



Product Overview

This chapter describes the features of the Catalyst 4840G SLB switch. It shows how a Catalyst 4840G SLB switch fits into the network, and lists the type of interfaces used in the switch. Selected features of Server Load Balancing (SLB), Firewall Load Balancing (FWLB) and Layer 3 functionality are described briefly. This chapter includes the following sections:

- SLB Overview, page 1-1
- Network Configuration Example, page 1-2
- SLB and Layer 3 Switching Interfaces, page 1-3
- SLB Features, page 1-3
- Layer 3 Switching Software Features, page 1-4
- Network Management Features, page 1-6
- Supported Key Features for SLB, page 1-7

SLB Overview

The Catalyst 4840G SLB switch is a high-performance switch for the campus LAN or intranet, providing both wire-speed Ethernet routing and switching for a load-balancing device. The Catalyst 4840G SLB switch provides secure and reliable web and application hosting services to your Internet or intranet clients.

The simplest and most often enabled method to address increasing website traffic and reliability requirements involves multiple web servers and SLB switches. Each server has identical website content and usually runs mirroring software to maintain duplicate content across all the servers in the server farm. The Catalyst 4840G SLB switch redistributes requests from clients evenly among all the servers in the server farm and achieves a balanced load for each server.

All of the physical (real) servers appear as one virtual server, resulting in only a single IP address and a single URL required for an entire server farm. By distributing client requests across a server farm, the Catalyst 4840G SLB switch optimizes responsiveness and system capacity while ensuring scalability. The Catalyst 4840G SLB switch also dramatically improves site reliability, by allowing individual servers to fail or be taken off line without a significant impact on content flow to the users.

In addition, the load-balancing switch performs the following three major functions:

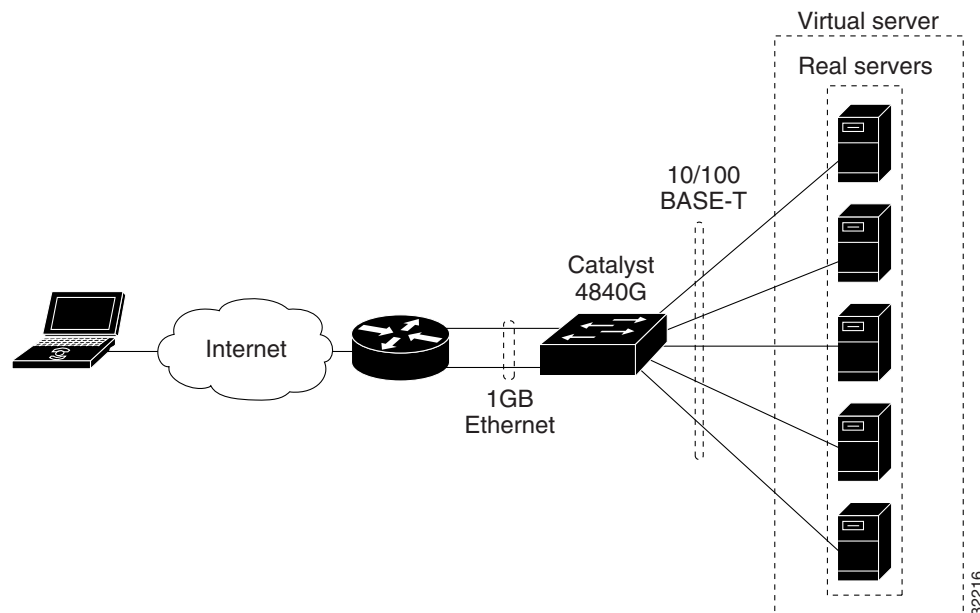
- Packet switching
- Route processing
- Intelligent network services

Compared to other routers, the Catalyst 4840G SLB switch processes more packets faster, by using ASIC hardware instead of microprocessor-based engines. Layer 3 switch routers also improve network performance with two software functions: route processing and intelligent network services.

Network Configuration Example

Figure 1-1 shows how you can use the Catalyst 4840G SLB switch in an enterprise-wide network.

Figure 1-1 Typical Network Configuration for the Catalyst 4840G SLB Switch



The Catalyst 4840G SLB switch intelligently load balances TCP/IP and UDP traffic across multiple servers. It appears as one virtual server to the requesting clients. All traffic is directed toward a virtual IP address (a virtual server) through DNS. Those requests are distributed over a series of real IP addresses on servers (real servers). A virtual IP address is an address that is in DNS and most likely has a domain name. A real IP address is physically located on a real server behind SLB.

Load balancing provides the following benefits:

- High performance is achieved by distributing client requests across a cluster of servers.
- Administration of server applications is easier. Clients know only about virtual servers; no administration is required for real server changes.
- Security of a real server is ensured because its address is never announced to the external network. Users are familiar with only the virtual IP address. In addition, filtering of unwanted traffic can be based on both IP addresses and TCP or UDP port numbers.
- Ease of maintenance with no downtime is achieved by allowing physical (real) servers to be transparently placed in or out of service.

For a detailed explanation of SLB, see Chapter 5, “Server Load Balancing.”

SLB and Layer 3 Switching Interfaces

The Catalyst 4840G SLB switch uses the following interfaces for load balancing and Layer 3 switching:

- Two Gigabit Ethernet interfaces for Layer 3 connections to the clients
- Forty Fast Ethernet interfaces that can be configured as Layer 2 or Layer 3 for attaching SLB servers

Fast Ethernet Client Capability

The Catalyst 4840G SLB switch allows both clients and servers to attach to the Gigabit Ethernet interfaces. The 40 Fast Ethernet 10/100 BASE-T ports are configured to connect only to servers. However, you can enable ports 37–40 for clients to attach to these ports. See Chapter 3, “Configuring the Catalyst 4840G Processor,” for information on enabling this capability.

SLB Features

This section lists Catalyst 4840G SLB switch software features.

Modes of Load Balancing

- Dispatched mode
- Directed mode (Network Address Translation of server addresses and client addresses is supported)

SLB Algorithms

- Weighted round robin
- Weighted least connections

Configurable SLB Features

- Sticky connections
- HTTP redirection load balancing
- Firewall load balancing (FWLB)
- Buddy groups
- Direct server access
- Clients and servers in the same bridged domain as the load balancer
- Clients and servers routed to the load balancer

Server/Application Availability Detection

- HTTP probe
- Ping probe
- Server Failure Detection (Dynamic Feedback Protocol [DFP])
- Interacton with DistributedDirector for global load balancing

Protocols Supported by Load Balancing

- Domain Name System (DNS)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol (HTTP)
- Hypertext Transfer Protocol-Server (HTTP-S)
- Internet Message Access Protocol (IMAP)
- Mapping of Airline Traffic over IP-Target A (MATIP-A)
- Network News Transfer Protocol (NNTP)
- Post Office Protocol 2 (POP2)
- Post Office Protocol 3 (POP3)
- Real audio and Real video using HTTP
- Remote Authentication Dial-In User Service (RADIUS)
- Simple Mail Transfer Protocol (SMTP)
- Telnet
- X.25 over TCP (XOT)

Layer 3 Switching Software Features

This section lists Layer 3 switching software features of the Catalyst 4840G SLB switch.

Layer 1 Features

- 10/100BASE-TX half duplex and full duplex
- 1000BASE-SX, LX, and long-haul full duplex

Layer 2 Bridging Features

- Layer 2 transparent bridging
- Layer 2 MAC learning, aging, and switching by hardware
- Spanning Tree Protocol (IEEE 802.1D) per bridge group
- A maximum of 16 active bridge groups
- Up to 4096 MAC addresses
- Integrated routing and bridging (IRB)
- 24K content-addressable memory (CAM) shared by Layer 2 entries and IP routing

Virtual LAN (VLAN) Features

- Inter-Switch Link (ISL)-based VLAN trunking (not supported in conjunction with FWLB; use IEEE 802.1Q)
- IEEE 802.1Q-based VLAN trunking

Layer 3 Routing, Switching, and Forwarding

- IP routing and switching between Ethernet ports
- Constrained multicast flooding (CMF)
- QoS-based forwarding founded on IP precedence-based forwarding
- Load balancing among equal cost paths based on source and destination IP addresses
- 24K CAM shared by Layer 2 entries and IP routing
- Up to 18000 IP routes
- Up to 20000 IP host entries

Supported Routing Protocols

- Routing Information Protocol (RIP and RIP II)
- Interior Gateway Routing Protocol (IGRP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Open Shortest Path First (OSPF)
- Protocol Independent Multicast (PIM)—sparse and dense modes
- Secondary addressing
- Static routes

Fast EtherChannel (FEC) Features

- Bundling of up to four Fast Ethernet ports
- Load sharing based on source and destination IP addresses of unicast packets
- Load sharing for bridge traffic based on MAC address
- ISL on the FEC
- IRB on the FEC
- IEEE 802.1Q trunking on the FEC
- Up to 10 active FEC port channels

Gigabit EtherChannel (GEC) Features

- Bundling of the two Gigabit Ethernet ports
- Load sharing based on source and destination IP addresses of unicast packets
- Load sharing for bridge traffic based on MAC address
- ISL on the GEC
- IRB on the GEC
- IEEE 802.1Q trunking on the GEC
- One active GEC port channel in one system

Additional Protocols and Features

- Bootstrap Protocol (BOOTP)
- Cisco Discovery Protocol (CDP) support on Ethernet ports
- Cisco Group Management Protocol (CGMP) server support

- Dynamic Host Configuration Protocol (DHCP) relay
- Hot Standby Routing Protocol (HSRP) over 10/100 Ethernet, Gigabit Ethernet, FEC, GEC, and Bridge Group Virtual Interface (BVI)
- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- IRB routing mode support
- Simple Network Management Protocol (SNMP)

Network Management Features

This section lists the network management features of the Catalyst 4840G SLB switch.

Configuration-Related Feature

- Support for all relevant configuration commands

Management

- HTTP server
- SNMP
- Telnet

Performance Monitoring

- RMON groups Statistics, History, Status, Event, and Alarm

Security Features

- TACACS+ (authentication only)

SNMP MIB Support

- Chassis MIB
- CISCO-CDP-MIB
- CISCO-ENVMON-MIB
- CISCO-SLB-MIB
- Ethernet MIB
- Flash MIB, ping MIB, image MIB, config man MIB, memory pool MIB
- MIB II
- RFC 1493 (Bridge MIB)
- RFC 2037 (Entity MIB)
- RS232 MIB, RMON alarm and event groups, new interface MIB (RFC1573)

Supported Key Features for SLB

This section briefly describes the key features that are supported in the SLB switching software and includes the following sections:

- Cisco Discovery Protocol, page 1-7
- Cisco Express Forwarding, page 1-7
- Content Flow Monitor Support, page 1-8
- Distributed Hardware Processing, page 1-8
- Fast EtherChannel, page 1-8
- Gigabit EtherChannel, page 1-9
- Integrated Routing and Bridging, page 1-9
- Network Class Redundancy, page 1-10
- PVST+ Spanning Tree Protocol, page 1-10
- QoS-Based Forwarding, page 1-10
- Remote Monitoring, page 1-10
- Routing Protocols, page 1-11
- SLB Algorithms, page 1-11
- SLB Nonstateful Backup—HSRP, page 1-11
- SLB Redirection Modes, page 1-12
- SLB Software Operation, page 1-12
- SLB Stateful Backup Feature, page 1-12
- Virtual LANs, page 1-13

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) is a device-discovery protocol that is both media and protocol independent. CDP is available on all Cisco products, including routers, switches, bridges, and access servers. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN. CDP enables Cisco products to exchange with each other information about their MAC addresses, IP addresses, and outgoing interfaces. CDP runs over the data link layer only, which allows two systems that support different network-layer protocols to learn about each other.

Each device configured for CDP sends periodic messages to a multicast address. Each device advertises at least one address at which it can receive SNMP messages.

Cisco Express Forwarding

Catalyst 4840G SLB software supports Cisco Express Forwarding (CEF), an advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions. Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP backbone switching.

CEF manages route distribution and forwarding by distributing routing information from the route processor to the individual Ethernet interfaces. This technology, used within an intranet, provides scalability in large campus core networks. CEF provides Layer 3 forwarding based on a topology map of the entire network, resulting in high-speed routing table lookups and forwarding.

One of the key benefits of CEF in Layer 3 switching is its routing convergence. Because the forwarding information base (FIB) is distributed to all interfaces, whenever a route goes away or is added, the FIB updates that information and provides it to the interfaces so that route processor interrupts are minimized. The interfaces receive the new topology very quickly and reconverge around a failed link based on the routing protocol being used.

**Caution**

We strongly recommend that you do *not* enter any CEF configuration commands. The CEF default settings should not be altered; doing so might adversely affect the performance of your system.

Content Flow Monitor Support

SLB supports the Cisco Content Flow Monitor (CFM), a Web-based status monitoring application within the CiscoWorks2000 product family. You can use CFM to manage Cisco server load-balancing devices. CFM runs on Windows NT and Solaris workstations and is accessed using a Web browser.

Distributed Hardware Processing

Load-balancing software uses a distributed architecture in which the control path and data path are relatively independent of one another. The control path code, such as SLB decision processes, runs on the route processor, while most of the data packets are forwarded by the Ethernet interface and the switching fabric.

The 4840G uses ten interface processors for the Fast Ethernet interfaces, with four ports supported per interface. To achieve maximum performance, you must distribute the servers evenly across the ten interface processors.

Each interface includes a microcoded processor that handles all packet forwarding. The following are the main functions of the control layer between the routing protocol and the firmware data path microcode:

- Managing the internal data and control circuits for the packet forwarding and control functions
- Extracting routing and packet forwarding control information from Layer 2 and Layer 3 bridging, Layer 4 load balancing, routing protocols, and configuration data, and then conveying the information to the interface to control the data path
- Collecting data path information, such as traffic statistics, from the interface to the route processor
- Handling certain data packets sent from the Ethernet interfaces to the route processor

Fast EtherChannel

Fast EtherChannel (FEC) establishes a high-bandwidth connection between two Layer 3 switch devices. You can use up to four Fast Ethernet connections as one Layer 3 forwarding path; this can provide up to 800 Mbps full-duplex aggregate capacity. If link detection determines a failure of any one link, the packets are switched to the remaining active links in the FEC.

FEC uses a source-destination IP address load-balancing scheme for up to four ports in a channel group. Each channel group has its own IP address. When a packet is queued to exit the port channel interface, the last two bits of the IP source and destination address determine which interface in the channel the packet will take.

To configure the EtherChannel, see Chapter 10, “Configuring EtherChannel.”

Gigabit EtherChannel

Gigabit EtherChannel (GEC) allows grouping of two Gigabit Ethernet interfaces into a single multigigabit logical EtherChannel link. GEC establishes a high-bandwidth connection between two Catalyst switch devices.

GEC uses a source-destination IP address load-balancing scheme for the two Gigabit Ethernet ports in the channel group. Each channel group has its own IP address. When a packet is queued to exit the port channel interface, the last two bits of the IP source and destination address determine which interface in the channel the packet will take.

As with all EtherChannel technologies, the traffic load is shared across all links within the bundled ports; convergence occurs within one second of a GEC failure.

To configure the EtherChannel, see Chapter 10, “Configuring EtherChannel.”

Integrated Routing and Bridging

Integrated routing and bridging (IRB) allows you to route a given protocol between routed interfaces and various bridge groups or between bridge groups within a single switch. Multiple ports in the switch can reside in one bridge group with one IP address and be routed to other switch interfaces with different IP addresses.

Specifically, local or unroutable traffic is bridged among the bridged interfaces in the same bridge group, while routable traffic is routed to other routed interfaces or bridge groups.

Layer 3 switching software supports IRB for IP only.

Here are some examples of when to use IRB:

- When you want to interconnect a bridged network with a routed network, the IRB feature enables the switch to act as a true router
For example, when you are migrating a bridged network to a routed network, or when the remote site does not have routing capabilities, you can use the switch to interconnect the bridged and routed networks.
- When you want to conserve IP addresses by connecting network segments with bridges and assigning each bridge group one network address
- When you want to break one large segment into several small segments to improve the performance of the end stations

To configure IRB, see the “Using IRB with BVI” section on page 9-4.

Network Class Redundancy

The redundancy of Catalyst software provides key network features, such as HSRP for both routing and load balancing, routing protocol convergence with RIP, OSPF, EIGRP, FEC, and load sharing across equal-cost Layer 3 paths and spanning trees (for Layer 2-based networks).

PVST+ Spanning Tree Protocol

The PVST+ Spanning Tree Protocol is a bridge protocol that enables a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange bridge protocol data unit (BPDU) messages with other bridges to detect loops and then remove the loops by shutting down selected bridge interfaces.

PVST+ is a technique for maintaining a network of multiple bridges or switches. When the topology changes, PVST+ transparently reconfigures bridges and switches to avoid the creation of loops, by placing ports in a forwarding or blocking state. Each VLAN is treated as a separate bridge, and a separate instance of PVST+ is applied to each.

Spanning Tree Protocol parameters are set for each VLAN. For each spanning tree instance, you configure a set of global options with a set of port parameters. The port parameter list contains only ports that are members of a given VLAN. A maximum of 64 spanning tree instances are supported, one for each VLAN.

To configure PVST+, refer to the *Cisco IOS Bridging and IBM Networking Configuration Guide*.

QoS-Based Forwarding

Quality of Service (QoS) includes technologies such as Resource ReSerVation Protocol (RSVP) and weighted round-robin (WRR), which help control bandwidth, network delay, jitter, and packet loss in congested networks. In the Catalyst 4840G SLB switch, QoS-based forwarding sorts traffic into a small number of classes and marks the packets accordingly. The QoS identifier provides specific treatment of traffic in different classes, so that a different quality of service is provided to each class.

Frame and packet scheduling and discarding policies are determined by the class to which the frames and packets belong. For example, the overall service given to frames and packets in the premium class will be better than the service given to the standard class; the premium class is expected to experience lower loss rate or delay.

The Catalyst 4840G SLB switch has QoS-based forwarding for IP traffic only. The implementation of QoS forwarding is based on local administrative policy and IP precedence. The mapping between the IP precedence field and the QoS field determines the delay priority of the packet.

Remote Monitoring

Catalyst 4840G software supports the first four remote monitoring (RMON) groups.

RMON is a network management protocol for gathering network information and monitoring traffic data within remote LAN segments from a central location. RMON allows you to monitor all nodes and their interaction on a LAN segment. RMON used in conjunction with the SNMP agent in the switch allows you to view both the traffic that flows through the switch and segment traffic not necessarily destined for the switch. Load balancing software combines RMON alarms and events with existing MIBs so that you can choose where monitoring will occur.

Routing Protocols

In addition to comprehensive load-balancing capability, Layer 3 switching software provides a comprehensive suite of routing protocols based on Catalyst 4840G software:

- RIP
- RIP II
- OSPF
- IGRP
- EIGRP

Many of the Catalyst 4840G software routing protocol features, such as route redistribution and load balancing over equal cost paths (for OSPF and EIGRP), are supported. Configuration of these routing protocols is identical to the configuration methods currently used on all Cisco routers.

To configure network and routing protocols, see Chapter 8, “Configuring Networking Protocols.”

SLB Algorithms

There are two algorithms for server load balancing or determining which server will receive a new connection request. Both of these algorithms are weighted.

- Round robin—This is the default algorithm; it directs the network connection to the next server and, if unweighted, treats all servers as equal, regardless of the number of connections or the response time.
- Least connections—The number of sessions assigned to a server is based on the number of current TCP connections.

**Note**

You can configure a real server with a weight relative to other real servers in the server farm, using the **weight (server farm)** real server configuration command.

SLB Nonstateful Backup—HSRP

HSRP provides high network availability, by routing IP traffic from hosts on Ethernet networks without relying on the availability of any single router. This feature is particularly useful for hosts that do not support a router discovery protocol, such as the Intermediate System-to-Intermediate System (IS-IS) Interdomain Routing Protocol (IDRP). It is also useful for hosts that do not have the functionality to switch to a new router when their selected router reloads or loses power.

Devices that are running HSRP detect a failure by sending and receiving multicast User Datagram Protocol (UDP) hello packets. When HSRP detects that the designated active router has failed, the selected backup router assumes control of the HSRP group MAC and IP addresses. (You can also select a new standby router at that time.)

The chosen MAC and IP addresses are unique and do not conflict with any others on the same network segment. The MAC address is selected from a pool of Cisco MAC addresses. You configure the last byte of the MAC address by configuring the HSRP group number. You also configure the unique virtual IP address. The IP address must be specified on a single router within the same group. When the HSRP is running, it selects an active router and instructs its device layer to listen on an additional (dummy) MAC address.

Load-balancing software supports HSRP over 10/100 Ethernet, Gigabit Ethernet, FEC, GEC, and BVI, thus ensuring that traffic from the load-balancing servers toward the clients goes through the active device using the HSRP IP address as a default gateway.

To configure HSRP, see the “Hot Standby Router Protocol (HSRP)” section on page 7-1.

SLB Redirection Modes

SLB provides for increasing website traffic and reliability by using multiple real Web servers. The SLB switch tracks network sessions and server load conditions in real time, directing each session to the most appropriate server.

The SLB switch can be configured to operate in one of two redirection modes: directed or dispatched. In directed mode the virtual server can be assigned an IP address that is not known to any of the real servers. SLB software translates packets exchanged between a client and a real server, translating the virtual server IP address to the real server address through network address translation (NAT). In dispatched mode the virtual server address is known to the real servers, and SLB redirects packets to the real servers at the MAC layer.

SLB Software Operation

SLB software watches the TCP connection setup and takedown handshake to identify the beginning and end of individual flows; consequently it can make the appropriate decisions as the conversation progresses.

The synchronize sequence number (SYN) is set in the first packet sent by the client to the IP address of the virtual server; in reality, this is the address of the SLB switch in behalf of all the real servers, so the packet is actually sent to the SLB software. The software maintains tables of all connections that are active, and because this is a new connection (identified by source and destination IP address as well as port numbers), the software creates a connection entry for this flow. The software then calculates which real server will handle this connection request, based on the configuration of the particular load-balancing algorithm.

After this decision has been made, the software then forwards the packet to the appropriate server by changing the destination IP or MAC address to that of the real server and sending the packet to the server. The real server receives the packet, gets the requested Web page, and sends the response back to the load-balancing software, which then sends the packet back to the original source workstation.

If multiple pages are requested during the user session, this process continues until the user asks for no more data and requests that the TCP connection be closed. At this point the load-balancing switch flushes its internal table of connection status for this particular session. It then calculates the server to receive the next new connection, based on the load-balancing algorithm.

SLB Stateful Backup Feature

The stateful backup feature enables SLB to incrementally back up its load-balancing decisions (keep state) between primary and backup Catalyst 4840G SLB switches. The backup switch has its virtual servers in a dormant state until failover is detected by HSRP; then SLB begins advertising virtual addresses and filtering traffic.

This enhancement provides SLB with a one-to-one stateful or idle backup scheme. Only one instance of SLB is handling client or server traffic at a given time and, at most, one backup platform for each active SLB switch. The state transfer between the primary and backup switches flows over the Content Aware Services Architecture (CASA) protocol, and the configuration of primary and backup switches is controlled at a lower level by HSRP groups and priorities.

To configure stateful backup, see the “SLB Stateful Backup” section on page 7-7.

Virtual LANs

A virtual LAN (VLAN) configures switches and routers according to logical rather than physical topologies. Using VLANs, a network administrator can combine any collection of LAN segments within an internetwork into an autonomous user group, which appears as a single LAN. VLANs logically segment the network into different broadcast domains, so that packets are switched only between ports within the VLAN. A VLAN usually corresponds to a particular subnet.

Catalyst 4840G software supports up to 255 VLANs per system. Because routing will take place, each VLAN is assumed to terminate at the load-balancing switch. Integrated routing and bridging (IRB) also is supported if this does happen. To configure IRB, see the “Integrated Routing and Bridging” section on page 9-3.

To configure VLANs, you define a subinterface at the interface, define a bridge group, and then map a VLAN to the subinterface.

To configure VLANs, see the “Using VLANs in SLB” section on page 4-6.

