provides a set 32-bit CRC. For more details about the interoperability of ONS Ethernet cards, including information on encapsulation, framing, and CRC, see the Chapter 6, "Configuring POS on the ML-Series Card."

The CE-MR-6 card supports GFP-F null mode. GFP-F CMFs are counted and discarded.

CE-MR-6 Loopback, J1 Path Trace, and SONET Alarms

The CE-MR-6 card supports terminal and facility loopbacks. It also reports SONET alarms and transmits and monitors the J1 Path Trace byte in the same manner as OC-N cards. Support for path termination functions includes:

- H1 and H2 concatenation indication
- C2 signal label
- Bit interleaved parity 3 (BIP-3) generation

- G1 path status indication
- C2 path signal label read/write
- Path level alarms and conditions, including loss of pointer (LOP), unequipped, payload mismatch, alarm indication signal (AIS) detection, and remote defect indication (RDI)
- J1 path trace for high-order CCAT paths
- J2 path trace for low-order VCAT circuits at the member level
- Extended signal label for the low-order paths

Terminal and Facility Loopback on LCAS Circuits In Split Fibre Routing

The following section lists guidelines to follow when the CE-MR-6 card includes a split fiber routing in a terminal and facility loopback on SW-LCAS circuits:

Note Make sure that you follow the guidelines and tasks listed in the following section. Not doing so will result in traffic going down on members passing through optical spans that do not have loopbacks.

- SW-LCAS circuit members must have J1 path trace set to manual.
- Transmit and receive traces must be unique.
- SW-LCAS circuits on CE-MR-6 must allow our of group (OOG) members on Trace Identifier Mismatch - Path (TIM-P).
- For members on split fiber routes, facility loopback must select the AIS option in CTC.
- Traffic hit is expected when loopback is applied. This is due to aysnchronous detection of VCAT defects and TIM-P detection on the other end of the circuit. This is acceptable since loopbacks are intrusive and affect traffic.

However, place members of an HW-LCAS circuit traversing an optical interface under maintenance in OOS,OOG (locked, outOfGroup) state before applying terminal/facility loopbacks.