

Cisco Catalyst 6500 Supervisor Engine 32 Architecture

Executive Overview

The Cisco® Catalyst® 6500 Supervisor Engine 32 is the latest addition to the Cisco Catalyst 6500 Family of supervisor modules. Designed primarily for access layer deployments, the Supervisor Engine 32 provides Layer 2 bridging and Layer 2 through 4 services with Layer 3 routing optional hardware-accelerated services. It provides connectivity into the classic 32-Gbps switching bus and provides Layer 2 and 3 switching performance up to 15 Mbps.

This new supervisor contains technology elements found in earlier supervisor models as well as a number of new technology enhancements. The Supervisor Engine 32 includes the following features:

- Two Supervisor Engine 32 models providing either of the following
 - 8 Gigabit Ethernet Small Form-Factor Pluggables (SFPs) and 1 10/100/1000TX uplink ports
 - 2 10-Gigabit Ethernet uplink ports and 1 10/100/1000 TX port
- Integrated policy feature card (PFC3B)
- Integrated multilayer switch feature card (MSFC2a)
- Classic 32-Gbps bus connection
- Two Universal Serial Bus (USB) ports on front panel
- Compact Flash slot
- Console port
- 256-MB bootflash
- 256-MB DRAM upgradable to 1 GB

The most obvious difference from earlier supervisor modules is the presence of eight SFP-based Gigabit Ethernet ports on the front of one of the Supervisor Engine 32 modules. These front ports can help reduce the need to use a chassis slot for a Gigabit Ethernet module where only a few Gigabit Ethernet ports are required. The other Supervisor Engine 32 model sees the introduction of 10-Gigabit Ethernet as an uplink technology for the first time.

Two USB ports are also on the front panel, one designated for host use (Type "A" USB port) and the other as a device port (Type "B" USB port). Host USB ports can be used to plug in devices such as a PC, while device ports can be used for attaching devices such as a Flash memory key. These ports will be enabled in a post First Customer Ship release of code.

The integrated PFC3B is included on the Supervisor Engine 32 to enhance its quality-of-service (QoS) and security capabilities. With the PFC3B, the Supervisor Engine 32 can support hardware-based QoS and security access control lists (ACLs) using Layer 2, 3, and 4 classification criteria to secure and prioritize target data. Standard PFC3B enhancements can also be utilized, allowing the Supervisor Engine 32 to take advantage of new hardware accelerated features such as CPU rate limiters, ACL hit counters, port access control lists (PACLs), and improvements in route and NetFlow capacities, to name a few.

The MSFC2a is also integrated into the Supervisor Engine 32 to enable it to be a full-fledged Layer 3 switch. At FCS, the Supervisor Engine 32 will act purely as a Layer 2 device. A Cisco IOS[®] Software option to enable Layer 3 services on the supervisor will become available at a future date. The forwarding architecture used by the MSFC2a and PFC3 is Cisco Express Forwarding, which is the Cisco Systems[®] architecture for providing Layer 3 switching in hardware. In the future, the option of adding a programmable intelligent services adapter (PISA) daughter card, which integrates the functionality of the MSFC2a into the card, will be provided along with hardware accelerate support for advanced Layer 4 to 7 services such as hardware-accelerated network-based application recognition (NBAR).

This document will explore the capabilities of the new Supervisor Engine 32. It will provide an insight into the hardware architecture of the new supervisor, the features it supports, the packet flow through the Supervisor Engine 32, and a summary of its performance characteristics

Hardware Architecture

There are numerous aspects to the architecture of the Supervisor Engine 32, which are explored in more detail in the following sections.

Supervisor Baseboard Features

The Supervisor Engine 32 baseboard is visually different from that of other Cisco Catalyst 6500 supervisor options. The initial Supervisor Engine 32 that was shipped provided for eight SFP-based Gigabit Ethernet ports and a single 10/100/1000TX RJ45 Gigabit Ethernet port. All nine ports are capable of the same bridging, routing, QoS, ACLs, and so on as separate Ethernet modules, and these ports can all be used at the same time. (See Figure 1.)

Figure 1. Supervisor Engine 32 with Eight SFP Gigabit Ethernet Ports



The Supervisor Engine 32-10GE is the second Supervisor Engine 32 option now shipping. It provides two front 10-Gigabit Ethernet ports (using XENPAK optics). Like its counterpart, both 10-Gigabit Ethernet ports can be used at the same time in addition to the 10/100/100TX port, which is available on the front panel. (See Figure 2.)

Figure 2. Supervisor Engine 32 with Two 10-Gigabit Ethernet Ports



Both Supervisor Engine 32 models incorporate a Compact Flash type II slot, located to the left on the front panel of the module. A console port is also available on the front panel should console access be required. At the far right of the front panel are two USB ports. These ports are disabled, but will be enabled in a future software release. The USB ports differ from an operational perspective in that one is designated as a host port and the other as a device port. (See Figure 3.)

Figure 3. Supervisor Engine 32 Front Panel



The Supervisor Engine 32 uses internal Compact Flash, which replaces the bootflash used in other supervisor modules. This internal Compact Flash is physically the same type of Compact Flash used externally. The default amount of bootflash available in the Supervisor Engine 32 is 256 MB. The internal Compact Flash can only be accessed when the module is removed from the chassis. This Compact Flash is able to store images that can be used to boot the system. The terminology used to access this Compact Flash from the command-line interface (CLI) differs slightly from what was used to access bootflash. This Compact Flash is referred to as "BOOTDISK," and changes to the CLI have been made to allow this to be referenced.

The supervisor can take advantage of the PFC3B, enabling a range of advanced hardware that can accelerate certain features. Some of these features are available at First Customer Ship, while others will be enabled in a future software release. Full Layer 3 functionality is now available with either a hybrid software implementation (Cisco Catalyst OS plus MSFC Cisco IOS Software) or native Cisco IOS Software. Both the PFC and MSFC options and the features they provide are discussed later in this paper.

The specifications of the Supervisor Engine 32 baseboard have also been upgraded when compared to earlier supervisor options. Table 1 compares the baseboard components available with the Supervisor Engine 32 to those in earlier wiring closet supervisors.

Table 1. Supervisor Engine 32 Baseboard Component Comparison

Feature	Cisco Catalyst 6500 Series Supervisor Engine 1A	Cisco Catalyst 6500 Series Supervisor Engine 2	Supervisor Engine 32
Backplane	32-Gbps bus	32-Gbps bus w/256-Gbps switch fabric module	32-Gbps bus
SP DRAM	128-MB default/max	128-MB default *recently updated to 256-MB default	256-MB default
SP DAM Upgrade Options	–	256 MB/512 MB	1 GB
SP NVRAM	512 KB	512 KB	2 MB
SP Bootflash	16 MB	32 MB	256-MB default
Removable Storage	PCMCIA	PCMCIA	Compact Flash
USB Ports	No	No	Yes
Uplink Ports	2 GBIC	2 GBIC	8 SFP Gigabit Ethernet + 1 10/100/1000TX or 2 10-Gigabit Ethernet + 1 10/100/1000TX

It is worthwhile noting that although the Supervisor Engine 32 supports a maximum of 1 GB of DRAM, this really doesn't add value over the default 512MB DRAM while the PFC3B is only limited to supporting up to 256K routes in the hardware forwarding tables.

Chassis Options

The Supervisor Engine 32 supports all existing Cisco Catalyst 6500 chassis as well as the new “E” series chassis. It is not supported in the Cisco Catalyst 6006 or 6009 Switch chassis. Like other Catalyst 6500 supervisors, it must be placed into a specific slot. That slot is dependent on the chassis model. Table 2 highlights which chassis slots support the Supervisor Engine 32.

Table 2. Supervisor Engine 32 Slot Dependencies

Slot	Cisco Catalyst 6503 Switch	Cisco Catalyst 6503-E Switch	Cisco Catalyst 6504-E Switch	Cisco Catalyst 6506 Switch	Cisco Catalyst 6506-E Switch	Cisco Catalyst 6509 Switch	Cisco Catalyst 6509-E Switch	Cisco Catalyst 6509-NEB Switch	Cisco Catalyst 6509-NEB-A Switch	Cisco Catalyst 6513 Switch
1	Yes	Yes	Yes	No	No	No	No	No	No	No
2	Yes	Yes	Yes	No	No	No	No	No	No	No
3	No	No	No	No	No	No	No	No	No	No
4			No	No	No	No	No	No	No	No
5				Yes	Yes	Yes	Yes	Yes	Yes	No
6				Yes	Yes	Yes	Yes	Yes	Yes	No
7						No	No	No	No	Yes
8						No	No	No	No	Yes
9						No	No	No	No	No
10										No
11										No
12										No
13										No

Supervisor Engine 32 Fan and Power Supply Requirements

With all chassis options, the presence of a Supervisor Engine 32 will require a new high-speed fan. For the “E” series chassis, this is not an issue as they all ship with a default high-speed fan. With the exception of the Cisco Catalyst 6509-NEB-A, the current (that is, non “E” series) chassis typically have two fan options, of which only one is compatible with the Supervisor Engine 32. The high-speed fan option needs to be selected for all chassis using the Supervisor Engine 32. The original fan and high-speed fan for the non “E” series chassis are both shown in Figure 4.

Figure 4. Catalyst 6500 Fans



For operational purposes, the fans in the right column in Table 3 must be used with the Supervisor Engine 32.

Table 3. Fan Options for Cisco Catalyst 6500 Chassis

Chassis	Original Fan (FAN1)	High-Speed Fan (FAN2)
Cisco Catalyst 6503	FAN-MOD-3	Fan-MOD-3HS
Cisco Catalyst 6503-E	–	WS-C6503-E-FAN
Cisco Catalyst 6504-E	–	WS-C6504-E-FAN
Cisco Catalyst 6506	WS-C6K-6SLOT-FAN	WS-C6K-6SLOT-FAN2
Cisco Catalyst 6506-E	–	WS-C6506-E-FAN
Cisco Catalyst 6509	WS-C6K-9SLOT-FAN	WS-C6K-9SLOT-FAN2
Cisco Catalyst 6509-E	–	WS-C6509-E-FAN
Cisco Catalyst 6509-NEBS	WS-C6509-NEB-FAN	WS-C6509-NEB-FAN2
Cisco Catalyst 6509-NEBS-A	–	FAN-MOD-09
Cisco Catalyst 6513	WS-C6K-13SLOT-FAN	WS-C6K-13SLOT-FAN2

In addition to the high-speed fan, the Supervisor Engine 32 also requires a minimum of a 2500W power supply or higher (AC or DC) to drive the new supervisor. If either a low-speed fan or a power supply less than 2500W is used, then a warning message will be displayed on bootup, and the switch will be shut down. For countries using 110V, both the 2500W and 3000W power supplies can still be used, but the power supply will run at approximately 50 percent capacity. To run those supplies at full capacity requires a 220V circuit.

The actual output power requirements of the baseboard and the individual components are listed in Table 4.

Table 4. Power Requirements

Component	Supervisor Engine 32-8GE	Supervisor Engine 32-10GE	PFC3B	MSFC2a
Power Requirement	1.89A	2.39A	1.47A	0.33A

The power requirements of the modules when combined together are detailed in Table 5 (figures @110V):

Table 5. AC Power Requirements of Supervisor Engine 32 with Components

Module	Output Current	Output Power	Heat Dissipation: AC (BTU/Hr)
WS-Sup32-GE-3B	3.69A	154.98W	661.57
WS-Sup32-10GE	4.19A	175.98W	751.21

In a DC environment, the power values change slightly and are detailed in Table 6 (at—48VDC).

Table 6. DC Power Requirements of Supervisor Engine 32 with Components

Module	Output Current	Output Power	Heat Dissipation: DC (BTU/Hr)
WS-Sup32-GE-3B	3.69A	154.98W	715.21
WS-Sup32-10GE	4.19A	175.98W	812.12

Supervisor Engine 32 and Line Cards

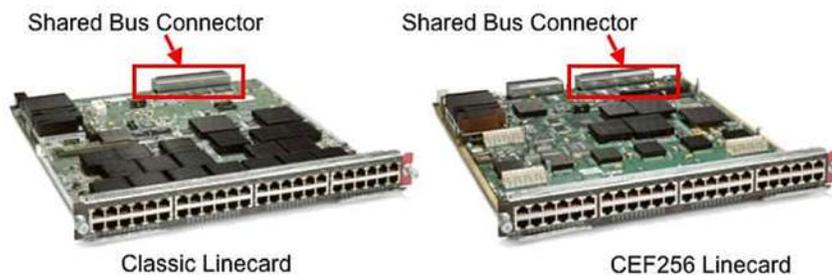
The Supervisor Engine 32 is a “classic” module, meaning it provides a connection to the “classic” 32-Gb bus to communicate with other line cards present in the chassis. Unlike some of the other supervisors, it has no built-in switch fabric, nor can it take advantage of a separate switch fabric module. This mode of operation thus defines the type of line cards that can work with this supervisor. Any line card that does not support data transfer over the classic bus cannot interoperate with the Supervisor Engine 32. A full list of the line card architectures supported with the Supervisor Engine 32 is in Table 7.

Table 7. Line-Card Architecture Compatibility with Supervisor Engine 32

Line-Card Architecture	Supported with Supervisor Engine 32
Classic	Yes
CEF256 (without DFC)	Yes
CEF256 (with DFC)	No
DCEF256 (WS-X6816)	No
CEF720 (67xx series)	No
Switch Fabric Module 1/2	No
Services modules	Yes (some exceptions)
Any DFC	No
OSM	Yes
FlexWAN	Yes
SIP/SPA	Yes (some exceptions)

As Table 7 shows, the Supervisor Engine 32 supports both CEF256 and classic line-card architectures. On both of these line cards there is a connector at the back of the line card that provides connectivity into the classic 32-Gb bus. This connector can be seen in Figure 5.

Figure 5. Connectors for Classic and CEF256 Line Cards



Supervisor Engine 32 Baseboard Architecture

The architecture of the Supervisor Engine 32 is similar to that of other Catalyst 6500 supervisors; however, it does have some unique aspects. The architecture of the Supervisor Engine 32 is shown in Figures 6 and 7.

Figure 6. Supervisor Engine 32-8GE Baseboard Architecture

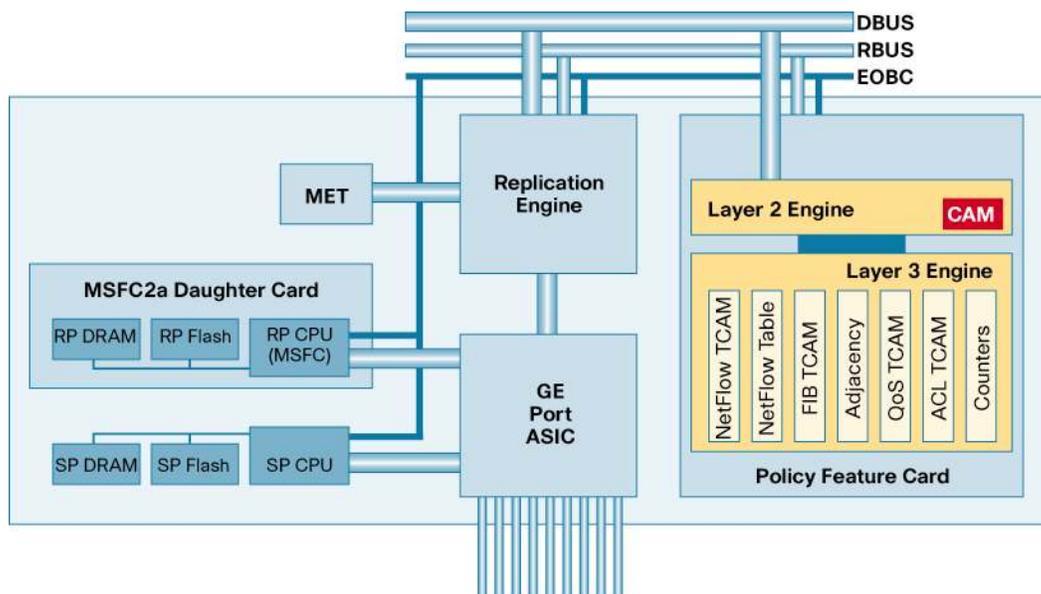
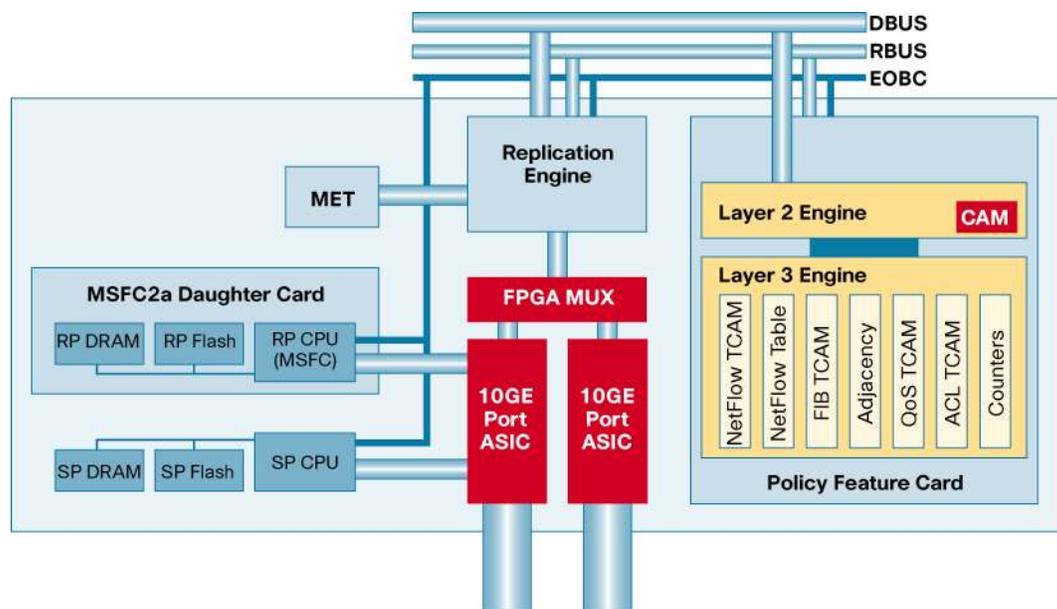


Figure 7. Supervisor Engine 32-10GE Baseboard Architecture



The Supervisor Engine 32 baseboard incorporates a default set of onboard application-specific integrated circuits (ASICs) for providing the foundation for Layer 2 through 4 services and interfacing into the 32-Gbps switching backplane. A single ASIC is used to connect the supervisor into the classic bus. This ASIC is also used for multicast replication and the switched port analyzer (SPAN) functionality. As can be seen in the diagram, this ASIC also interfaces into the multicast expansion table (MET), which provides the switch with an understanding of multicast group membership.

A switch processor (SP CPU) is used to perform all Layer 2 control plane processes, such as Cisco Discovery Protocol, Spanning Tree Protocol, and VLAN Trunking Protocol (VTP). The SP CPU has its own set of DRAM (256 MB), which is upgradable to 1 GB of bootflash and nonvolatile RAM (NVRAM is also dedicated to the SP CPU). Internal Compact Flash (256 MB) is used for the bootflash, and the SP CPU has access to 2 MB of NVRAM.

An onboard port ASIC is used to drive the front nine Gigabit Ethernet ports or the two 10-Gigabit Ethernet ports. All of the Gigabit Ethernet ports use SFP gigabit interface converters (GBICs). Different SFP options are available depending on the distance requirements. In redundant mode, all ports on both the primary and redundant supervisor are active. In a fully redundant chassis with two Supervisor Engine 32 modules, a total of 18 active Gigabit Ethernet ports will be available for use.

Integrated support for the PFC3B brings a range of advanced hardware accelerated features to the Supervisor Engine 32 and places the functionality of this supervisor well ahead of its wiring closet predecessors. The PFC3B provides a host of hardware-enabled features, allowing the Supervisor Engine 32 to functionally interoperate with the Cisco Catalyst 6500 Series Supervisor Engine 720. Features such as PACLs, ACL hit counters, CPU rate limiters, QoS and security ACLs, and more are all built into the hardware.

The shared 32-Gbps bus allows all connected ports (both the supervisor and line cards) to both transmit and receive data. The switching bus is actually composed of three discrete buses, each serving a specific function in the switching operation: the data bus (DBus), the results bus (RBus), and the Ethernet out-of-band control bus (EOBC).

The DBus is the main bus over which all data is transmitted. This bus is 256 bits wide and is clocked at 62.5 Mhz, which yields bandwidth of 16 Gbps. The common practice in the industry is to state performance in full duplex numbers; as this bus can send 16 G of data and receive 16 G of data, Cisco refers to the bus as a 32-Gbps bus. The RBus is the bus that is used by the supervisor engine to forward the result of the forwarding operation to each of the attached line cards. This bus also operates at 62.5 Mhz but is only 64 bits wide but is never the performance bottleneck. Finally, the EOBC is the control bus that relays control information between the line cards and the switch processor CPU.

The Supervisor Engine 32 supports the integrated MSFC2a. This MSFC option is functionally equivalent to the MSFC2 found on the Supervisor Engine 2. The only exception is that the MSFC2a supports up to 1 GB DRAM compared to 512 MB on the Supervisor Engine 2 MSFC2 (WS-F6K-MSFC2).

The Route Processor (RP) CPU has 64 MB of bootflash available to it along with 2 MB of NVRAM. A full duplex, 1-Gbps in-band connection allows the MSFC2a to communicate with other components on the Supervisor Engine 32 baseboard.

In the future, the Supervisor Engine 32 will support an optional PISA. This modular daughter card will include the functionality of the MSFC2a as well as incorporating a programmable network processor, which will allow the Supervisor Engine 32 to provide advanced Layer 4 to 7 services in hardware.

The initial target of PISA will be to provide hardware-based NBAR at speeds up around 1 to 2 Gbps. The network processor architecture found on the PISA is flexible enough for Cisco to provide other Layer 4 to 7 services in the future.

Supervisor Engine 32 Feature Review

The Supervisor Engine 32 provides a host of features that can be categorized into the following sections.

Supervisor Engine 32 with PFC3B Layer 2 Features

There are a number of Layer 2 features in the PFC3B that differentiate the Supervisor Engine 32 from earlier supervisor models. From a capacity perspective, the Supervisor Engine 32 is similar to the Supervisor Engine 2 in terms of its support for ACLs and MAC addresses. It does differ in that it provides some of the new features previously only found in the Supervisor Engine 720. A comparison of the Supervisor Engine 32 Layer 2 features with the Supervisor Engine 1 and Supervisor Engine 2 is summarized in Table 8. Some features will be enabled at a later date via new software.

Table 8. Wiring Closet Supervisor Layer 2 Feature Comparison

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
Layer 2 MAC Addresses	128 K (32 K effective)	128 K (32 K effective)	64 K (32 K effective)
Security VACLs	16 K VACLs	32 K VACLs	32 K VACLs
Security PACLs	No	No	32 K PACLs
ACE Hit Counters	No	No	Yes
SP Rate Limiters	No	No	Yes
SPAN	Yes	Yes	Yes

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
Number SPAN Sessions (RX or Both)	2	2	2
Number SPAN Sessions (TX)	4	4	4
R-SPAN	Yes	Yes	Yes
Number R-SPAN Source Sessions	1	1	2
Number R-SPAN Destination Sessions	24	24	24
ER-SPAN	No	No	Yes
Number ER-SPAN Source Sessions	–	–	2
Number ER-SPAN Destination Sessions	–	–	24
DHCP Snooping	Yes	Yes	Yes
Dynamic ARP Inspection	No	Yes	Yes
IP Source Guard	No	No	Yes

Among the many Layer 2 features available to the Supervisor Engine 32 via the PFC3B, there are a few unique features not available on earlier wiring closet supervisors. Port ACLs provide the functionality of a VLAN ACL (VACL) but can be applied on a single Layer 2 switch port within a VLAN (unlike a VACL, which is applied to the entire VLAN). The PACL can be applied on ingress traffic and will be processed prior to any VACLs that may be associated with the switch port.

ACL hit counters provide a way to monitor the number of times a specific access control entry (ACE) within an ACL has been used on traffic passing through the interface. Understanding hit patterns allows administrators to tune their ACLs to be more effective on the traffic they are applied to.

Enhanced Remote SPAN (ERSPAN) is a way to forward a copy of data to a destination SPAN port over multiple Layer 3 hops. Typically, ERSPAN might be applied in a campus where the source is on one subnet and the destination SPAN port is located at another subnet. ERSPAN uses generic routing encapsulation to carry the traffic over the Layer 3 network.

IP Source Guard is one of a number of new features now available as part of the Cisco Catalyst Integrated Security Toolkit (CIST). While making up part of this toolkit, IP Source Guard does, however, depend on hardware in the PFC3B to perform its functionality. The main benefit IP Source Guard provides is to protect against spoofed packets. Spoofed packets are a way for hackers to gain entry into a network by changing their source IP address to one that is recognized by the network as an “internal” address or a “secure” address. IP Source Guard uses Dynamic Host Configuration Protocol (DHCP) snooping to snoop on DHCP requests and build a dynamic PACL that denies all packets that do not match the assigned DHCP address. This PACL is applied at the interface level. This provides a level of protection against “spoofed” packets not found in earlier supervisors.

Supervisor Engine 32 with PFC3B Layer 3 Features

The PFC3B extends Layer 3 services for the Supervisor Engine 32 well beyond what earlier wiring closet supervisors have provided. The new Layer 3 features that are supported on a Supervisor Engine 32 are detailed in Table 9.

Table 9. Supervisor Engine 32 with PFC3 Layer 3 Features

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
Forwarding Architecture	Flow based	Cisco Express Forwarding based	Cisco Express Forwarding based
FIB TCAM	–	256 K	256 K
Adjacency Table	–	256 K	1 M
NetFlow Table	128 K (64 K)	128 K (64 K)	128 K (115 K)
IPv6	No	No	Yes
MPLS	No	No	Yes
VRF lite	No	No	Yes
GRE in Hardware	No	No	Yes
NAT in Hardware	No	No	Yes
PAT in Hardware	No	No	Yes
Multipath URPF	No	No	Yes
EoMPLS	No	No	Yes
HSRP/VRRP Groups	255	16	255

Cisco Express Forwarding is the forwarding architecture used by the Supervisor Engine 32. Cisco Express Forwarding utilizes both the MSFC2a and the PFC3B to build a forwarding information base (FIB) that provides the hardware with a view of the network topology. The MSFC2a will use configured routing protocols (Open Shortest Path First [OSPF], Enhanced Interior Gateway Routing Protocol [EIGRP], border Gateway Protocol [BGP], and so on) to interact with its Layer 3 peers and collect routing information about the network topology. The MSFC2a uses this information to build the FIB, which is then passed to the PFC. The PFC stores this information in a FIB ternary content addressable memory (TCAM). The FIB TCAM is located on the PFC daughter card and provides the system with very high-speed memory to enable fast route lookups during forwarding operations.

The PFC3B adds a number of hardware-based features to the portfolio of features available on the Supervisor Engine 32. The most significant of these are IPv6 and Multiprotocol Label Switching (MPLS), both of which were not available in the earlier Supervisor Engine 1 and Supervisor Engine 2. Enabling IPv6 and/or MPLS on the Supervisor Engine 32 allows any Ethernet port in the chassis to receive and transmit these packet types. Both of these features significantly increase the deployment options available to networks using the Supervisor Engine 32. Support for these and other Layer 3 options will be enabled by a future software release.

Supervisor Engine 32 Control Plane Protection

From a security perspective, the incorporation of CPU rate limiters in the hardware strengthens the Supervisor Engine 32 from attacks that can compromise the operational running and performance of the entire switch. The control plane is a critical part of the processing capabilities for the switch, handling functions such as Spanning Tree Protocol, logging operations, handling Simple Network Management Protocol (SNMP) events, and managing the CLI. CPU rate limiters define a set of rate limiters that can be applied to different forms of traffic destined to the control plane. With the FCS version of the Supervisor Engine 32, the list of CPU rate limiters in Table 10 can be enabled.

Table 10. CPU Rate Limiters

Form of Rate Limiter	Rate Limiter	Function
Unicast Rate Limiter	VACL logging	CLI notification of VACL denied packets
Layer 2 Rate Limiter	Layer 2 protocol tunneling	L2PT encapsulation/deencapsulation
	Protocol data units	Process L2 PDUs
General Rate Limiter	MTU failure	Packets requiring fragmentation

Also built into the Supervisor Engine 32 hardware are a number of Layer 3 CPU rate limiters that will be enabled in a future software release. These CPU rate limiters are designed to protect the performance of the Layer 3 control plane (or route processor) from being compromised. The Layer 3 rate limiters are grouped in Table 11 into two sections: unicast and multicast rate limiters.

Table 11. Layer 3 CPU Rate Limiters

Form of Rate Limiter	Rate Limiter	Function
Unicast Rate Limiter	Cisco Express Forwarding glean	ARP packets
	Cisco Express Forwarding receive	Traffic destined to the router (MSFC)
	Cisco Express Forwarding no route	Packets with no route in the FIB table
	IP errors	Packets with IP checksum or length errors
	ICMP redirect	Packets requiring ICMP redirect
	ICMP no route	ICMP unreachable for unroutable packets
	ICMP ACL drop	ICMP unreachables for admin deny packets
	RPF failure	Packets failing URPF check
	L3 security	Packets needing CBAC, IPSec, or authentication proxy processing
	ACL input	Logs on ACLs for TCP intercept, NAT or Reflexive ACLs
	ACL output	Logs on ACLs for TCP Intercept, NAT, or reflexive ACLs
	IP options	Unicast packets with IP options set
	Capture	Used with optimized ACL logging
Multicast Rate Limiter	Multicast FIB miss	Packets with no mroute in the FIB
	IGMP	For IGMP packets
	Partial shortcut	For partial shortcut entries
	Directly connected	Local multicast on connected interface
	IP options	Multicast traffic with IP options set
	IPv6 directly connected	Packets with no mroute in the FIB
	IPv6 *,G M bridge	starg-m-bridge packets matching (*,G/m)SM, (FF::/8)
	IPv6 *,G bridge	Packets matching (*,G/128)
	IPv6 S,G bridge	Packets matching (S,G)RP-RPF post-so, (*,FFx2/16)
	IPv6 route control	Packets matching (*,FF02::X/128)
	IPv6 default drop	L3 drops on (*,G/m)SSM entries, (*,G/m)SSM non-RPF
	IPv6 secondary drop	L3 drops on (*,G/128) SPT threshold is infinity
General Rate Limiter	TTL failure	Packets with TTL <= 1

Another feature built into the hardware is control plane policing. This capability introduces a new control plane interface that can have a QoS policy applied to it to limit the total amount of traffic that is forwarded to the control plane. This feature will be enabled via a future software release.

Supervisor Engine 32 QoS Features

QoS features built into the Supervisor Engine 32 are listed in Table 12.

Table 12. Supervisor Engine 32 QoS Features

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
TX uplink Queue Structure	1p2q2t	1p2q2t	1p3q8t
RX Uplink Queue Structure	1p1q4t	1p1q4t	2q8t
Uplink Buffer per Port	512 KB	512 KB	9.5 MB for GE ports, 100 MB for 10GE
Uplink Port Scheduler	WRR	WRR	DWRR/SRR
QoS ACEs	16 K	32 K	32 K
Aggregate Policers	1023	1023	1023
Unique Microflow Policing Rates	63	63	63
Microflow Flow Mask	Full flow	Full flow	Supports 2 flow masks (SRC-ONLY and DEST-ONLY)
User-Based Rate Limiting	No	No	Yes
DSCP Transparency	No	No	Yes
Egress Aggregate Policing	No	No	Yes

Many of the QoS features in the Supervisor Engine 32 are consistent with other supervisor module QoS capabilities. The Supervisor Engine 32 carries across QoS features from the Supervisor Engine 2 as well as inheriting some of the QoS features from the Supervisor Engine 720. It also includes a new QoS feature (Shaped Round Robin [SRR]) not yet found in other modules that is available on its uplink ports.

The queue structure found in the uplink ports is quite different from that in earlier supervisors. On the transmit side, each Gigabit Ethernet uplink port now contains a single strict priority queue along with three normal queues. Each of the normal transmit queues on the Gigabit Ethernet port is primed with eight thresholds, which allow the port congestion management algorithm to provide very granular congestion control. On the receive side, there are two normal queues, each of which has eight thresholds for congestion management. There is no strict priority queue on the ingress port. This QoS port type is consistent with the latest Ethernet modules.

Per-port buffering has also been significantly increased for each of the uplink ports on the Supervisor Engine 32 front panel. While the earlier supervisor uplink ports supported 512 Kb of per-port buffering, each of the Supervisor Engine 32 Gigabit Ethernet uplink ports is provided with 9.5 MB of buffering. The 10-Gigabit Ethernet ports will have 100 MB of per-port buffering. This amount of buffering is of particular importance to those networks wishing to deploy bursty applications or high data volume applications (for example, network video), which can use the extra buffering should they need it.

As with the Supervisor Engine 720, new features such as differentiated services code point (DSCP) transparency are supported on the Supervisor Engine 32. DSCP transparency is a new mechanism that maintains the integrity of the DSCP as it transits the switch. Certain situations can arise where, say, a packet arrives on an untrusted port and the switch will assign a zero class-of-service (CoS) value to the packet. From this CoS value, an internal priority is derived that is used by the switch to write the type of service (ToS) and CoS priority bits on egress. DSCP transparency protects against this situation, and others, by negating the option of using the internal priority to derive the egress DSCP value. Rather, the ingress DSCP value will simply be written on egress.

SRR and Deficit Weighted Round Robin (DWRR) are two scheduling mechanisms that can be configured on the Supervisor Engine 32 Gigabit Ethernet ports. SRR is an alternative scheduling mechanism to DWRR. SRR introduces the concept of allowing the administrator to define the maximum amount of bandwidth that each queue is allowed to use. This is unique to SRR, and there is no equivalent function within DWRR to provide the same facility. Configuration of this feature, like DWRR, still requires a weight to be configured on each of the queues, but the manner in which the SRR algorithm uses the “weight” value is different. After assigning a “weight” to each queue, the total of the weights is normalized by the SRR algorithm to 100 percent. Depending on the hardware granularity of the given queue, a bandwidth value is derived and assigned to the queue. The outbound flow of data will then be shaped to this bandwidth value. Unlike DWRR, a given queue that is shaped will not be able to exceed the defined bandwidth value. Traffic in excess of the defined rate will be buffered, resulting in a “smoothing out” of the traffic over a given period of time.

DWRR tries to provide a fairer allocation of bandwidth between the queues, more so than normal Weighted Round Robin (WRR). While the weights determine what allocation of bandwidth each queue is allowed access to, the DWRR algorithm maintains a counter of excess bandwidth used by each queue. For example, say a given queue has used up all but 500 bytes of its allocation, but has another packet in the queue (that is a full 1500 bytes in size); then the packet is scheduled (transmitted), but the queue has used 1000 bytes of bandwidth in excess of its allocation on that pass of the queue. The DWRR algorithm “remembers” the extra 1000 bytes used and deducts this from the queue’s bandwidth allocation the next time the queue is serviced. Over a period of time, the queues will statistically get a lot closer to using their allocated portion of bandwidth.

Supervisor Engine 32 Multicast Features

Multicast features supported on the Supervisor Engine 32 are listed in Table 13.

Table 13. Supervisor Engine 32 Multicast Features

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
Multicast Forwarding Performance	15 Mpps	15 Mpps	15 Mpps
Number of MROUTES	16 K	16 K	32 K
S,G Forwarding in Hardware	Yes	Yes	Yes
*,G Forwarding in Hardware	Yes	Yes	Yes
Bidirectional PIM in Hardware	No	No	Yes
Outgoing Interfaces (OIF)	64 K	64 K	64 K
Central Replication	Yes	Yes	Yes
IGMP V1 Snooping	Yes	Yes	Yes
IGMP V2 Snooping	Yes	Yes	Yes

Feature	Supervisor Engine 1A with PFC	Supervisor Engine 2 with PFC2	Supervisor Engine 32 with PFC3B
IGMP V3 Snooping	No	Yes	Yes
IGMP Querier	Yes	Yes	Yes
Future MVPN Support	No	No	Yes

The Supervisor Engine 32 extends support for multicast in a number of areas. The most significant of these is the increased capacity for storing multicast routes (MROUTES). The earlier supervisor models supported a maximum of 16 K S, G, or *,G entries. The Supervisor Engine 32 supports up to 32 K MROUTES, more than doubling the capacity found on previous supervisors.

Bidirectional Protocol Independent Multicast (PIM) is also another new feature, which is enabled with the presence of a PFC3B on the Supervisor Engine 32. Bidirectional PIM builds a bidirectional multicast distribution tree, which supports bidirectional traffic flow. The advantage that bidirectional shared trees provide is that many multicast sources can send on the same tree without Layer 3 devices having to explicitly keep state for each source. This has the added benefit of reducing the load on the supervisor's CPU and memory.

The PFC3B supports MPLS natively in the hardware, which will allow the Supervisor Engine 32 to take advantage of multicast over MPLS VPN (MVPN) when it becomes available in a future release of software. MVPN extends multicast support across MPLS networks, allowing more deployment options for customers looking to deploy multicast.

Like other Layer 3 features, most of the Layer 3 multicast features will be enabled in a future release of software.

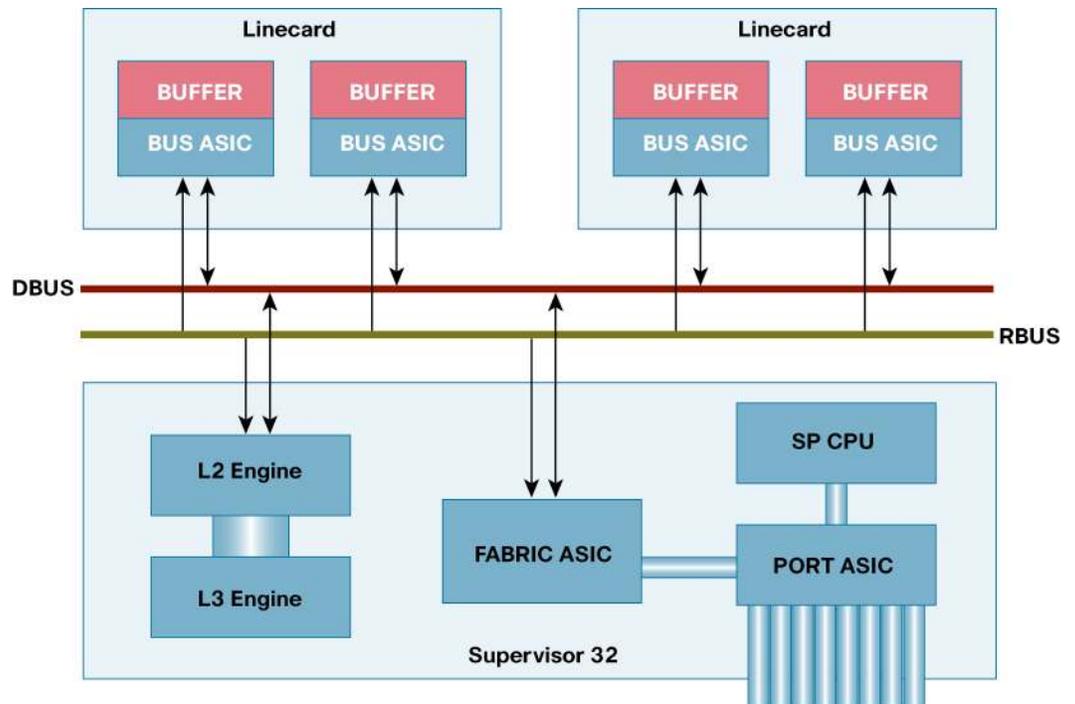
Packet Flow Through the Supervisor Engine 32

This section will describe the packet flow through the Supervisor Engine 32.

Packet Flow—Shared Bus

The packet walk will use the diagram in Figure 8 to explain the steps taken to pass a packet through the Supervisor Engine 32.

Figure 8. Base Packet Flow Architecture

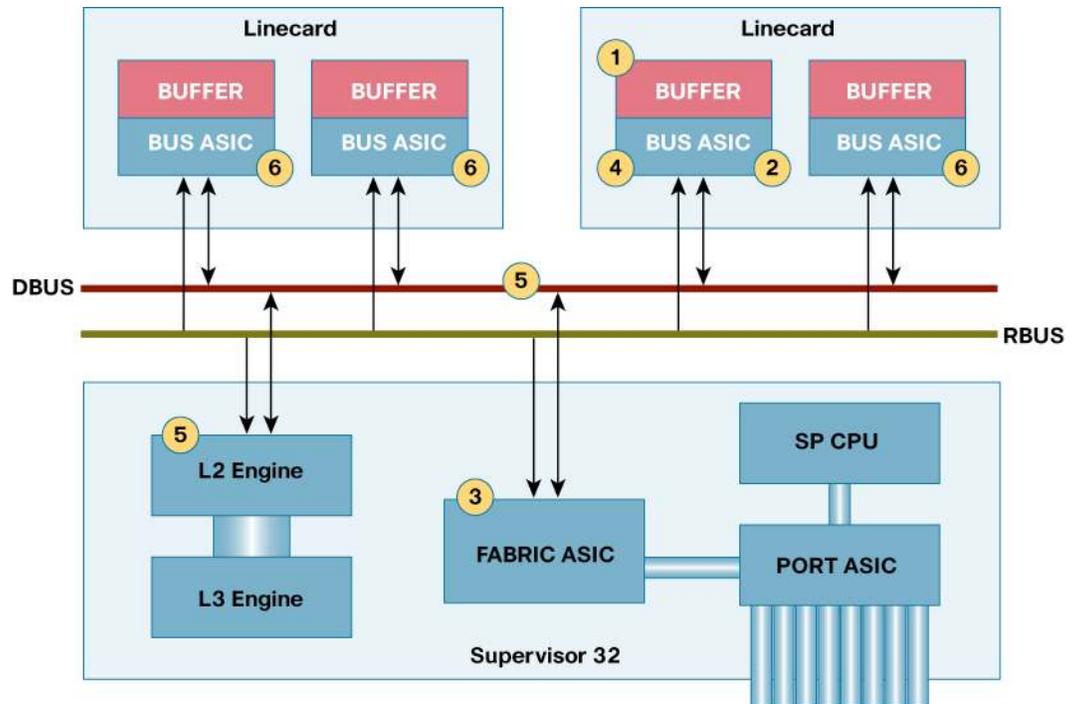


Packet Flow Steps

1. The packet initially arrives at the port and is placed temporarily in an input buffer. Using information in the existing packet header, the port ASIC will build an internal header containing information that the central forwarding engine will use to perform a forwarding table lookup and apply QoS and security policies (if configured). Ingress QoS can also be performed here if configured.
2. The local port ASIC on the line card will arbitrate for access to the bus to perform packet transmission. There is a local arbitration mechanism built into each line card that communicates with the central arbitration process running on the Supervisor Engine 32.
3. If the bus is not in use, then the central arbitration mechanism on the supervisor will forward a message to the line card arbitration mechanism indicating that it is able to begin transmission.
4. When the arbitration process on the module receives acknowledgement from the central arbiter on the supervisor, the port ASIC forwards the packet on the local shared bus.
5. The packet is forwarded over the bus to the supervisor and will be received by the Layer 2 forwarding engine on the PFC3B.
6. As the bus is a shared medium, all other line cards connected to the shared bus will see the packet and store that packet temporarily in their transmit buffers. This packet will stay in those buffers until the supervisor instructs the line card to either forward or drop the packet.

Figure 9 illustrates steps 1 through 6.

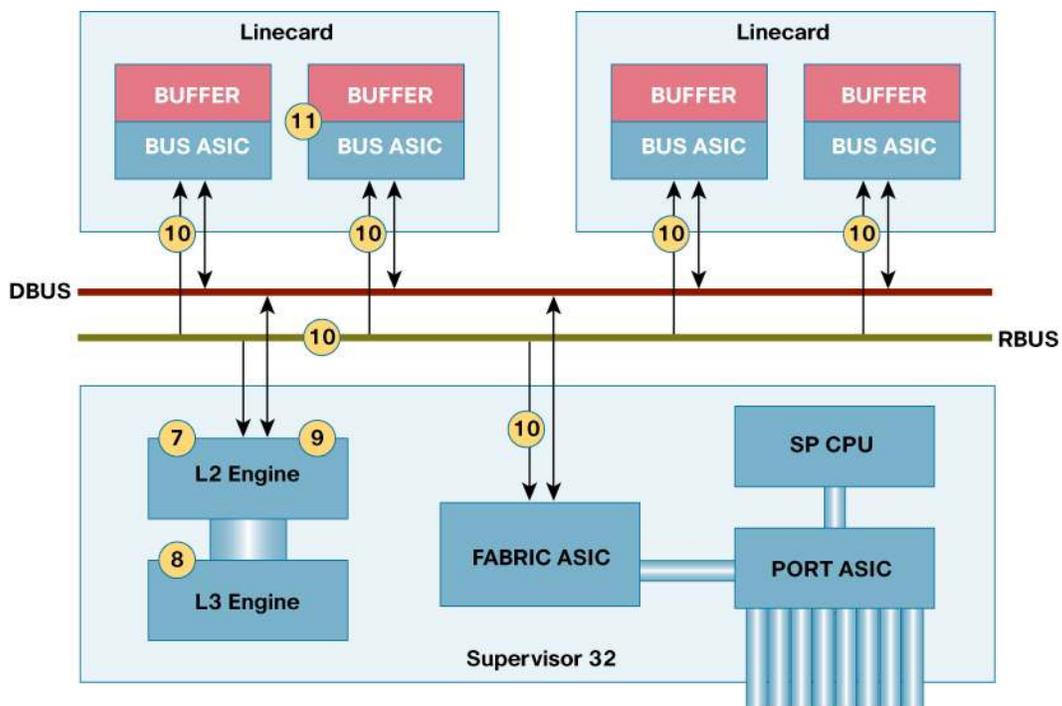
Figure 9. Steps 1 Through 6 of the Packet Flow



7. The Layer 2 engine on the PFC will perform a Layer 2 lookup using the destination MAC address. Following this operation, the packet is passed to the Layer 3 engine for further processing.
8. The Layer 3 engine on the PFC then performs a number of processes in parallel. If an MSFC was present, it will use Cisco Express Forwarding to populate the FIB TCAM on the PFC with a view of the network topology. The FIB is built from the master routing tables that are located on the MSFC. The Layer 3 engine would perform a lookup in the FIB if the forwarding operation were deemed a Layer 3 switching operation. In parallel to this, a lookup is also performed on the QoS and security ACLs to see if any of the ACLs need to be applied to this packet. NetFlow statistics will also be updated for the flow that this packet is a part of.
9. The results of all lookup operations are then pulled together by the PFC. This result contains the following information:
 - a. Instructions to either forward or drop the packet
 - b. MAC rewrite information necessary to modify the Layer 2 MAC destination address so the packet arrives at its correct next hop destination
 - c. QoS information necessary to place the packet into its correct queue and any rewrite information necessary for adjusting CoS and ToS values
10. The result is forwarded over the results bus to all destination ports.
11. The destination port will receive the results information and use this to build the Ethernet header for the packet. The packet is retrieved from its local buffer and forwarded out the physical interface.

Figure 10 illustrates steps 7 through 11.

Figure 10. Steps 7 Through 11 of the Packet Flow



Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
 www.cisco.com
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems, Inc.
 168 Robinson Road
 #28-01 Capital Tower
 Singapore 068912
 www.cisco.com
 Tel: +65 6317 7777
 Fax: +65 6317 7799

Europe Headquarters
 Cisco Systems International BV
 Haarlerbergpark
 Haarlerbergweg 13-19
 1101 CH Amsterdam
 The Netherlands
 www-europe.cisco.com
 Tel: +31 0 800 020 0791
 Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)