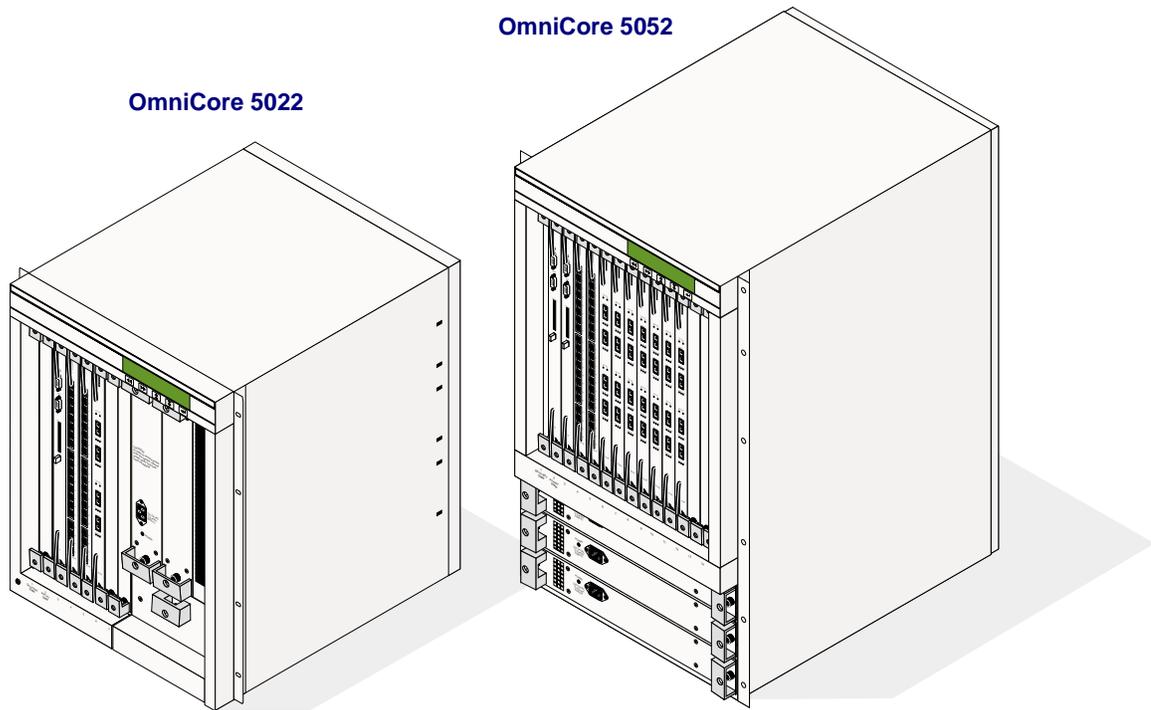


1 OmniCore 5052 and 5022 Overview

The OmniCore 5052 and 5022 routing switches are designed to meet the current and future demands of large enterprise backbones, data centers, and other high-performance network environments.



OmniCore 5052 and 5022 Routing Switch

OmniCore routing switches provide wire-speed, packet-by-packet routing throughput of greater than 37 million packets (5052) and 15 million packets (5022) per second, and its architecture can scale to more than four times this throughput. The OmniCore routing switches' extensive routing capabilities, flexible connectivity options, and policy-based management simplify network design and administration, thereby reducing total cost of ownership. Superior forwarding rates, scalability, and mission-critical resiliency features make OmniCore routing switches the ideal solution for scaling the performance of your local, metropolitan, regional, or wide-area network.

Feature Summary

Each OmniCore routing switch provides a complete solution with Infinite Routing™, switching, high density and flexible connectivity, mission-critical reliability, application-aware networking, and network management.

Infinite Routing Features

The OmniCore routing switch supports wire-speed IP, IPX, and IPv6 routing, static routes, and a host of IP capabilities that reduce the complexities of networking.

IP Routing

The OmniCore routing switch supports IP routing with RIPv1 and RIPv2 and provides support for OSPF and BGPv4. There is support for ICMP and ICMP Router Discovery Protocol (IRDP). RIP is a distance vector protocol that uses the number of hops between routers as its means of determining the best route through the network. Text key authentication is supported for RIPv2, and MD5 authentication is supported for OSPF. Support for security services such as Access Control, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System “plus” (TACACS+) are also provided.

Addresses

There is support for up to 64,000 longest prefix match routes at each gigabit interface or group of ten Fast Ethernet interfaces. These routes are populated and managed from the Enterprise Management Module.

Multinet

Multiple logical sub-networks can exist within a port-based VLAN.

Multiple Route Support

The OmniCore routing switch supports multiple routes per interface and use the optimal route when making a routing decision. However the longer routes are still kept in memory for use if the optimal route fails. This capability allows network managers to create resilient topologies.

Honoring TOS Precedence

Packets entering the OmniCore routing switch have their TOS precedence values strictly honored within the system. For example, a packet with a TOS precedence value of 5 is treated as priority 5 within the system for queuing purposes and the switch will route that packet accordingly.

Priority Assignment

An inbound physical port or router interface (VLAN) can be assigned a priority (0-7) that is used internal to the system for queuing. This priority can be mapped into the TOS precedence field and/or the 802.1p field at the outbound port for communication with other network devices. Support for classification and assignment of TOS or 802.1p/Q values will be extended to include user-defined policies based on Layer 4, Layer 3, and Layer 2 attributes.

Static Routes with TOS Precedence

Static and default routes are supported, however, host routes are not supported. Multiple entries to the same route destination can also be entered and the system use the lowest cost route first. In the event of a link failure, the next-best route is used. Static routes take precedence over routes learned by dynamic protocols such as RIP and OSPF. For each static route entered, an entry for Type of Service (TOS) can be specified which will cause both a path selection and a priority queue selection. The result is that packets can be forwarded to different destinations based on the TOS value and will flow through the system with different priorities.

TOS-Based Routing

Type of Service (TOS)-based routing provides the ability to make forwarding decisions based on the destination address and a desired Quality of Service (QoS). When multiple links are available, a decision can be made based on QoS. This feature is supported in static routes and OSPF.

Not-So-Stubby-Area

As defined in RFC 1587, the OmniCore routing switch supports Unsays, allowing the creation of OSPF areas similar to the stub area, with the additional capability of importing AS external routes in a limited fashion.

IPX Routing

The OmniCore routing switch supports routing of IPX packets with RIP and SAP. It offers two types of IPX routing. The first type of IPX routing is in hardware and at wire-speed (also known as fast-path). The second way is traditional software-based routing. In this mode, IPX packets are forwarded to the EMM where they are processed by its software routing engine. This is commonly referred to as slow-path routing. There are very specific configurations that dictate when a packet is handled via fast-path or slow-path. These conditions are outlined below. However software-based routing and hardware routing cannot coexist on the same interface.

Multiple Route Support

The OmniCore routing switch supports multiple routes, and uses the optimal route as determined by IPX RIP/SAP. However, the longer route is still kept in memory for use if the optimal route fails. This capability allows network managers to create resilient topologies.

Multiple Encapsulation or Frame Type Support

The OmniCore routing switch's IPX implementation supports the following IPX frame types:

- Ethernet version II (ETHERNET_II)
- Ethernet 802.3 raw frames with no LLC layer (ETHERNET_802.3)
- Ethernet 802.3 frames with LLC (ETHERNET_802.2)
- SNAP frames (ETHERNET_SNAP)

The switch supports translational bridging for these different encapsulation types, but does not do so at wire-speed. Instead, translations between different encapsulation types are handled via slow path.

IPX – VLAN to VLAN connectivity

The OmniCore routing switch performs wire-speed IPX routing when each port is configured as a Virtual Router with its own VLAN and a separate IPX network number. Note that in this configuration, the Virtual Router interface or VLAN cannot span multiple ports.

Static Routes

Static and default routes can be defined for IPX. These routes are redistributed via IPX RIP.

IPv6 Routing

The OmniCore routing switch will provide support for IPv6 RIPng.

IP Assist Capabilities (IP Helper)

In addition to its ability to route IPv4 packets at wire speed, the OmniCore routing switch also provides capabilities to assist in the administration of IP networks.

DHCP Relay Agent

Dynamic Host Configuration Protocol (DHCP) allows devices to gain their IP address through a DHCP server, which dynamically assigns IP address according to the device's location or user name. The OmniCore routing switch supports DHCP by forwarding DHCP discover or BOOTP request packets to a DHCP server.

The OmniCore routing switch supports multiple DHCP server addresses and forwards DHCP discover packets and BOOTP requests to multiple DHCP servers for redundancy. It stamps every incoming DHCP discover or BOOTP request packet indicating which scope (subnet) the packet originated from. The DHCP server then uses that information to determine the IP address.

To support multinet, the OmniCore routing switch stamps the gateway field of the DHCP request packet with the router interface's IP address. In multinet configurations, a round robin scheme is used for each logical subnet.

Switching Features

The OmniCore routing switch supports Spanning Tree, port trunking, VLANs, IEEE 802.1Q tagging, and IEEE 802.1p priority.

Addresses

The OmniCore routing switch is capable of learning 64,000 addresses per gigabit port or per group of ten Fast Ethernet ports.

Spanning Tree

The OmniCore routing switch supports bridging as defined in the IEEE 802.1d standard and Fast STP (FSTP) as defined in the IEEE 802.1w standard. FSTP protocol is active in the routing switch when STP is enabled. Bridge mode can be configured to specify a single spanning tree or a per-VLAN spanning tree.

Port Trunking

The OmniCore port trunking algorithm is designed for inter-switched links. OmniCore port trunking allows four ports of the same type and configuration to be combined to form a single logical trunk. Up to four trunk groupings can be created in a single OmniCore chassis. Spanning Tree (802.1d) is not turned off for OmniCore trunked ports.

Link Aggregation Control Protocol (LACP)

The OmniCore routing switch supports LACP as defined in the IEEE 802.3ad standard. LACP is the control protocol which establishes and maintains link aggregation. Link aggregation is a method of combining multiple links between systems to increase bandwidth, enhance resiliency, and provide load sharing.

Virtual LANs (VLANs)

The OmniCore routing switch supports 512 VLANs based on port, protocol, MAC, IEEE 802.1Q tag, or Layer 3 address. These VLANs are defined as follows:

- **Port-based** Defined as a port on a switch or a group of ports on a switch.
- **Protocol** A group of users using the same protocol.
- **MAC** A group of MAC addresses.
- **Tagged** Using the IEEE 802.1Q standard for VLAN tagging, the switch can map users into another vendor's VLANs, or create its own VLAN tags and pass them on to other devices. There is support for the full range of IEEE 802.1Q tags, from 0 to 4,095. When tagging is enabled, multiple VLANs can be trunked together over a single link and the switch separates the different VLANs at their ingress. However, if tagging is not enabled, then the switch drops tagged packets.
- **Layer 3 address** Based on the subnet or IPX network number, the switch can create a VLAN.

Hot-Standby Routing Protocol

The OmniCore routing switch supports HSRP (Hot Standby Routing Protocol) providing hosts with static routes protection against router failure. The HSRP protocol detects when the designated active router fails. It then selects a standby router to assume control of the HSRP group's MAC and IP addresses and, at the same time, selects a new standby router. Multiple HSRP groups can be configured on an interface, up to 256 groups per VLAN.

High Density and Flexible Connectivity Features

The OmniCore routing switch supports 10/100/1000 Mbps Ethernet, Fiberoptic Digital Data Interface (FDDI), and Packet-Over-SONET (POS) interface modules, and will support ATM.

10/100/1000 Mbps Ethernet Interface Modules

The OmniCore routing switch supports up to 240 (5052) or 100 (5022) auto sensing 10/100 BASE-TX ports, and up to 120 (5052) or 50 (5022) 100BASE-FX (MMF) ports.

The ports of the 10BASE and 100BASE modules support auto-negotiation in both half- and full-duplex, the ports of the 1000BASE modules support auto-negotiation full-duplex mode, while the 100BASE-FX ports support 100 megabit operation only with manual full/half duplex configuration. The Gigabit Ethernet interfaces run at 1000 Mbps full-duplex (per the IEEE 802.3z standard).

All modules, whether running at 10 Mbps, 100 Mbps, or 1000 Mbps, provide flow control (IEEE 802.3x) using pause frames when running in full-duplex mode, and jam frames when

running in simplex mode. In addition, all modules support Layer 2 switching and Layer 3 routing on a packet-by-packet basis. By default, flow control is not enabled. See [Interface Modules](#) on page 1-15 for a complete list of OmniCore interface modules and part numbers.

FDDI Interface Modules

The OmniCore routing switch supports a 2-port Dual Attach Station (2-Port DAS) interface module. This module also functions as a 2-port Single Attach Station (SAS).

The Fiberoptic Digital Data Interface (FDDI) module supports wire-speed routing of IPv4, IPv6, and IPX packets. When FDDI frames enter the module, they are fragmented into Ethernet frames and switched or routed. Ethernet frames are translated into FDDI frames, but not defragmented, meaning that 1,518 byte frames are put on the FDDI ring as a 1,518 byte frames and not aggregated with other Ethernet frames to make a single large FDDI frame. See [Interface Modules](#) on page 1-15 for a complete list of OmniCore interface modules and part numbers.

POS Interface Modules

The OmniCore routing switch supports interface modules for Packet-over-SONET (POS) at OC-3c (155 Mbps) and OC-12c (622 Mbps) rates. These modules support RFC 1619, PPP over SONET/SDH, RFC 1662, PPP and HDLC Framing, and will support SONET Automatic Protection Switching (APS). These modules also utilize the same Application Specific Integrated Circuits (ASICs) found on the 10/100/1000 interface modules to provide equivalent features and performance for POS links. See [Interface Modules](#) on page 1-15 for a complete list of OmniCore interface modules and part numbers.

Mission-Critical Reliability Features

The OmniCore routing switch has advanced resiliency features that ensure its availability—even in the event of a component failure. These features include:

- Load sharing, hot-swappable power supplies.
- Redundant physical interfaces.
- Load sharing, hot-swappable management modules.
- Redundant, hot-swappable cooling devices (fans).
- Trunking of interfaces with graceful failover in the event of link failure.
- Hot-swappable, N+1 redundant memory fabrics (5052 only).

Redundant Memory Fabric (5052 Only)

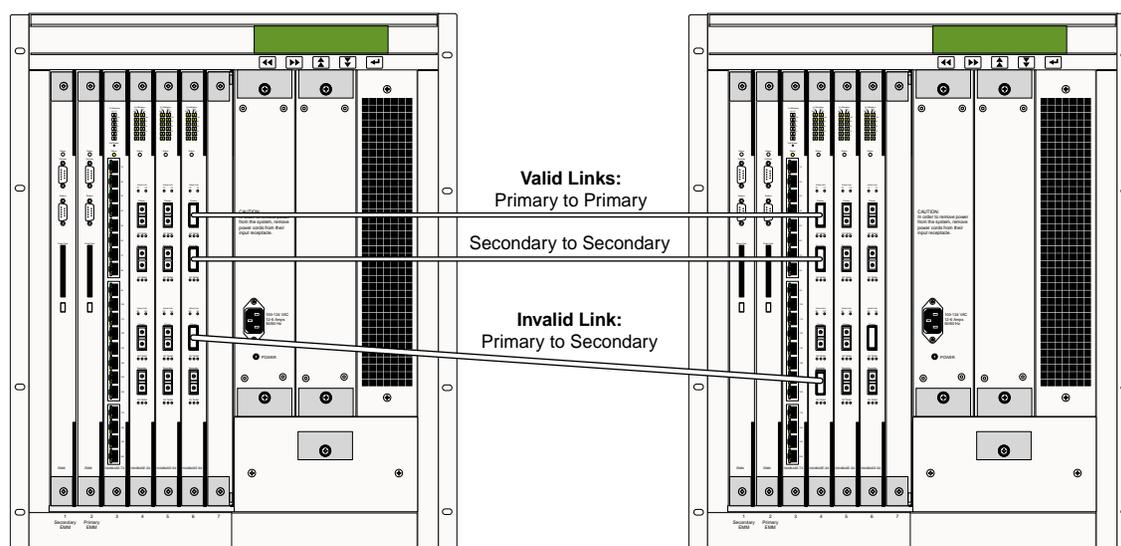
The memory fabric of the OmniCore routing switch provides N+1 redundancy. In the event that one of the eight memory modules fails, the redundant module will become live in less than 400 milliseconds. The switch will generate an alarm if one of these modules fails. To replace a failed module, the rear panel of the chassis can be removed exposing the memory fabric of the switch. Each memory module has an LED indicating its operating status. To replace the memory fabric (rear cards), see [Appendix A, "Servicing Components"](#).

Redundant Enterprise Management Module (EMM)

The OmniCore routing switch's slots 1 and 2 are reserved for Enterprise Management Modules (EMMs). One EMM is required for operation and the second EMM provides redundancy. Only slot 2 can be populated with an EMM with a Gigabit Ethernet interface. Note that in the 5052 switch only, should the EMM in slot 1 take over, the gigabit port on the module in slot 2 will continue to be fully functional. The redundant EMM works by continually polling the switch's control bus. If the primary EMM fails, the secondary EMM assumes control in a matter of milliseconds, restarts the switch, and resumes normal operation.

Redundant Physical Interfaces

The OmniCore routing switch supports Gigabit Ethernet interface modules with redundant physical links. These redundant links share the same MAC address. The signal transmitted on redundant links is identical to the signal on the primary interface. The primary Gigabit Ethernet interface is continually checking itself for link integrity pulses. If the primary interface loses link integrity, the secondary interface takes over in less than 100 milliseconds. In this case, the primary link can be serviced without affecting traffic on the secondary link. Please note that redundant links should connect only to other redundant links as shown below.



Valid Use of Redundant Links

Application-Aware Networking Features

The OmniCore routing switch protects applications with a combination of application-based filtering and forwarding. This application-aware functionality is an extension of traditional QoS capabilities. The OmniCore switch's QoS and traffic management features are outlined below. (Note that many of the OmniCore switch's specific QoS features such as honoring TOS precedence, priority assignment, and TOS routing are discussed in [IP Routing](#) on page 1-35.)

QoS Policy Management

The OmniCore routing switch has advanced Quality of Service (QoS) capabilities embedded in hardware to allow wire-speed analysis of Layer 2, Layer 3, Layer 4, and application attributes. It can route traffic according to a broad set of management policies using these capabilities. The OmniCore routing switch will provide predefined templates to ease configuration of these policies.

Layer 2 QoS 802.1p/Q

The OmniCore routing switch is capable of interpreting Layer 2 VLAN (802.1Q) and priority (802.1p) tags. It uses this information to place packets into its eight priority-based queues. It will be able to look up to 256 bytes deep into a packet and apply any QoS policies related to specific applications. These capabilities can be turned on using the CLI or activated from TrackView, Alcatel's graphical user interface for configuring and administering policies on OmniCore routing switches.

Layer 3 QoS IP Type of Service (TOS)/Precedence

The Type Of Service (TOS) field is found in all IP packets and is the only standards-based way of providing QoS across IP networks. TOS-based Quality of Service is used to select the best route through an IP network when more than one path exists between a source and a destination. Routing table convergence is automated through the use of OSPF, which automatically maps TOS bits to form a unique, dynamic topology. In addition, TOS can use static route information for determining the best route through the network.

In addition to the TOS field, IP packets also have a 3-bit precedence field that defines eight different priority queues. These precedence queues work very well with the OmniCore routing switch's eight priority queues per port. Similar to the systems' ability to stamp 802.1p at line rate, the OmniCore routing switch can also pass on precedence priority to existing routers. This can currently be done on a port or router interface (VLAN) basis and will be extended to applications.

Queuing Mechanisms

An OmniCore routing switch can service its eight-per-port priority-queues in three different ways: WFQ, Strict, and PBQ. WFQ (Weighted Fair Queuing) and Strict priority modes provide "canned" settings for priority queue bandwidths, burst enabling, and burst bandwidth for the port. PBQ (Programmable Bandwidth Queuing), available on the latest generation of interface modules, allows the user to modify the default WFQ settings for each of the eight individual priority queues.

Network Management

The OmniCore routing switch provides proactive network management using SNMPv1, four-groups of RMON (statistics, history, alarms, and events), and wire-speed port mirroring to aid in troubleshooting the network. In addition, routing switch configuration and management can occur through the following mechanisms: Telnet, Java-based GUI with TrackView management, a Command Line Interface (CLI), and HP OpenView. All switch management is completely tied to SNMP since the switch uses public and private MIBs for management.

Command Line Interface (CLI)

The CLI can be accessed via the RS-232 console port found on an EMM interface module or by using Telnet. The CLI allows users to configure services and monitor the operation of the system. By default, up to five Telnet sessions can operate concurrently.

The OmniCore routing switch provides a full suite of commands that allow you to create and modify user names for access to the CLI. Support for security services such as Access Control, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System "plus" (TACACS+) are also provided.

TrackView Network Management

TrackView is a Java-based application for managing the OmniCore routing switch. By using Java, TrackView management is platform-independent and can be run on almost any platform. TrackView management extends upon the functionality found in the CLI, simplifying configuration tasks, such as hardware configuration, routing configuration, bridging configuration, VLAN membership, and policy definition. TrackView management will also be capable of being run within web-based browsers.

Simple Network Management Protocol

The OmniCore routing switch can be managed entirely through public and private Simple Network Management Protocol (SNMP) MIBs. The private MIBs are Alcatel's proprietary extensions while the public MIBs are:

- RFC 1157 SNMP
- RFC 1213 MIB-II
- RFC 1493 Bridge MIB
- RFC 1512 FDDI MIB
- RFC 1567 BGP-4 MIB
- RFC 1643 Ethernet MIB
- RFC 1724 RIPv2 MIB
- RFC 1742 AppleTalk MIB
- RFC 1757 RMON (4 groups)
- RFC 1850 OSPF MIB
- RFC 2021 RMON2 MIB
- RFC 2037 Entity MIB
- RFC 2096 IP Forwarding Table MIBs
- RFC 2206 RSVP MIB
- Multicast Routing MIBs: IGMPv2, DVMRP, PIM-SM

All MIBs, including Alcatel's proprietary MIBs, can be viewed from any MIB browser.

RMON

The OmniCore routing switch supports four groups of RMON I on a per-port basis. These groups are described below.

Statistics

Statistics provides the ability to report on performance metrics, packet size distribution, and errors.

History

The OmniCore provides history for short-term and long-term statistics. The short term group and long-term group have 50 separate buckets that can be allocated to record different statistics at user-defined intervals (between 1 second and 1 hour). This information can be used to monitor network performance and activity to provide proactive network management.

Alarms

Each metric taken in a statistics group can be configured for thresholds that trigger alarms. For instance, if too many runt packet sizes are generated over a period of time, an alarm can be triggered. The alarm triggers an SNMP trap, which is logged for future review. In addition, an alarm can trigger any executable, such as a paging script to alert a network manager.

Events

Events allow OmniCore routing switch users to configure and monitor what generates an alarm on the switch. It allows users to create their own thresholds that triggers alarms based on any statistical information.

RMON Probe (Port Mirroring)

One of the benefits of *Parallel Access Shared Memory* is the ability to provide wire-speed port mirroring without affecting switch performance. Therefore, an RMON probe can be used to gather RMON statistics other than Statistics, History, Alarms, and Events.

Network Time Protocol and Timezone

The OmniCore routing switch supports Network Time Protocol (NTP) used to synchronize the time of a computer client or server to another server or reference time source that is synchronized to Coordinated Universal Time (UTC). Time zone information is required for the translation of UTC time to local time. Alcatel's CLI Timezone implementation provides the interface for establishing any time zone in the world.

OmniCore 5052 and 5022 Relevant Standards

- IEEE 802.3z 100BASE-X Gigabit Ethernet
- IEEE 802.3u 100BASE-TX, 100BASE-FX Fast Ethernet
- IEEE 802.3 10BASE-T Ethernet
- IEEE 802.3x Full-Duplex with Flow Control
- IEEE 802.1D Spanning Tree Protocol
- IEEE 802.1Q VLAN Tagging, GVRP
- IEEE 802.1D-1998 Priority and Dynamic Multicast Filtering, GARP, GMRP
- RFC 768 UDP
- RFC 791 IP
- RFC 793 TCP
- RFC 854 Telnet
- RFC 1035 DNS
- RFC 1058, 1723 RIP and RIPv2
- RFC 1122 Host Requirements
- RFC 1157 SNMP
- RFC 1213 MIB-II
- RFC 1256 RDP
- RFC 1332 IPCP
- RFC 1350 TFTP
- RFC 1390 FDDI
- RFC 1493 Bridge MIB
- RFC 1512 FDDI MIB

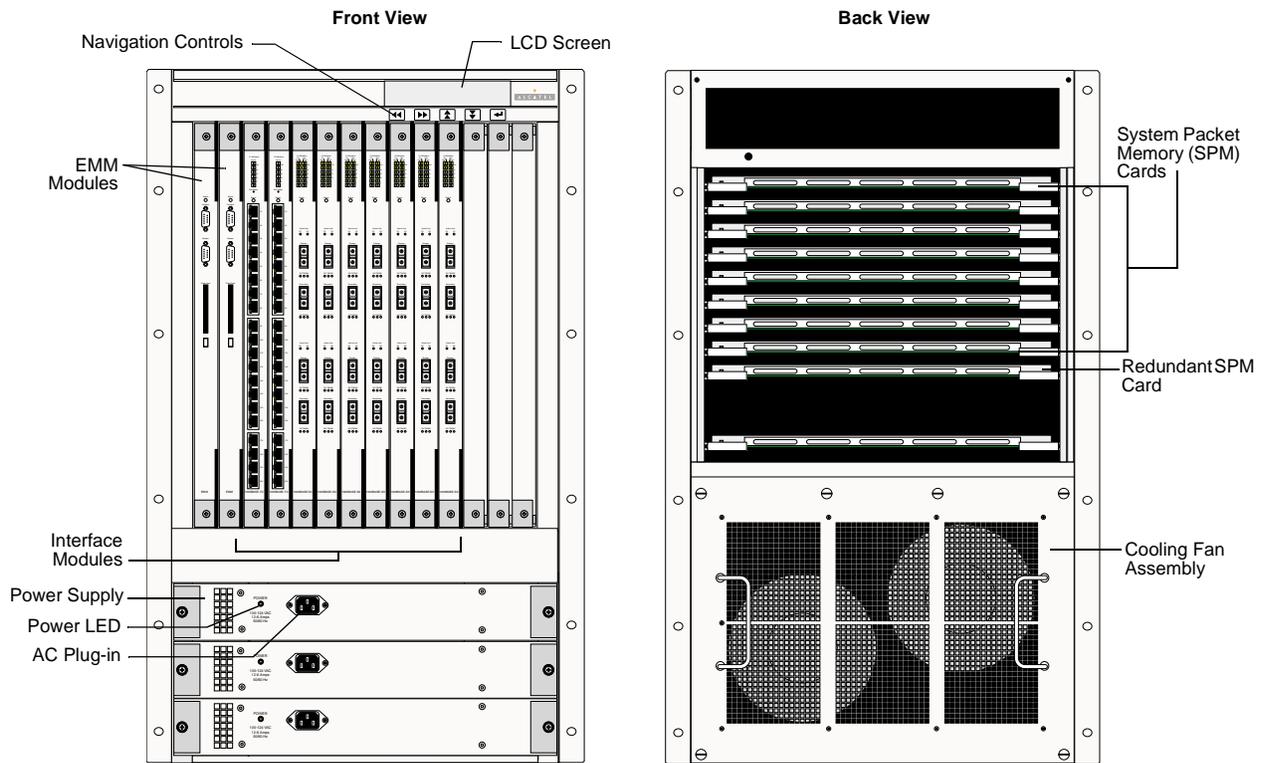
- RFC 1542 BOOTP
- RFC 1583, 1587, 2178 OSPF
- RFC 1619 PPP over SONET/SDH
- RFC 1643 Ethernet MIB
- RFC 1657 BGP-4 MIB
- RFC 1661 The Point-to-Point Protocol (PPP)
- RFC 1662 PPP in HDLC-like Framing
- RFC 1724 RIPv2 MIB
- RFC 1742 AppleTalk MIB
- RFC 1757 RMON (4 groups)
- RFC 1812 Router Requirements
- RFC 1850 OSPF MIB
- RFC 2021 RMON2 MIB
- RFC 2131 DHCP
- RFC 2138 RADIUS
- RFC 2037 Entity MIB
- RFC 2080 RIPng
- RFC 2096 IP Forwarding Table MIBs
- RFC 2206 RSVP MIB
- RFC 2236 IGMP and IGMPv2
- Multicast Routing MIBs: IGMPv2, DVMRP, PIM-SM, MOSPF IP
- ISO 9314-1 (PHY), 9314-2 (PMD), 9314-3 (MAC) FDDI
- ANSI FDDI X3T9.5 SMF-PMD Standard, SMT Specification Version 7.3
- ANSI T1.105-1995 SONET – Basic Description including Multiplexing Structure
- ANSI T1.105.02-1995 SONET – Payload Mappings
- Bellcore GR-253-CORE SONET Transport Systems Common Generic Criteria
- ITU-T G.709 Synchronous Multiplexing Structure

Hardware Descriptions

This section features hardware descriptions for both the OmniCore 5052 and 5022 routing switches.

OmniCore 5052

The OmniCore 5052 routing switch is a chassis-based product with 14 slots: 12 slots for interface modules, and two for Enterprise Management Modules (EMMs). The OmniCore 5052 switch also supports a maximum of three power supplies. The following illustrations display the major components of the OmniCore 5052.



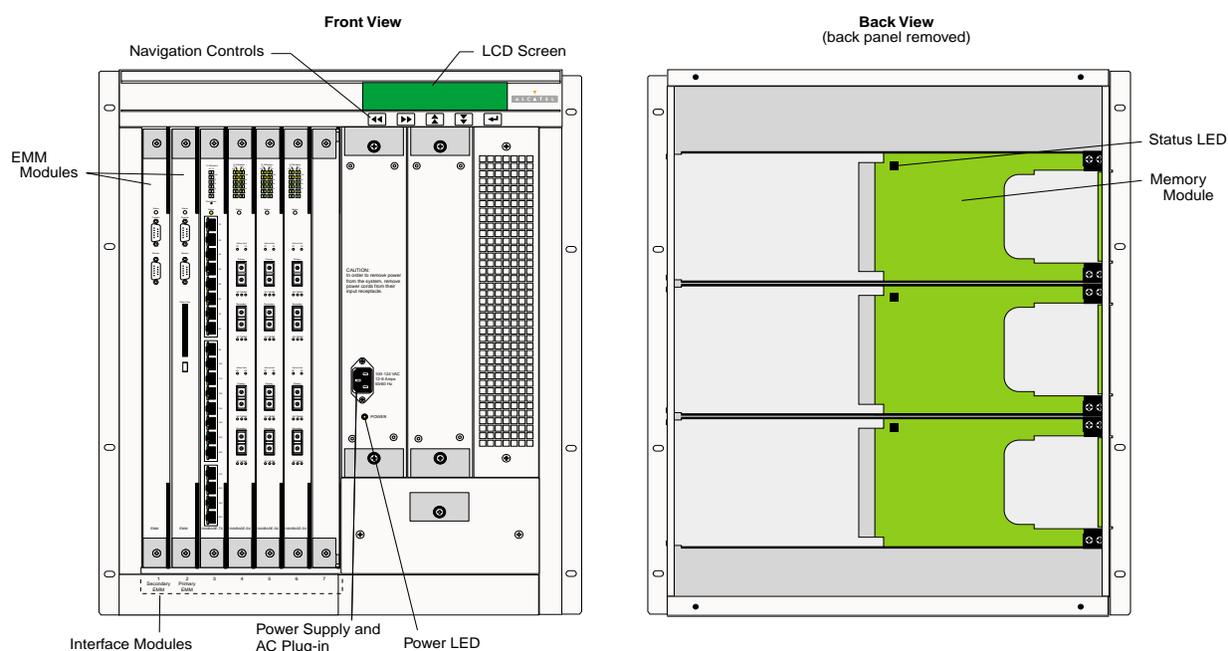
OmniCore 5052: Front and Back Views

OmniCore 5052 System Packet Memory Cards

The OmniCore 5052 routing switch has nine System Packet Memory (SPM) slots. These are situated in the rear of the switch. A total of eight SPM modules are required for proper operation; the ninth slot is provided for redundancy. All SPM cards are hot-swappable and include a status LED to indicate the operating status of each card. See previous figure to view System Packet Memory card placement.

OmniCore 5022

The OmniCore 5022 routing switch is a chassis-based product with seven slots: five slots for interface modules, and two for Enterprise Management Modules (EMMs). The OmniCore 5022 switch also supports a maximum of two power supplies. The following illustration displays the major components of the OmniCore 5022 routing switch.



OmniCore 5022: Front and Back Views

OmniCore 5022 System Memory Modules

The OmniCore 5022 routing switch has two memory modules. These are situated in the rear of the switch. Both memory modules include a status LED to indicate the operational status of each module. See figure [OmniCore 5022: Front and Back Views](#) on page 1-13 to view system memory module placement.

Power Supply

The OmniCore routing switch comes equipped with up to three (5052) or two (5022) AC or DC power supplies. The OmniCore 5052 switch requires a minimum of two power supplies, and a third power supply can be installed for redundancy. For the OmniCore 5022 switch, a minimum of one power supply is required, and a second power supply can be installed for redundancy. OmniCore power supplies act in a load-sharing manner and are hot-swappable. Each power supply also includes a status LED indicating power supply operational status.

System Safety Features

The OmniCore routing switch provides additional features to keep the unit operating safely. These are described below.

Environmental Controls

The OmniCore routing switch employs thermocouples to monitor the system's temperature. You can view the system's internal temperature using the LCD display (an optional feature for the OmniCore 5022 switch). Note that the internal temperature is not the ambient air operating temperature specified in [Appendix C, "Technical Specifications"](#) and that the internal temperature is always higher than the operating temperature.

Over-Temperature Automatic Shutdown

Should the OmniCore routing switch exceed its maximum operating temperature, the unit will automatically begin a shut down process. Once the switch has cooled to normal operating temperature, it will power up again. If the over-temperature condition has not been corrected, the switch will continue to cycle the power off and on to avoid hardware damage.

Cooling Fans

The OmniCore routing switch features a hot-swappable fan unit which contains two fans. Only one of these fans is required for proper operation; the second is supplied for redundancy.

Interface Modules

The OmniCore routing switch supports a broad range of hot-swappable interface modules. The following table lists all available interface modules with their part numbers. Any combination of interface modules can be plugged into the OmniCore chassis. However, slots 1 and 2 are reserved for Enterprise Management Modules (EMMs).

◆ Note ◆

On rare occasions an interface module may need to be removed and reinstalled. When doing so, remove its copper or fiber links (connectors) before removing the line card. Remove the line card and wait at least 5 seconds before reinstalling the interface module, then reinsert the links.

For versatility, some of the interface modules used in the OmniCore routing switch implement Gigabit Interface Converter (GBIC) modules for fiber optic coupling. For more information, see [Gigabit Interface Converter \(GBIC\) Modules](#) on page 1-17.

OmniCore Interface Modules

Interface Module	Part Number	For Details
<i>1000BASE Interface Modules</i>		
2-Port 1000BASE-SX Gigabit Ethernet	OC-5000-1058	see page 1-18
2-Port/2-Redundant Port 1000BASE-SX Gigabit Ethernet	OC-5000-1061	see page 1-19
2-Port 1000BASE-LX Gigabit Ethernet	OC-5000-1063	see page 1-21
2-Port/2-Redundant Port 1000BASE-LX Gigabit Ethernet	OC-5000-1062	see page 1-22
2-Port 1000BASE-LH Gigabit Ethernet	OC-5000-1067	see page 1-23
2 GBIC-Port/2-Redundant GBIC-Port, 1000BASE Gigabit Ethernet ^a , Large Table (LT)	OC-5000-1079	see page 1-30
2 GBIC-Port/2-Redundant GBIC-Port, 1000BASE Gigabit Ethernet ^a , Small Table (ST)	OC-5000-1099	see page 1-30
6-Port 1000BASE-SX Gigabit Server (legacy)	OC-5000-1068	see page 1-20
6-Port 1000BASE-SX Gigabit Server (version 2.6.0 or newer software)	OC-5000-1069	see page 1-20
6 GBIC-Port 1000BASE Gigabit Server ^a , ST	OC-5000-1104	see page 1-31
6 GBIC-Port 1000BASE Gigabit Server ^a , LT	OC-5000-1093	see page 1-31
<i>10/100BASE and 100BASE Interface Modules</i>		
10-Port 100BASE-FX Fast Ethernet, Multimode Fiber (MMF)	OC-5000-1065	see page 1-24
10-Port 100BASE-FX Fast Ethernet, Single Mode Fiber (SMF)	OC-5000-1072	see page 1-24
20-Port 100BASE-FX Fast Ethernet, MMF	OC-5000-1080	see page 1-25
20-Port 100BASE-FX Fast Ethernet, MMF, Large Table (LT)	OC-5000-1097	see page 1-25
20-Port 100BASE-FX Fast Ethernet, MMF, Small Table (ST)	OC-5000-1101	see page 1-25

OmniCore Interface Modules (Continued)

20-Port 100BASE-FX Fast Ethernet, SMF, LT	OC-5000-1102	see page 1-25
20-Port 100BASE-FX Fast Ethernet, SMF, ST	OC-5000-1103	see page 1-25
20-Port 10/100BASE-TX Fast Ethernet (legacy)	OC-5000-1056	see page 1-26
20-Port 10/100BASE-TX Fast Ethernet (version 2.6.0 or newer software)	OC-5000-1057	see page 1-26
20-Port 10/100BASE-TX, LT	OC-5000-1096	see page 1-26
20-Port 10/100BASE-TX, ST	OC-5000-1100	see page 1-26
OC-3c POS Interface Modules		
2-Port OC-3c POS, Single Mode Fiber–Intermediate Reach (SMF-IR)	OC-5000-1059	see page 1-27
2-Port OC-3c POS, SMF-IR, LT	OC-5000-1115	see page 1-27
2-Port OC-3c POS, Multimode Fiber–Short Reach (MMF-SR)	OC-5000-1071	see page 1-27
2-Port OC-3c POS, MMF-SR, LT	OC-5000-1116	see page 1-27
OC-12c POS Interface Modules		
2-Port OC-12c POS, SMF-IR	OC-5000-1066	see page 1-28
2-Port OC-12c POS, MMF-SR	OC-5000-1070	see page 1-28
2-Port OC-12c POS, MMF-SR, LT	OC-5000-1098	see page 1-28
2-Port OC-12c POS, SMF-IR, LT	OC-5000-1114	see page 1-28
DAS/SAS Interface Module		
2-Port DAS/SAS FDDI	OC-5000-1064	see page 1-29
Enterprise Management Modules (EMM)		
EMM, 64 MB	OC-5000-1060	see page 1-32
EMM, 256 MB	OC-5000-1119	see page 1-32
EMM with 1000BASE-SX Gigabit Ethernet (non-GBIC) Port, 64 MB	OC-5052-1055	see page 1-33
EMM with 1000BASE-SX Gigabit Ethernet (non-GBIC) Port, 256 MB	OC-5052-1121	see page 1-33
EMM with Gigabit Ethernet GBIC-Port ^b , 256 MB	OC-5000-1095	see page 1-33

a. GBIC modules are available in SX, LX, and LH configurations and are hot-swappable. See [Gigabit Interface Converter \(GBIC\) Modules](#) on page 1-17.

b. See Footnote a. The SX GBIC is shipped standard. The LX and LH are by special order only.

Gigabit Interface Converter (GBIC) Modules

A GBIC module is an optical interface module designed to allow hot plugging into a host interface card without damage to the GBIC module or the host circuitry. Any GBIC module will fit in a socket designed for any other GBIC module, however, customers should use only Alcatel approved GBIC modules. Third party GBIC modules are not guaranteed to work properly.

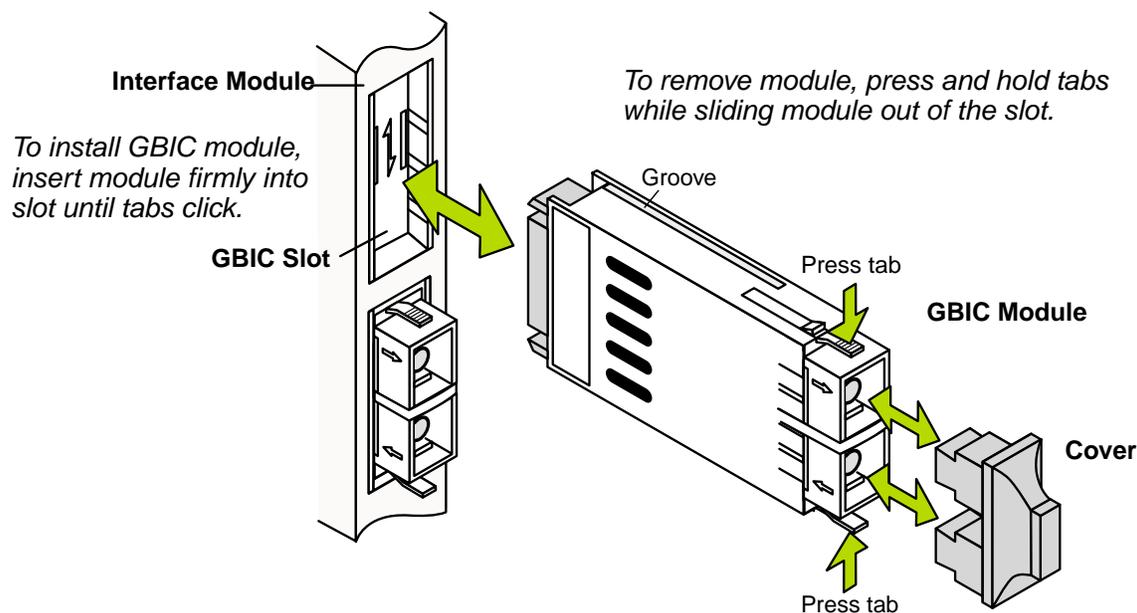
GBIC modules use Duplex-SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using short wavelength (800 nm) lasers. Maximum link distance is 550 meters over 50 μm multimode fiber and 275 meters using 62.5 μm multimode fiber.

The following GBIC modules are available from Alcatel:

- SX, 850NM, MMF, P.N. OC-5000-1109
- LX, 1300NM, SMF, P.N. OC-5000-1110
- LH, 1550NM, SMF, P.N. OC-5000-1111

◆ Caution ◆

Before handling the module, you must discharge all static electricity on your person to avoid Electrostatic Discharge (ESD) damage. Using the wrist strap that came with your OmniCore routing switch, ensure the wrist strap touches your skin. Attach the other end of the strap to the chassis or one of the grounding clips provided on the front or rear panel of the OmniCore chassis.

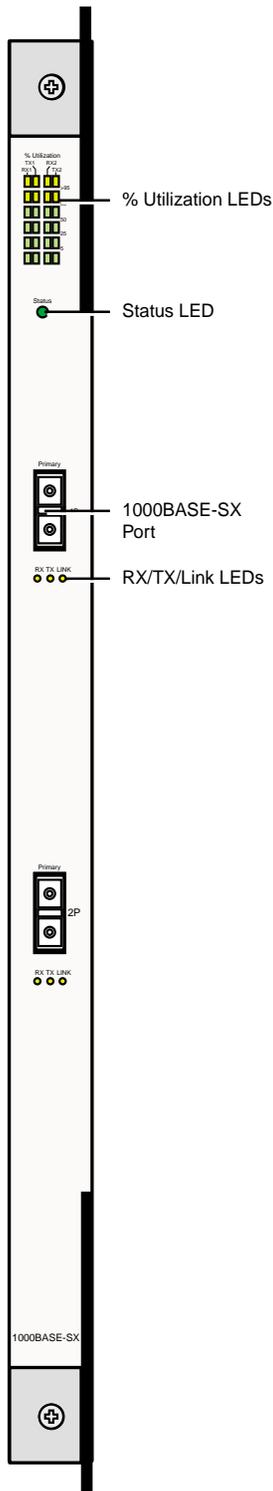


Gigabit Interface Converter (GBIC) Module

2-Port 1000BASE-SX Gigabit Ethernet Interface Module

The 2-port 1000BASE-SX interface module is hot-swappable and features the LED indicators listed in the table below.

The 1000BASE-SX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using short wavelength (800 nm) lasers. Maximum link distance is 550 meters over 50 μm multimode fiber and 275 meters using 62.5 μm multimode fiber.



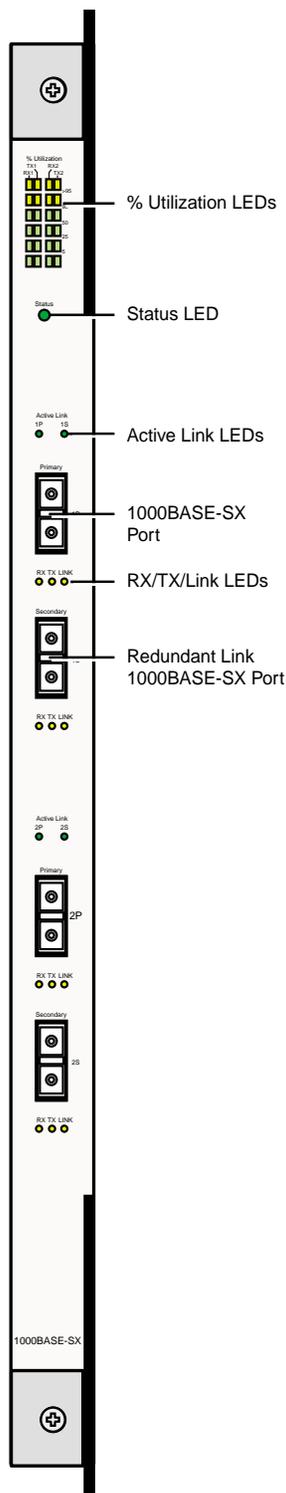
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port/2-Redundant Port 1000BASE-SX Gigabit Ethernet Interface Module

The 2-port redundant link 1000BASE-SX interface module provides two additional ports for redundancy. This module is hot-swappable and features the LED indicators listed in the table below.

The 1000BASE-SX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using short wavelength (800 nm) lasers. Maximum link distance is 550 meters over 50 μm multimode fiber and 275 meters using 62.5 μm multimode fiber.



LED Descriptions

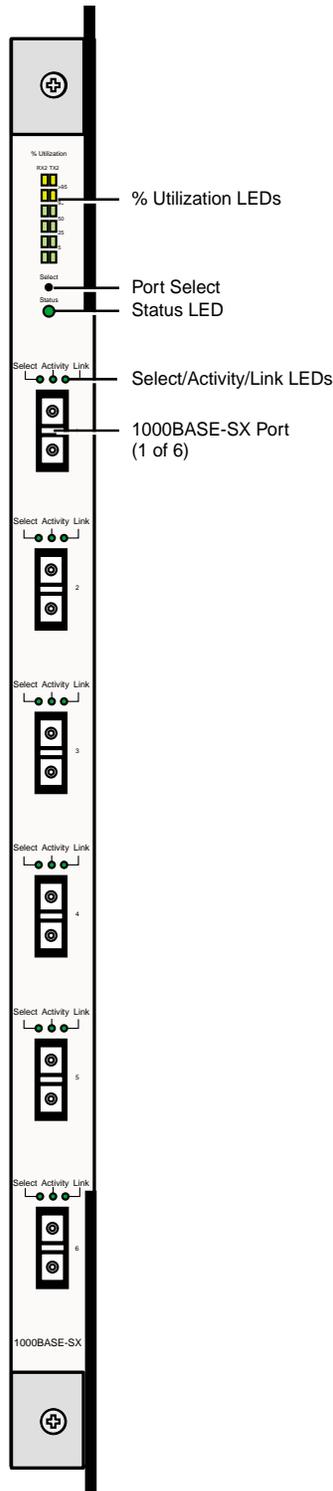
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

These Gigabit Ethernet links share the same gigabit MAC. The signal transmitted on redundant links is identical to the signal on the primary interface. The primary Gigabit Ethernet interface is continually checking itself for link integrity pulses. If the primary interface loses link integrity, the secondary interface takes over in less than 10 milliseconds. In this case the primary link can be serviced without affecting traffic on the secondary link.

6-Port 1000BASE-SX Gigabit Ethernet Interface Module

The 6-port 1000BASE-SX server interface module functions as a port aggregator, is hot-swappable, and features the LED indicators listed in the table below.

The 1000BASE-SX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using short wavelength (800 nm) lasers. Maximum link distance is 550 meters over 50 μm multimode fiber and 275 meters using 62.5 μm multimode fiber.



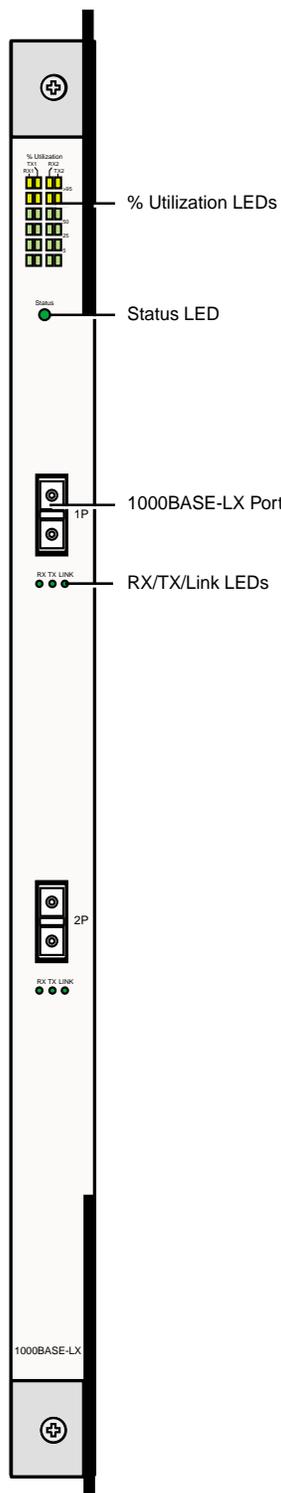
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port 1000BASE-LX Gigabit Ethernet Interface Module

The 2-port 1000BASE-LX interface module is hot-swappable and features the LED indicators listed in the table below.

The 1000BASE-LX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using long wavelength (1300 nm) lasers. Maximum link distance is 5 km over single-mode fiber.



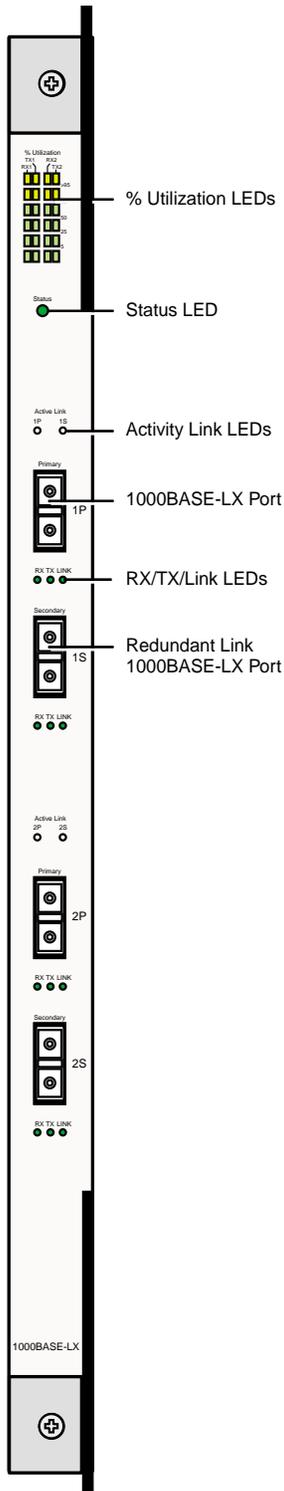
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port/2-Redundant Port 1000BASE-LX Gigabit Ethernet Interface Module

The 2-port redundant link 1000BASE-LX interface module provides two additional ports for redundancy. This module is hot-swappable and features the LED indicators listed in the table below.

The 1000BASE-LX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using long wavelength (1300 nm) lasers. Maximum link distance is 5 km over single-mode fiber.



LED Descriptions

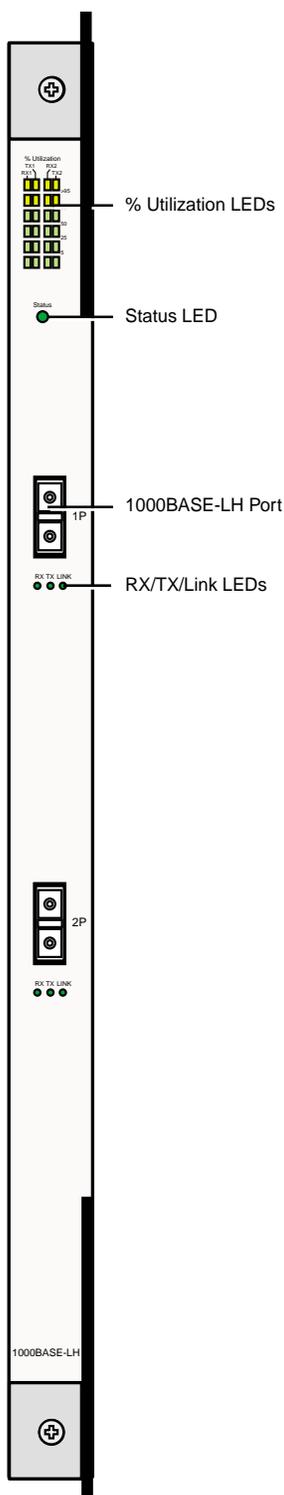
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Active Link	Green None	Port is operational on an interface. Port is not operational on an interface.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

These Gigabit Ethernet links share the same gigabit MAC. The signal transmitted on redundant links is identical to the signal on the primary interface. The primary Gigabit Ethernet interface is continually checking itself for link integrity pulses. If the primary interface loses link integrity, the secondary interface takes over in less than 10 milliseconds. In this case the primary link can be serviced without affecting traffic on the secondary link.

2-Port 1000BASE-LH Gigabit Ethernet Interface Module

The 2-port 1000BASE-LH (long haul) interface module is hot-swappable and features the LED indicators listed in the table below.

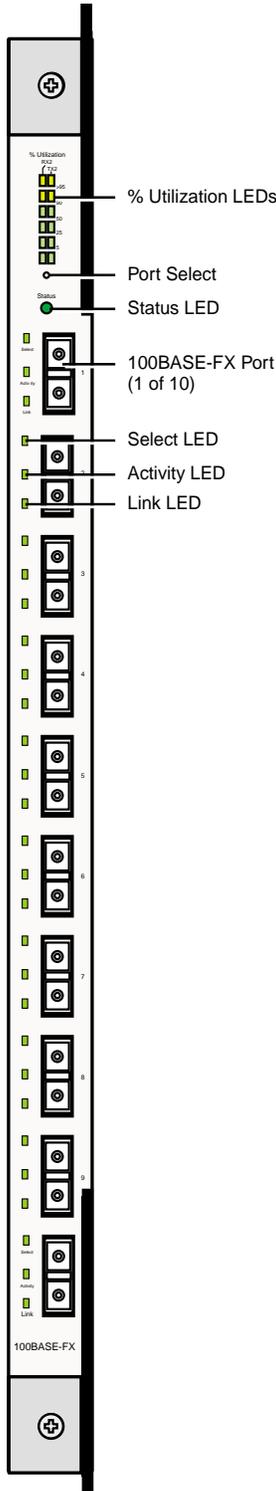
The 1000BASE-LH ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using long wavelength (1550 nm) lasers. Maximum link distance is 50 km over single-mode fiber.



LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

10-Port 100BASE-FX Fast Ethernet Interface Module



The 10-port 100BASE-FX interface module is hot-swappable and features the LED indicators listed in the table below.

The 100BASE-FX ports use Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the IEEE 802.3u Fast Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using long wavelength (1300 nm) lasers. Maximum link distance is 2 km over multimode fiber.

This interface module is also available in a single-mode fiber version, which provides a maximum link distance of 14 km.

LED Descriptions

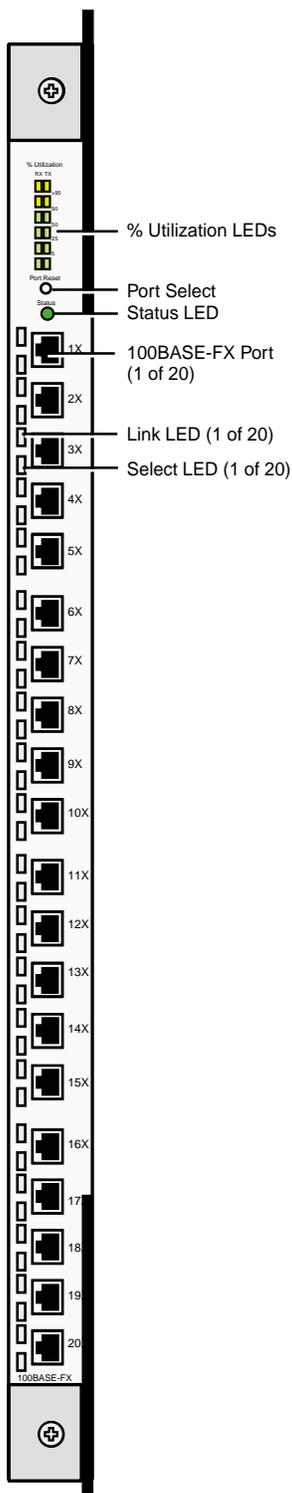
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
Activity	Green None	Port is receiving or transmitting data. Port is not receiving or transmitting data.
Select	Yellow None	Port activity is currently being displayed through the % Utilization LEDs. Port activity is not being displayed through the % Utilization LEDs.
Link	Green None	Port is operational on interface. Port is not operational on interface.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

The Gigabit Ethernet links share the same gigabit MAC. The signal transmitted on redundant links is identical to the signal on the primary interface. The primary Gigabit Ethernet interface is continually checking itself for link integrity pulses. If the primary interface loses link integrity, the secondary interface takes over in less than 10 milliseconds. In this case the primary link can be serviced without affecting traffic on the secondary link.

20-Port 100BASE-FX Fast Ethernet Interface Module

The 20-port 100BASE-FX interface module is hot-swappable and features the LED indicators listed in the table below. The 100BASE-FX ports use MT-RJ connectors and operate in full-duplex mode.

These ports are fully compliant with the IEEE 802.3u Fast Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. These ports transmit using long wavelength (1300 nm) lasers. Maximum link distance is 2 km over multimode fiber.



LED Descriptions

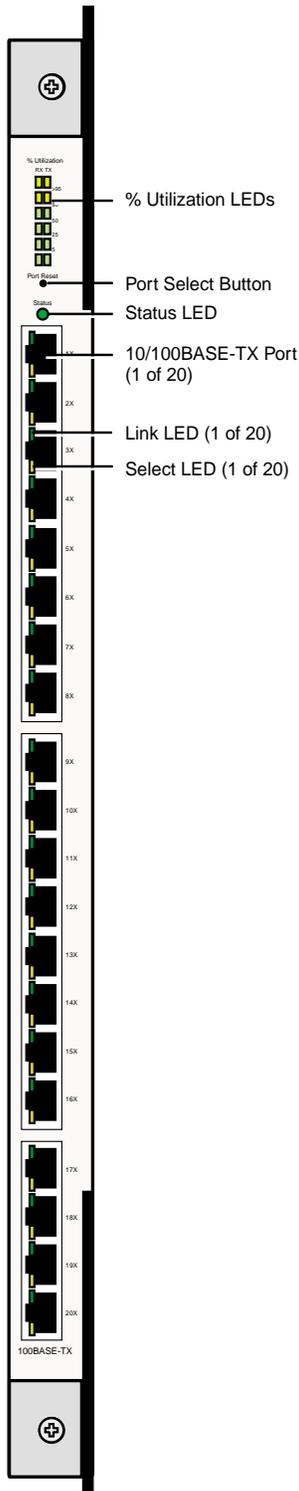
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
Select	Yellow None	Port activity is currently being displayed through the % Utilization LEDs. Port activity is not being displayed through the % Utilization LEDs.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

20-Port 10/100BASE-TX Fast Ethernet Interface Module

The 20-port 10/100BASE-TX interface module is hot-swappable and features the LED indicators listed in the table below.

The ports use RJ-45 connectors and support auto-negotiation. These ports are fully compliant with the IEEE 802.3 Ethernet, IEEE 802.3u Fast Ethernet, and IEEE 802.3x Full-Duplex with Flow Control standards.

Maximum link distance is 100 meters over category 5 (100BASE-TX) or category 3 (10BASE-T) unshielded twisted-pair cabling.



LED Descriptions

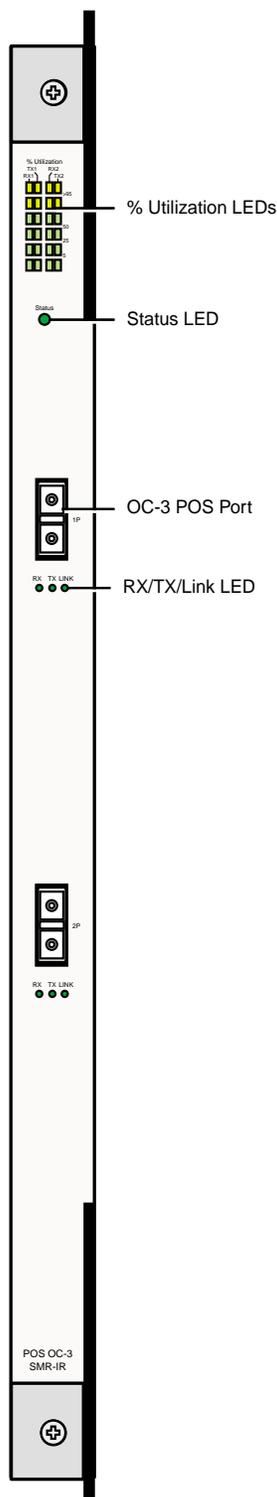
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity. Port is selected by pressing the Port Select button. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
Select	Yellow None	Port activity is currently being displayed on the % Utilization LEDs. Port activity is not being displayed on the % Utilization LEDs.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port OC-3c Packet-Over-SONET (POS) Interface Module

The 2-port OC-3c POS interface module is hot-swappable and features the LED indicators listed in the table below.

The OC-3c POS ports utilize Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the ITU-T G.709 (Synchronous Multiplexing Structure) and ANSI T1.105-1995 and ANSI T1.105.02-1995 SONET standards. RFC 1619 (PPP over SONET/SDH) and RFC 1662 (PPP in HDLC-like framing) are also fully supported. Each port provides a line rate of 155 Mbps.

This interface module is available in either a single-mode fiber-Intermediate Reach (SMF-IR) version, or a multimode fiber-Short Reach (MMF-SR) version. Maximum link distance is 15 km over single-mode fiber and 2 km over multimode fiber.



LED Descriptions

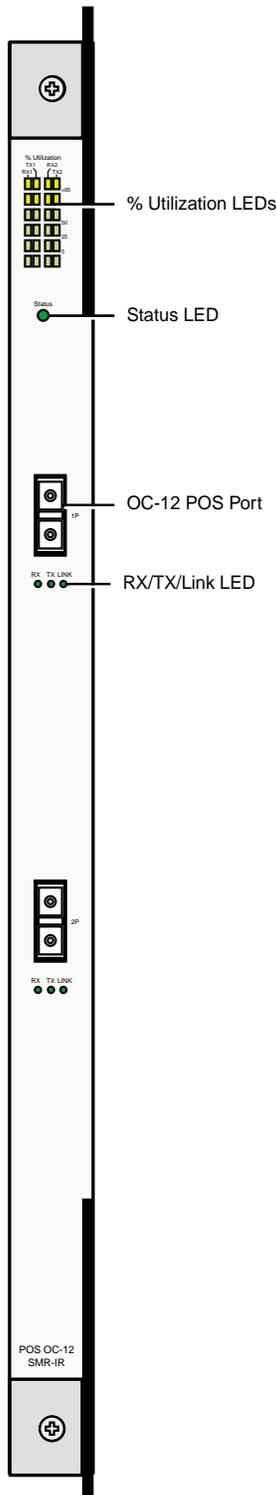
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port OC-12c Packet-Over-SONET (POS) Interface Module

The 2-port OC-12c POS interface module is hot-swappable and features the LED indicators listed in the table below.

The OC-12c POS ports utilize Duplex SC connectors and operate in full-duplex mode. These ports are fully compliant with the ITU-T G.709 (Synchronous Multiplexing Structure) and ANSI T1.105-1995 and ANSI T1.105.02-1995 SONET standards. RFC 1619 (PPP over SONET/SDH) and RFC 1662 (PPP in HDLC-like framing) are also fully supported. Each port provides a line rate of 622 Mbps.

This interface module is available in either a single-mode fiber-Intermediate Reach (SMF-IR) version, or a multimode fiber-Short Reach (MMF-SR) version. Maximum link distance is 15 km over single-mode fiber and 500 meters over multimode fiber.



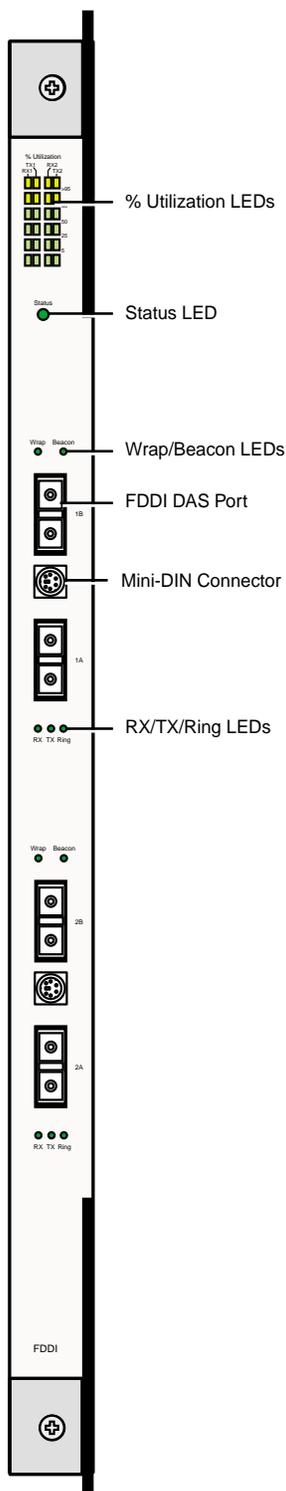
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

2-Port DAS/SAS FDDI Interface Module

The 2-port Dual Access Station (DAS) or Single Access Station (SAS) FDDI interface module is hot-swappable and features the LED indicators listed in the table below.

The DAS/SAS ports utilize Duplex SC connectors and support FDDI-to-Ethernet fragmentation on a port-by-port basis. In addition, two 6-pin mini-DIN connectors are supplied, one for each DAS connection. The mini-DIN connectors allow for the attachment of an external optical bypass switch to prevent routing switch or module failure from interrupting the FDDI network. These ports are fully compliant with the ANSI X3T9.5 standards for FDDI. Maximum link distance is 550 meters over multimode fiber.



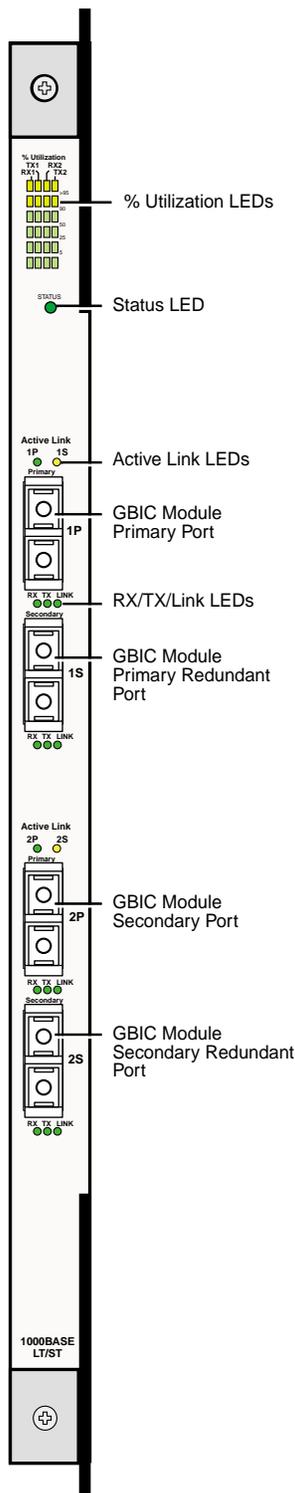
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
Wrap	Green None	Indicates a broken ring and the secondary ring is being used (DAS), or that the link is an SAS connection. Indicates that both DAS rings are operational.
Beacon	Green None	Ring is attempting to reestablish itself. Ring condition is normal.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Ring	Green None	Connection is valid and operational. Connection is not valid or operational.

2 GBIC-Port/2-Redundant GBIC-Port Gigabit Line Card

This interface module is hot-swappable and features the LED indicators listed in the table below. It serves as a replacement for the following Ethernet Gigabit interface modules:

- 2-Port 1000BASE-SX Gigabit Ethernet
- 2-Port Redundant Link 1000BASE-SX Gigabit Ethernet
- 2-Port 1000BASE-LX Gigabit Ethernet.
- 2-Port Redundant Link 1000BASE-LX Gigabit Ethernet
- 2-Port 1000BASE-LH Gigabit Ethernet



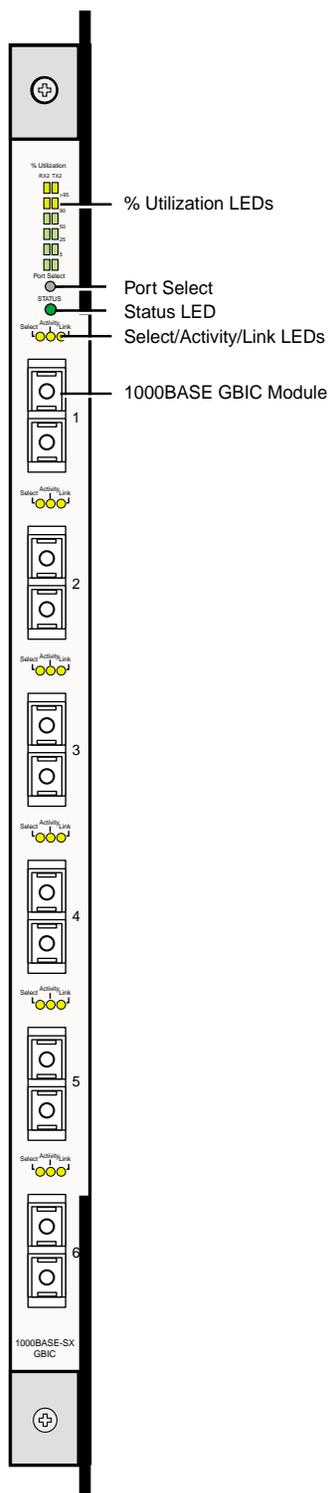
LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

GBIC modules are available in SX, LX, and LH configurations and are hot-swappable. See [Gigabit Interface Converter \(GBIC\) Modules](#) on page 1-17.

6 GBIC-Port 1000BASE Gigabit Server

This 1000BASE interface module functions as a port aggregator, is hot-swappable, and features the LED indicators listed in the table below. This interface board does not support redundant links.

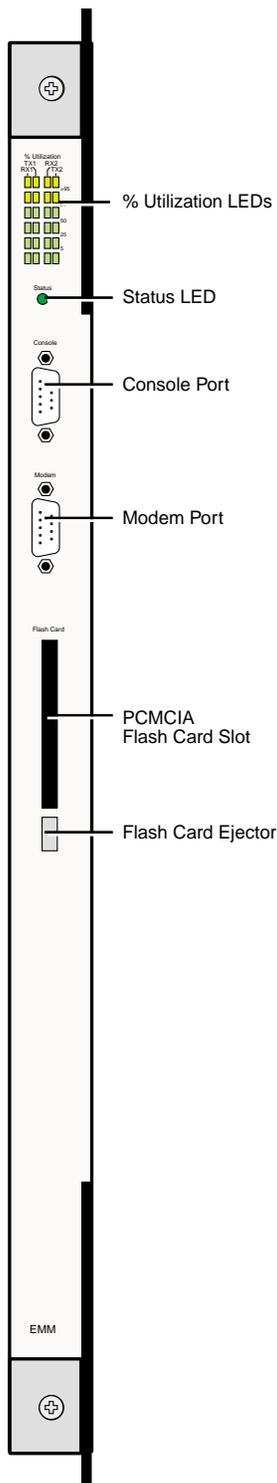


LED Descriptions

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None	Module is operational. Module is not operating properly. Module is not operational.
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

GBIC modules are available in SX, LX, and LH configurations and are hot-swappable. See [Gigabit Interface Converter \(GBIC\) Modules](#) on page 1-17.

Enterprise Management Module (EMM)



The Enterprise Management Module (EMM) is used to control all aspects of the OmniCore 5052 and 5022 routing switches. It will operate properly in either slot 1 or slot 2 (slots 1 and 2 are reserved for EMMs). One EMM card is required for operation, the other is installed for redundancy. All EMM interface modules are hot-swappable. The following table describes the Status LED on the EMM.

Each EMM features two serial ports: a Console port and a Modem port. The Console port is used for switch management and a terminal or PC can be attached to this port. Both serial ports require a standard null modem cable with a DB-9 connector (supplied). For specific serial port pin-out details, see [Appendix C, "Technical Specifications"](#). One PCMCIA flash card slot is also provided.

Status LED

LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None Flashing Green	Module is operational. Module is not operating properly. Module is not operational. Redundant card is ready (on standby).

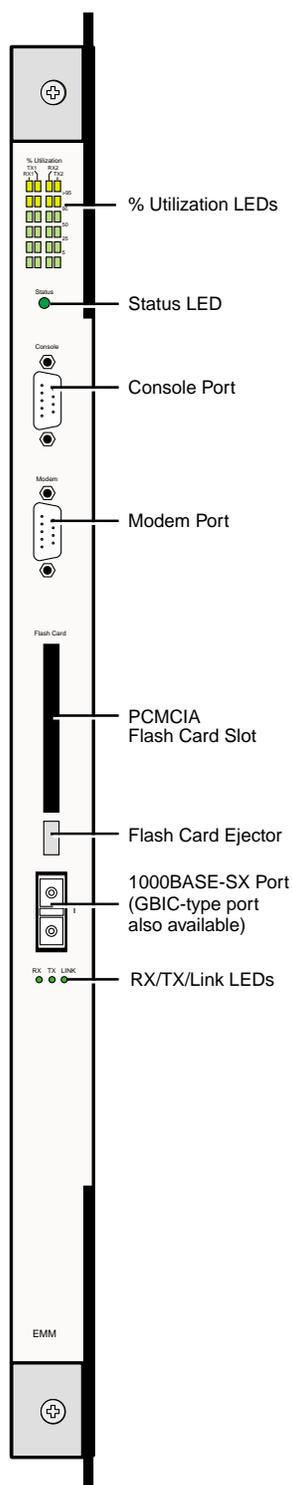
This EMM module is available with 64 MB or 256 MB of RAM.

EMM Modules

EMM Module	Part Number
EMM, 64 MB	OC-5000-1060
EMM, 256 MB	OC-5000-1119

Enterprise Management Module with Gigabit Ethernet Port

This EMM module is the same as described in the previous section except that it can be ordered with either a fixed 1000BASE-SX port or with a GBIC port.



EMM Modules

EMM Module	Part Number
EMM with 1000BASE-SX Gigabit Ethernet Port, 64 MB	OC-5052-1055
EMM with 1000BASE-SX Gigabit Ethernet Port, 256 MB	OC-5052-1121
EMM with GBIC Port, 256 MB	OC-5000-1095

LED Descriptions

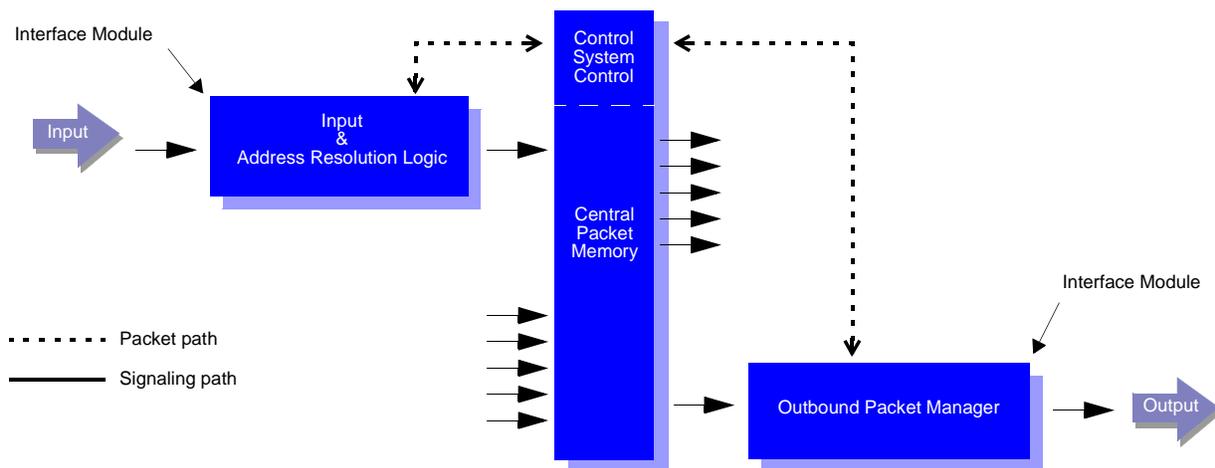
LED	Color	Description
% Utilization	Green/Yellow None	Current port activity is displayed. Port currently has no activity.
Status	Green Amber None Flashing Green	Module is operational. Module is not operating properly. Module is not operational. Redundant card is ready (on standby).
RX	Green None	Port is receiving data. Port is not receiving data.
TX	Green None	Port is transmitting data. Port is not transmitting data.
Link (Integrity)	Green None	Connection is valid and operational. Connection is not valid or operational.

The fixed 1000BASE-SX port uses a Duplex SC connector and operates in full-duplex mode. It is fully compliant with the IEEE 802.3z Gigabit Ethernet standard and the IEEE 802.3x Full-Duplex with flow control standard. The port transmits using a short wavelength (800 nm) laser. Maximum link distance is 550 meters over 50 μm multimode fiber and 275 meters using 62.5 μm multimode fiber.

GBIC modules are available in SX, LX, and LH configurations and are hot-swappable. See [Gigabit Interface Converter \(GBIC\) Modules](#) on page 1-17. However, the SX GBIC is shipped standard. The LX and LH are by special order only.

Architecture

The OmniCore routing switch's packet routing path is shown in the following diagram.



OmniCore Packet Routing Path

1. As a packet enters an interface module, Address Resolution Logic examines all Layer 2 information (addressing, frame type, VLAN, priority, etc.) and all Layer 3 information (addressing, protocol, TOS, etc.).
2. The switch evaluates Layer 2 address information to determine the type of traffic being sent to the routing switch (e.g., is it an SNMP frame, a frame that should be routed, a frame that should be switched, etc.)
3. The switch then evaluates the packet's Layer 3 information.
4. In parallel with examining the Layer 2 and Layer 3 information, the Address Resolution Logic examines Layer 4 attributes relative to any Layer 4 policies.
5. The Address Resolution Logic notifies the system control circuitry that a packet is coming.
6. The central system control reserves a memory location and signals the appropriate outbound ports.
7. With the Layer 2, Layer 3, and Layer 4 information in hand, the Address Resolution Logic performs one or more parallel functions:
 - Passes the traffic on to the central memory
 - Associates traffic with the appropriate priority and the appropriate egress ports
 - Performs the appropriate accounting functions
 - Redirects the traffic to a different (overridden) destination
 - Drops the traffic according to a security policy
8. Once the traffic has been passed to the central memory, the system control circuitry signals the appropriate outbound ports to indicate that they now control the packets flow.
9. The outbound port pulls the data from the central location when it needs to. The destination port can pull the packet from the central memory using any of three different queuing methods: WFQ, Strict, and PBQ. WFQ (Weighted Fair Queuing) and Strict priority modes provide "canned" settings for priority queue bandwidths, burst enabling, and burst

bandwidth for the port while Weighted Fair Queuing (WFQ) drains queues based on a sophisticated smoothing algorithm that gives high priority queues more bandwidth and lower priority queues less bandwidth. This entire process happens extremely fast and in a deterministic manner. The routing latency within the switch for gigabit speed links is less than 10 microseconds.

Operation

The OmniCore routing switch can operate as a Layer 2 switch or a combination of switch and router (i.e., routing switch). The system is shipped for operation as a switch with router operation requiring ports to be configured for routing. The OmniCore routing switch provides routing support for IPv4, IPX, IPv6, and additional protocols.

Switching

The OmniCore routing switch is a store-and-forward device. This provides granular packet-by-packet switching (and routing) performance. When devices are connected, the OmniCore routing switch immediately begins learning their Layer 2 addresses. Since it can act as a switch-only device, many of the same requirements for Layer 2 switches also apply to the OmniCore routing switch. One of the most important of these requirements is support for Spanning Tree (802.1d), which prevents looped network topologies. The system also supports IEEE 802.1Q tags and IEEE 802.1p priority in this mode of operation.

IP Routing

The OmniCore routing switch routes a packet when the packet's destination MAC address is equal to the MAC address of the inbound port. The OmniCore routing switch performs the following functions when it routes IP packets:

- Decrements the time to live (TTL) value.
- Verifies that the TTL is greater than zero. If it is, then the packet is routed.
- Substitutes the MAC address.
- Updates the checksum of the packet to reflect any MAC address substitutions for the next hop router or end stations.

The OmniCore routing switch uses routing update information received and sent using IP routing protocols to generate routing tables that are stored on an as needed basis in every port. A master routing table is stored in the EMM. This central forwarding information is then distributed to the appropriate ports. As packets that need to be routed arrive, they are forwarded at the ingress interface using the packet's Layer 3 information.

The OmniCore routing switch uses a longest prefix match algorithm to determine the best route for each packet. Longest prefix match uses the longest route match in its routing tables when making a routing decision. This provides very deterministic route selection with minimal latency.

The OmniCore routing switch supports routing by class or by Classless Interdomain Routing (CIDR) address. The use of CIDR addressing enables the deployment of variable length subnet masks (VLSMs) and supernets. The system also supports multinet, which allow multiple IP networks to exist on one interface.

The OmniCore routing switch provide hardware support for the IP TOS precedence field, by using its information for priority queuing and route calculation.

IPX Routing

The OmniCore routing switch is capable of routing IPX packets at wire speed. To do so, the switch exchanges IPX RIP/SAP updates with other routers. These updates are processed by the OmniCore routing switch's EMM and then distributed to the interface modules, as they are required.

When an IPX packet arrives, the switch looks at its IPX network number and also determines the frame encapsulation type. The OmniCore routing switch then uses the RIP and SAP information it has exchanged with other IPX routers to determine the best route for a packet. The switch will perform wire-speed, hardware-based routing if the IPX network number is a known unique IPX network number. If there are multiple IPX networks attached to a port, or if a single IPX network spans multiple ports, then the switch routes IPX packets via the EMM.

Other Routing Protocols

The OmniCore routing switch has the ability to route up to eight additional protocols (including AppleTalk and DECnet) via hardware assist. When the system receives a packet using these protocols, it is immediately identified as not being IPv4 or IPX, and it is forwarded to the EMM. The EMM then applies software-based routing techniques to route each packet.

Multicast Support

The OmniCore routing switch architecture lends itself to wire-speed multicast performance. This is directly related to the system's *Parallel Access* Shared Memory, which provides each port with a full-duplex path into and out of the switch's central memory fabric. Because each port has a full-duplex path into and out of system memory, the system overcomes the problems associated with packet replication.

The system supports routed multicast protocols on ports configured for routing and supports IGMP snooping on ports configured for switching to provide a complete multicast solution. The OmniCore routing switch provides support for multicasting using the following protocols.

Layer 2 Multicast

- Internet Group Management Protocol (IGMP) "Snooping"

Layer 3 Multicast

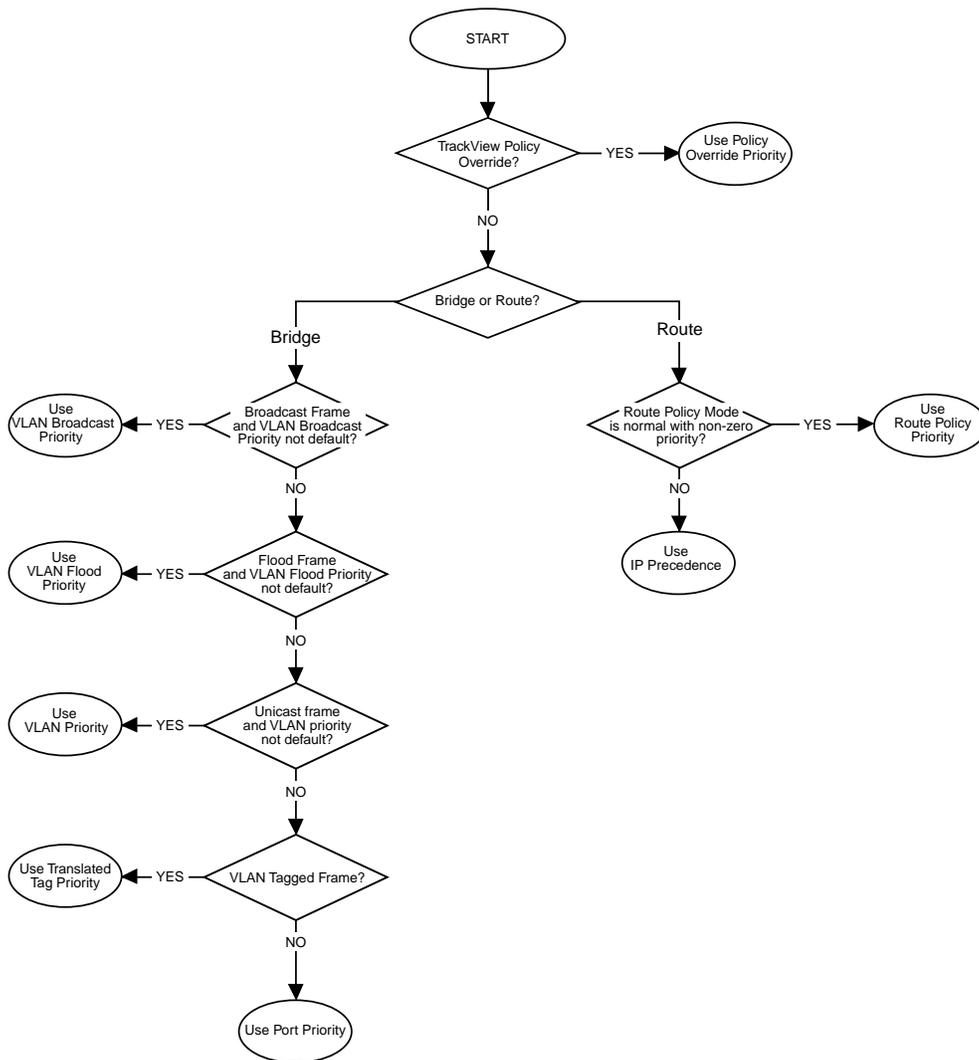
- Internet Group Management Protocol (IGMP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Protocol-Independent Multicast Sparse-Mode (PIM-SM)

OmniCore Frame Prioritization

The OmniCore routing switch implements a complex set of frame prioritization rules. A frame's priority can come from a number of places:

- Every ingress port has a default priority.
- A VLAN tagged frame contains an indicated priority.
- VLANs have configurable flood, broadcast and normal traffic priorities.
- The IP Precedence field may be used.
- An IP route may be given a policy.
- TrackView may configure a classification-based policy.

These sources are organized in a hierarchy, as shown in the following diagram, with most-specific classifications taking preference over least-specific classifications. For example, a priority indicated from TrackView policy manager (a very traffic-specific source) has preference over the priority indicated by the ingress port (a less traffic-specific source)



Frame Priority Determination Flow Chart

Port Priority

A received frame is given an initial priority classification based on the priority configured at the receive port. A port's priority may be configured using the priority command within the slot/port menu, as follows:

```
OmniCore> slot 3 port 1 priority 5
```

Following this command, all frames received on slot 3, port 1 will be given priority 5 unless acted upon by a higher preference source.

VLAN Tagged Frame Priority

Frames received on a tagged port contain a priority indication in the VLAN tag. The IEEE802.1Q standard defines a translation between this tag encoded priority and the effective internal traffic class (priority). This translation is shown in the following table.

Ingress Tag Priority	Internal Priority (Egress Queue)	Egress Tag Priority
7	7	7 (Highest Priority)
6	6	6
5	5	5
4	4	4
3	3	3
0	2	0 (Default)
2	1	2
1	0	1 (Lowest Priority)

IEEE 802.1p
CLI and SNMP values for
VLAN, flood, broadcast, and port.
IEEE 802.1p

Internal and IEEE 802.1p Priorities

The consideration behind the non-linear translation is that it is expected that normal best-efforts traffic will carry a priority indication of zero. Rather than treat this normal traffic as lowest priority, the translation allows for two lesser classes of traffic below it. The idea is that normal traffic will have a Tag priority indication of zero, and that normal, best efforts, traffic should not be considered the least important traffic. The translation allows for two levels of traffic of a lower class than "best efforts" to be defined.

Ingress Translation

When a frame is received on a VLAN tagged port, the priority as indicated in the VLAN tag is translated according to the preceding table. Thus a frame with an indicated priority of zero receives an internal priority of 2, unless superseded by a higher-preference rule.

Egress Translation

The translation is also performed on frames that egress a VLAN tagged port. A route-policy, for example, that raises a frame's priority to 2, will egress a tagged port with an indicated priority of 0, while a non-prioritized frame will egress with an indicated tag priority of 2.

◆ Note ◆

The management view of priorities is always linear 0 to 7, and represents a frame's internal priority. The internal priority represents the egress queue for the frame. Tag priority translation is only performed at frame ingress/egress on tagged ports.

VLAN Priorities

Three classes of traffic may be individually prioritized during bridging operations:

- Flooded traffic
- Broadcast traffic
- Unicast traffic

Flooded frames are bridged frames whose MAC Destination Address are not found in the Bridge Filtering Database. Flood frame priority is set using the *flood-priority* VLAN parameter.

Broadcast frames are bridged frames whose MAC Destination Address is the Ethernet broadcast address FF:FF:FF:FF:FF:FF. Broadcast frame priority is set using the *bcast-priority* VLAN parameter.

Unicast frames are bridged frames whose MAC Destination Address is found the Bridge Filtering Database. This is normal traffic. Unicast frame priority is set using the *priority* VLAN parameter.

VLAN priorities may only be set when a VLAN is created or when it is disabled. VLAN priorities take the values, *default* and *0* to *7*. The *default* value indicates that no explicit priority is required for a particular class of VLAN traffic. This enables the lesser preference rules (port priority and VLAN tag priority) to be active.

An example of how to create a VLAN 2 with flood priority 0, broadcast priority 1 and normal priority default is as follows:

```
OmniCore> vlan 2 tag 2 flood-priority 0 bcast-priority 1 priority default create
```

◆ Note ◆

VLAN priorities are only effective for bridged traffic.

IP Precedence

The IP header contains a 3-bit precedence field. For IP routed traffic, this field is used to determine the priority. The IP precedence value is a linear 0 to 7, with 0 being lowest priority and 7 being highest.

IP Route Policy

The *ip route-policy* command may be used to bind specific QoS actions to routes. The *qos-mode* parameter of a route policy can select between using the precedence field in the IP header and a priority explicitly configured in the policy itself.

The following command creates a route-policy for 10.0.0.0/8:

```
OmniCore> ip route-policy 10.0.0.0/8 0 create
```

The policy is initially in normal QoS mode which, unless an explicit priority is configured, means that the IP precedence field in the packet will be used to select a frame priority.

An explicit priority can be set with the priority command:

```
OmniCore> ip route-policy 10.0.0.0/8 0 priority 5 create
```

Route policy commands also support IP Precedence re-stamping, which is covered in the command documentation.

TrackView Policies

The Policy Manager within TrackView can be used to apply a priority to frames matching a configured classification rule. These policies implement the highest preference prioritization rules.