



CHAPTER 10

Ethernet Operation

This chapter describes the operation of the Cisco ONS 15600 ASAP Ethernet card.

For Ethernet card specifications, refer to [Appendix A, “Hardware Specifications.”](#) For step-by-step Ethernet card circuit configuration procedures, refer to the *Cisco ONS 15600 Procedure Guide*. Refer to the *Cisco ONS SONET TLI Command Guide* for TLI provisioning commands.

Chapter topics include:

- [10.1 Any Service Any Port Card Application, page 10-1](#)
- [10.2 Transport Functionality, page 10-2](#)
- [10.3 Ethernet Rates and Mapping, page 10-4](#)
- [10.4 Protocols over Ethernet, page 10-5](#)
- [10.5 Buffering and Flow Control, page 10-6](#)
- [10.6 Autonegotiation, page 10-7](#)
- [10.7 Gigabit EtherChannel/IEEE 802.3ad Link Aggregation, page 10-8](#)

10.1 Any Service Any Port Card Application

The Any Service Any Port (ASAP) Carrier Card plugs into any of the eight available I/O slots for the ONS 15600. Each ASAP carrier card has four slots available for up to and including four ASAP Pluggable Input/output Modules (PIMs). There are four slots available on the PIM for Pluggable Port Modules (PPM), which are Small Form-factor Pluggable (SFP) optics. Four PPMs per ASAP PIM are allowed, which means that there can be as many as 16 PPM optical network interfaces for each ASAP carrier card. Each PPM can support any of the following single-rate SFP optics:

- OC-3 LR-2
- OC-12 LR-2
- OC-48 LR-2

Each of the PPM ports also supports the following multirate SFP optics:

- OC-3-SR-1
- OC-12-SR-1
- OC-48 IR-1
- GE LX

Each ASAP 4-port I/O (ASAP_4PIO) PIM is hot-pluggable while other ports on other PIMs are functioning. Each SFP is also hot-pluggable.

Layer 1 Ethernet transport is implemented for Gigabit Ethernet (GE) interfaces. GE traffic is encapsulated by generic framing procedure (GFP), ITU X.86, or Cisco high-level data link control/LAN extension (HDLC/LEX) and mapped into a SONET payload.

The data path processing consists of:

- Optical-to-electrical conversion (O/E) and electrical to optical (E/O) on the ASAP_4PIO PIM
- Gigabit Ethernet to Ethernet over SONET (EoS) mapping (for the GE ports only)



Note

For a single flow of Gigabit Ethernet, traffic can be mapped to an STS-48c. This choice limits maximum usable Ethernet ports to eight when all are mapped to STS-48c. Any time an STS-48c is used for an Ethernet connection, one-eighth of the total card bandwidth is consumed on the EoS mapper.

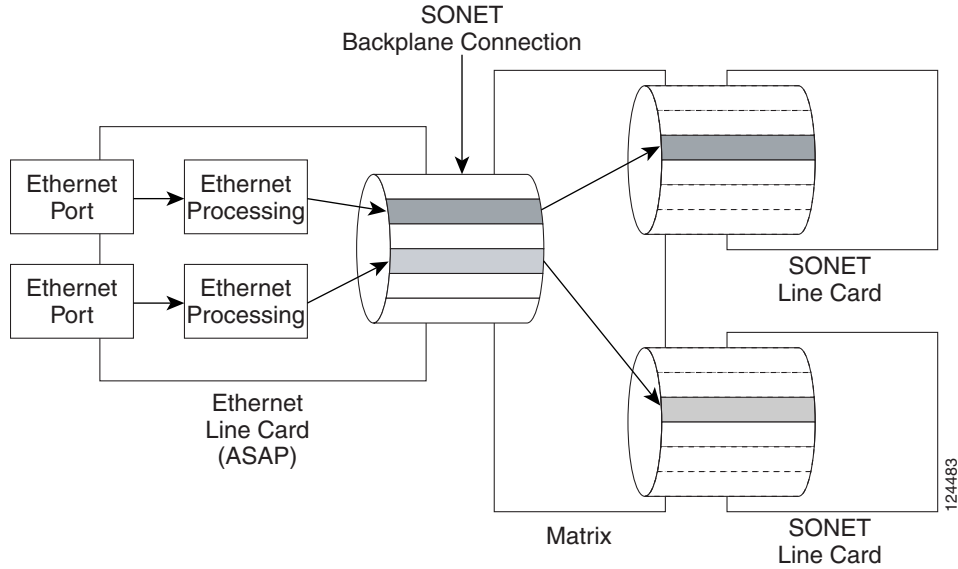
Ethernet facilities on the ASAP card include:

- Ethernet ports can be provisioned on a multirate PPM.
- Encapsulation methods allowed for the Ethernet ports are:
 - Cisco HDLC/LEX
 - GFP
 - ITU X.86
- Provisioning of an Ethernet port automatically provisions a POS port for the connection and circuit.

10.2 Transport Functionality

Figure 10-1 shows the transport of Ethernet frames into a SONET path using the ASAP carrier card in the ONS 15600.

Figure 10-1 ONS 15600 Ethernet Frame Transport

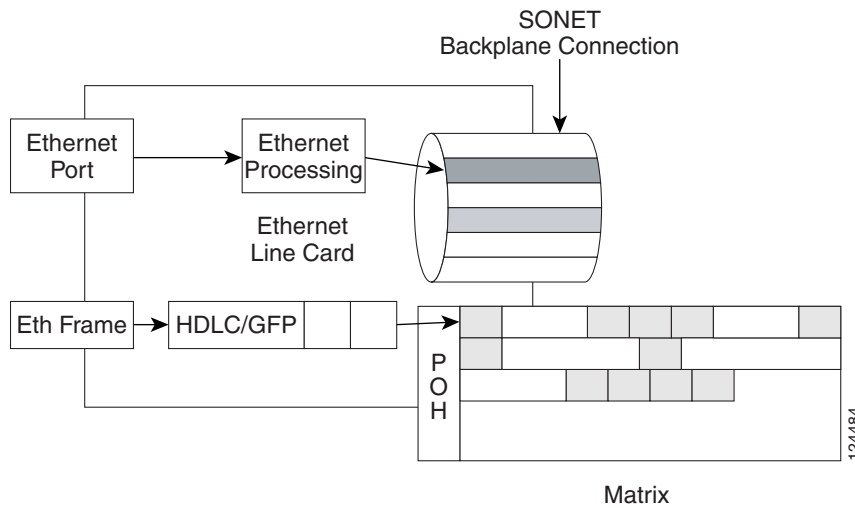


The ONS 15600 provides transport functionality of all frames arriving on the Ethernet interfaces. Frames arriving on each Ethernet interface are mapped onto their own SONET path. In the reverse direction, frames arriving on a SONET path are mapped to one and only one Ethernet interface. Frames arriving on different Ethernet interfaces are not mapped onto the same SONET path.

Valid received Ethernet frames are encapsulated and transported across the SONET network. There are no modifications made to any bits in the frame that is delivered at the destination Ethernet port. Because invalid frames are dropped, this transport is not transparent.

Figure 10-2 shows Ethernet framing with HDLC/GFP header/trailers and mapping into a synchronous payload envelope (SPE) for transport over a SONET network.

Figure 10-2 Ethernet Framing



A valid frame is defined as one with the following attributes:

- Valid preamble + synchronization
- Valid size
- Valid cyclic redundancy check (CRC)
- Valid interpacket gap

10.3 Ethernet Rates and Mapping

This section explains the Ethernet frame format, encapsulation methods, path and circuit configurations, and oversubscription.

10.3.1 Frame Size

The IEEE 802.3 frame format is supported on the GE interfaces. The minimum frame size supported by the system is 64 bytes; the interface drops smaller sized frames. The maximum frame size per interface is 9600 bytes and can be provisioned independently. 9600 bytes is the maximum transmission unit (MTU).

10.3.2 Encapsulations

The Ethernet ports can be provisioned for encapsulation in frame-mapped GFP (GFP-F) as per ITU G.7041, Cisco HDLC/LEX as per RFC 1841, or Ethernet over Link Access Protocol over SDH (LAPS) as per ITU X.86. LEX with CRC-32 is the default encapsulation, as it is supported by most Cisco data cards.

The value of the C2 byte in a high-order SONET path is set to the following values for each encapsulation type:

- 0x1B for GFP-F
- 0x18 for X.86
- 0x05 for LEX with CRC-16
- 0x01 for LEX with CRC-32



Note

Bridge Control Protocol (BCP), HDLC, LEX, and ITU X.86 are all variations of Point-to-Point Protocol (PPP) in HDLC framing. To the Ethernet processor, the packets look essentially the same. GFP encapsulation is drastically different.

10.3.3 Path and Circuit Sizes

GE interfaces can be mapped to all supported SONET paths of sizes STS-1, STS-3c, STS-6c, STS-9c, STS-12c, STS-24c, and STS-48c.



Note

Mapping depends on the particular concatenation being supported on the card. It also depends on the concatenation being supported on all of the network elements (NEs) that the end-to-end circuit is being built through.

For full line-rate mapping to ensure that no frames are dropped, a GE interface is mapped to either an STS-24c or an STS-48c. At the STS-24c rate, the effects of bandwidth expansion due to HDLC/GFP/ITU X.86 encapsulations need to be taken into account to ensure no packet loss.

10.3.4 Oversubscription

Oversubscription of data interfaces is the mapping of an interface to a path with a lower bandwidth than the rate of the interface. When oversubscribing, the operator is relying on the fact that on data interfaces, packet transmission rates and utilization is not always 100 percent. For example, a customer with a GE interface might transmit at an average rate of only 100 Mbps.

To ensure that packet loss does not occur or that any ensuing packet loss is minimized, system engineers need to carefully estimate the amount of memory in the NE that is required to sustain any traffic bursts. Traffic bursts are characterized by several different models that take into account such parameters as the mean rate, peak rates, length of burst, and so on. Performing these calculations is more important in switching applications than in transport applications.

In the ONS 15600, oversubscription of interfaces is supported. This occurs when the GE interface is mapped to any of the STS-1, STS-3, STS-6, STS-9, and STS-12 (c) rates. No guarantees on frame loss or delay due to oversubscription is provided by the system. Actual loss and delay values depend on incoming Ethernet traffic patterns.

10.4 Protocols over Ethernet

The ASAP card supports the BCP, PPP Half Bridge, and VLAN protocols, described in the following subsections.

10.4.1 Bridge Control Protocol

Support of BCP is relevant between two routers or bridges. An Ethernet card providing transport functionality is neither a router nor a bridge. Therefore, even BCP encapsulation is not supported on the Ethernet ports of the ASAP transport line card.

The ASAP line card Ethernet function is transparent to any Layer 2 and above protocol packets (the entire control plane). Any BCP control packets are transported out to the SONET interface.

10.4.2 PPP Half Bridge

For situations in which a routed network needs connectivity to a remote bridged Ethernet network, a serial or integrated services digital network (ISDN) interface can be configured to function as a PPP half-bridge. The line to the remote bridge functions as a virtual Ethernet interface, and the router's serial or ISDN interface functions as a node on the same Ethernet subnetwork as the remote network.

The bridge sends bridge packets to the PPP half-bridge, which converts them to routed packets and forwards them to other router processes. Likewise, the PPP half-bridge converts routed packets to Ethernet bridge packets and sends them to the bridge on the same Ethernet subnetwork.

The ASAP line card is transparent to any Layer 2 and above protocol packets (the entire control plane). Any PPP Half Bridge control packets are transported out to the SONET interface.

10.4.3 VLAN

The ASAP line card is transparent to VLAN tag information in the Ethernet frame.



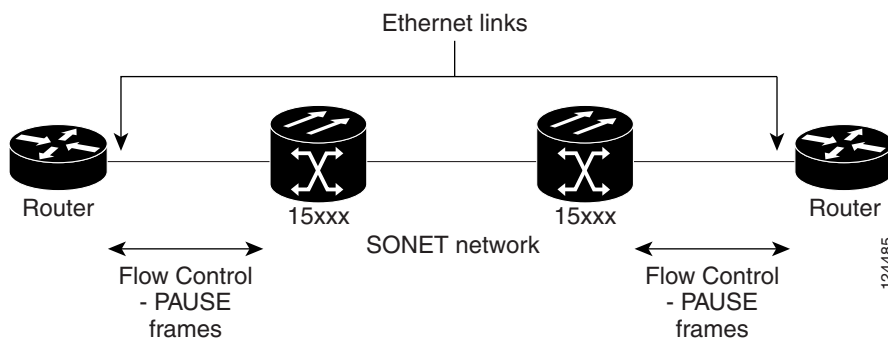
Note

VLANS are tunneled, and not terminated.

10.5 Buffering and Flow Control

Because the STS circuits can often be oversubscribed (have a bandwidth lower than that needed to support the traffic on the Ethernet port), a combination of buffering and local flow control is supported. Initially, frames that arrive on an ingress interface that cannot be transmitted immediately on the egress interface are placed in an ingress buffer. When this buffer starts filling up and is in danger of overflowing, flow control is employed. Local flow control is flow control between the Ethernet interface and the router or switch it is connected to, over the single Ethernet link. This is depicted in [Figure 10-3](#).

Figure 10-3 Buffering and Flow Control



To prevent dropping of frames on an ingress Ethernet interface due to buffer congestion, an Ethernet interface sends a PAUSE frame to its peer on its egress Ethernet interface. An Ethernet interface might be capable of both sending and receiving PAUSE frames (symmetric flow control). On the other hand, it might be capable of performing only one of the two (asymmetric flow control). Two factors determine what an Ethernet interface ultimately ends up supporting:

- Flow control capability of the interface as provisioned by the operator (fixed at OFF or ON)
- Flow control capability of the peer as determined by AutoNegotiation

The Ethernet interfaces are capable of using symmetric or asymmetric flow control to meter packet receptions according to the requirements specified in IEEE 802.3x. This is done to avoid dropping packets internally due to output or input queue congestion. When an Ethernet interface uses symmetric flow control, it generates and obeys pause frames. When it uses asymmetric flow control, it only generates pause frames. Through management control, the Ethernet interface allows an operator to turn flow control ON or OFF (the default is ON).

When an Ethernet interface is set to be capable of generating PAUSE frames, it starts generating them at a rate of X PAUSE frames per second when the ingress buffer passes above a certain level called the High Water Mark. A PAUSE_TIME is specified in each PAUSE frame. The values of the PAUSE_TIME parameter and the rate X parameter are determined by the system and are not user settable.

It is expected that the peer Ethernet interface stops sending new frames after receiving a PAUSE frame for the amount of time as specified in the PAUSE_TIME. When an Ethernet interface is actively generating PAUSE frames, it stops generating them when the ingress buffer passes below a certain level, called the Low Water Mark. It is expected that the peer Ethernet interface starts sending new frames after the PAUSE frames stop.

When an Ethernet interface is set to be capable of receiving PAUSE frames, it temporarily suspends transmission of any new Ethernet frames upon reception of a PAUSE frame from its peer. Note that the transmission of any Ethernet frames already begun are completed.

When an Ethernet interface has temporarily PAUSED transmission of Ethernet frames, it resumes transmission of the frames when no new PAUSE frame is received within the PAUSE_TIME specified in the last received PAUSE frame.

The High Water Mark and Low Water Mark values are both operator provisionable. Legal values are as follows:

- Low Water Mark: from 10 kB to 74 kB
- High Water Mark: from 11 kB to 75 kB

The Low Water Mark is always lower than the High Water Mark.


The ASAP line card utilizes tail drop when frames need to be dropped. This means that the last frames received are the first to be dropped. This is implemented with a Last In First Out (LIFO) buffer.

10.6 Autonegotiation

The Ethernet interfaces on the ASAP card are capable of autonegotiating for full duplex (only) operation as per IEEE 802.3z. The default provisioning is to autonegotiate the duplex mode and persists in memory. Autonegotiation can be provisioned to be ON or OFF.

When flow control is turned on (and so is autonegotiation), the Ethernet interface autonegotiates and advertises all modes that it supports. The mode used by the Ethernet interface is the one agreed to by the Ethernet peers. If an agreement on the supported flow control mode cannot be reached, the transmitter is turned off. Autonegotiation is depicted in [Figure 10-4](#).

Figure 10-4 Autonegotiation




10.7 Gigabit EtherChannel/IEEE 802.3ad Link Aggregation

The end-to-end Ethernet link integrity feature can be used in combination with Gigabit EtherChannel (GEC) 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 ASAP card supports all forms of link aggregation technologies including GEC, which is a Cisco proprietary standard, and the IEEE 802.3ad standard. The end-to-end link integrity feature ASAP card allows a circuit to emulate an Ethernet link. This allows all flavors of Layer 2 and Layer 3 rerouting to work correctly with the ASAP card. [Figure 10-5](#) illustrates GEC support.

Figure 10-5 ASAP Gigabit EtherChannel (GEC) Support



Although the ASAP 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 the ASAP card to an ONS network, the ONS SONET side network is transparent to the EtherChannel devices. The EtherChannel devices operate as if they are directly connected to each other. Any combination of 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 data links together to provide more aggregated bandwidth. Spanning Tree Protocol (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 nonblocked path. GEC can also provide ASAP 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.