CHAPTER 4

Configuring Interfaces

This chapter describes how to configure interfaces on Catalyst 4840G SLB switches, and how to configure VLAN encapsulation for bridging.

For further information about the commands used in this chapter, refer to the command reference publications in the Cisco IOS documentation set and to Appendix A, "Command Reference."

This chapter includes the following sections:

Note You are at Step 2 in the suggested process for configuring your Catalyst 4840G SLB switch. See the "Switch Configuration Steps" section on page 2-1.

In order for a Catalyst 4840G SLB switch to relay packets from one data link to another, you must define the characteristics of the interfaces through which the packets are received and sent. Interface characteristics include, but are not limited to, IP address, address of the port, data encapsulation method, and media type.

- Each IP physical and virtual interface must be configured with an IP address and an IP subnet mask.
- The virtual interfaces supported by Catalyst 4840G SLB switches can include subinterfaces.

  A subinterface allows a single physical interface to support multiple logical interfaces or networks—that is, several logical interfaces or networks can be associated with a single hardware interface.

About MAC Addresses and Port ID Names

Layer 3 interfaces use both a MAC address and an interface port ID. The switch uses these designators to route traffic. The MAC address, also referred to as the hardware address, is required for every port or device that connects to a network. Other devices in the network use MAC addresses to locate specific ports in the network and to create and update routing tables and data structures.

Tips To find the MAC address for a device, enter the **show interfaces** command.

The interface port ID designates the physical number of the SLB or Layer 3 interface. This is the name that you use to identify the interface when you configure it. The software uses interface port IDs to control activity within the Catalyst 4840G SLB switch and to display status information. Interface port IDs are not used by other devices in the network; they are specific to the individual load-balancing switch and its internal components and software.

Fast Ethernet interface ports are numbered from 1 through 40 and the two Gigabit Ethernet interface ports are numbered 41 and 42. For example, the first Fast Ethernet interface ID is fastethernet1 and the first Gigabit Ethernet interface ID is gigabitethernet41.

You can identify module ports by physically checking the port number on the back of the SLB switch. You can also use the **show** commands to display information about a specific interface, or all the interfaces, in the switch.

Interface Configuration Steps

The Catalyst 4840G SLB switch supports 40 Fast Ethernet and 2 Gigabit Ethernet interfaces. This section describes the general configuration steps that apply to all interfaces and provides some example configurations for both interface types.

The following general configuration steps apply to all interfaces. Follow these steps beginning in global configuration mode:

Step 1 Enter the **configure** EXEC command at the privileged EXEC prompt to enter the global configuration mode.

```
SLB-Switch> enable
SLB-Switch# configure terminal
SLB-Switch(config)#
```

Step 2 From global configuration mode, enter the **interface** command, followed by the interface type (for example, Fast Ethernet or Gigabit Ethernet) and its interface port ID.

For example, to configure the Gigabit Ethernet port 41, enter this command:

```
SLB-Switch(config)# interface gigabitethernet 41
```

Step 3 Follow each **interface** command with the interface configuration commands required for your particular interface.

The commands you enter define the protocols and applications that will run on the interface. The commands are collected and applied to the **interface** command until you enter another **interface** command, a command that is not an interface configuration command, or you enter **end** to return to privileged EXEC mode.

Step 4 After an interface is configured, check its status by entering the EXEC **show** commands.

```
SLB-Switch# show interfaces gigabitEthernet 41
GigabitEthernet41 is administratively down, line protocol is down
  Hardware is xpif_port, address is 0050.3e7e.f107 (bia 0050.3e7e.f107)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  Full-duplex, 1000Mb/s, 1000Base-SX, Auto-negotiation
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/475, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 watchdog, 0 multicast
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns(0/0/0)
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
SLB-Switch#
```

Configuring Gigabit Ethernet Interfaces

To configure an IP address and autonegotiation on a Gigabit Ethernet interface, perform this task starting in global configuration mode:

	Command	Purpose
Step 1	`SLB-Switch(config)# interface gigabitethernet` *interface_number* `SLB-Switch(config-if)#`	Enter Ethernet interface configuration mode to configure the Gigabit Ethernet interface.
Step 2	`SLB-Switch(config-if)# [no] negotiation auto`	Specify the negotiation mode. When you set negotiation mode to **auto**, the Gigabit Ethernet port attempts to negotiate the link (the duplex setting) with the partner port. When you set the Gigabit Ethernet interface to **no negotiation auto**, the port forces the link up no matter what the partner port setting is. This brings up the link with 1000 Mbps and full duplex only.
Step 3	`SLB-Switch(config-if)# ip address` *ip-address subnet-mask*	Specify the IP address and IP subnet mask to be assigned to the Gigabit Ethernet interface.
Step 4	`SLB-Switch(config-if)# exit` `SLB-Switch(config)#`	Return to global configuration mode. Repeat Steps 1 to 3 to configure the second Gigabit Ethernet interface.

	Command	Purpose
Step 5	`SLB-Switch(config)# end`	Return to privileged EXEC mode.
Step 6	`SLB-Switch# copy system:running-config nvram:startup-config`	Save your configuration changes to NVRAM.

This example shows how to configure a Gigabit Ethernet interface with autonegotiation and an IP address:

```
SLB-Switch(config)# interface gigabitethernet 41
SLB-Switch(config-if)# negotiation auto
SLB-Switch(config-if)# ip address 10.1.2.3 255.0.0.0
SLB-Switch(config-if)# exit
SLB-Switch(config)# end
SLB-Switch# copy system:running-config nvram:startup-config
```

Configuring Fast Ethernet Interfaces

Use the following procedure to assign an IP address to the Fast Ethernet 10BASE-T or 100BASE-T interface of your Catalyst 4840G SLB switch so that it can be recognized as a device on the Ethernet LAN. The Fast Ethernet interface supports 10-Mbps and 100-Mbps speeds with Cisco 10BASE-T and 100BASE-T routers, hubs, switches, and Catalyst 4840G SLB switches.

	Command	Purpose
Step 1	`SLB-Switch(config)# interface fastethernet` *interface_number* `SLB-Switch(config-if)#`	Enter Ethernet interface configuration mode to configure the Fast Ethernet interfaces.
Step 2	`SLB-Switch(config-if)# ip address` *ip-address subnet-mask*	Specify the IP address and IP subnet mask to be assigned to the Fast Ethernet interface.
Step 3	`SLB-Switch(config-if)# [no] speed [10 | 100 | auto]`	Configure the transmission speed for 10 or 100 Mbps, or for autonegotiation (the default). If you set the speed to **auto**, you enable autonegotiation, and the Catalyst 4840G SLB switch matches the speed of the partner node.
Step 4	`SLB-Switch(config-if)# [no] duplex [full | half | auto]`	Configure duplex mode. If you set duplex for **auto**, the Catalyst 4840G SLB switch matches the duplex setting of the partner node.
Step 5	`SLB-Switch(config-if)# end` `SLB-Switch#`	Return to privileged EXEC mode.
Step 6	`SLB-Switch# copy system:running-config nvram:startup-config`	Save your configuration changes to NVRAM.

This example shows how to configure a Fast Ethernet interface with an IP address and autonegotiated speed and duplex mode:

```
SLB-Switch(config)# interface fastethernet 1
SLB-Switch(config-if)# ip address 10.1.2.4 255.0.0.0
SLB-Switch(config-if)# speed auto
SLB-Switch(config-if)# duplex auto
SLB-Switch(config-if)# end
SLB-Switch# copy system:running-config nvram:startup-config
```

Monitoring Ethernet Interfaces

To verify the settings after you have configured Gigabit Ethernet or Ethernet 10/100 BASE-T operation, perform one of these tasks:

Command	Purpose
show interface gigabitethernet *interface_number*	Display the status and global parameters of the Gigabit Ethernet interface.
show interface fastethernet *interface_number*	Display the status and global parameters of the Fast Ethernet interface.

This example shows sample output from the **show interfaces gigabitethernet** command:

```
SLB-Switch# show interfaces gigabitethernet 41
GigabitEthernet41 is down, line protocol is down
  Hardware is xpif_port, address is 0030.40d6.4d07 (bia 0030.40d6.4d07)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  Full-duplex, 1000Mb/s, 1000Base-SX, Auto-negotiation
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 1d19h, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/475, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     4217 packets input, 1475664 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 watchdog, 4217 multicast
     0 input packets with dribble condition detected
     4397 packets output, 1465255 bytes, 0 underruns(0/0/0)
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
SLB-Switch#
```

This example shows sample output from the **show interface fastethernet** command:

```
SLB-Switch# show interfaces fastEthernet 1
FastEthernet1 is up, line protocol is up
  Hardware is epif_port, address is 0050.3e7e.f007 (bia 0050.3e7e.f007)
  Internet address is 172.20.52.9/27
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  Auto-duplex, Auto Speed, 100BaseTX
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 1000 bits/sec, 2 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     1200042 packets input, 90451466 bytes, 0 no buffer
     Received 21167 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 watchdog, 1177312 multicast
     0 input packets with dribble condition detected
```

```
     10938 packets output, 3651546 bytes, 0 underruns(0/0/0)
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
SLB-Switch#
```

Using VLANs in SLB

VLANs enable network managers to group users logically rather than by physical location. A VLAN is an emulation of a standard LAN that allows data transfer and communication to occur without the traditional restraints placed on the network. It can also be considered a broadcast domain set up within a switch. With VLANs, switches can support more than one subnet (or VLAN) on each switch, and give routers and switches the opportunity to support multiple subnets on a single physical link. A group of devices on a LAN are configured so that they communicate as if they were attached to the same LAN segment, when they are actually located on different segments. Layer 3 switching supports up to 255 VLANs per system.

VLANs enable efficient traffic separation and provide excellent bandwidth utilization. VLANs logically segment the physical LAN structure into different subnetworks so that packets are switched only between ports within the same VLAN.

Layer 3 switching software supports a port-based VLAN on a trunk port, which is a port that carries the traffic of multiple VLANs. Each frame transmitted on a trunk link is tagged as belonging to only one VLAN.

Layer 3 switching software supports VLAN frame encapsulation through the Inter-Switch Link (ISL) protocol and the 802.1Q standard.

Figure 4-1 shows a network topology where two VLANs span a Catalyst 5500 switch and Catalyst 4840G SLB switch. Both VLANs in this topology are bridged using the Inter-Switch Link (ISL) protocol.

Figure 4-1 VLAN Spanning Devices in a Network

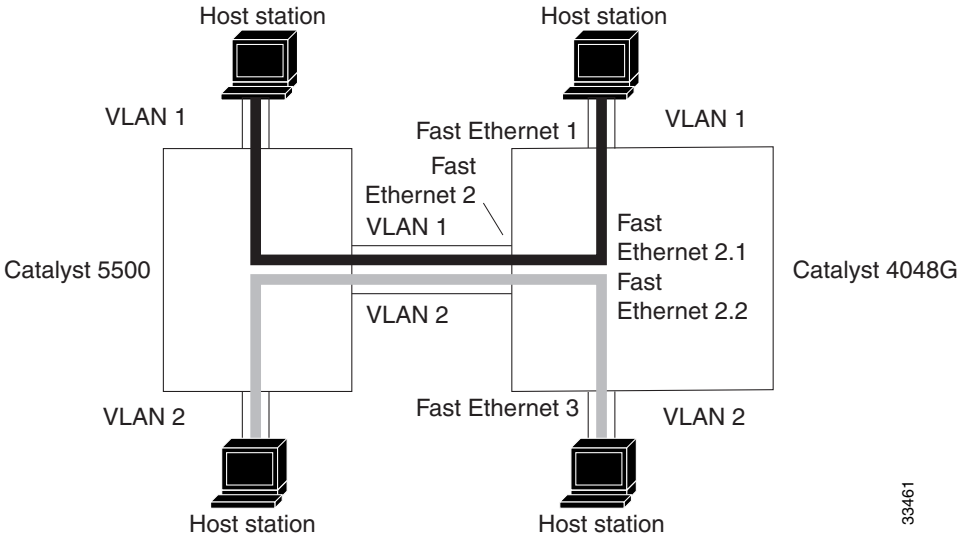

Note The four adjacent ports (such as ports 1 through 4 or ports 5 through 8) on the 10/100 interfaces must all use the same VLAN encapsulation; that is, either 802.1Q and native, or ISL and native.

Configuring ISL VLAN Encapsulation

ISL is a Cisco protocol for interconnecting multiple switches and maintaining VLAN information as traffic travels between switches.

When configuring ISL with IP, you cannot configure IP addresses on a subinterface unless the VLANs are already configured (that is, you must have already entered the **encapsulation isl** or **encapsulation dot1q** command).

The maximum VLAN bridge group values are as follows:

- Maximum number of bridge groups: 64
- Maximum number of interfaces per bridge group: 128
- Maximum number of subinterfaces per system: 255

The VLAN configuration example shown in Figure 4-2 depicts the following:

- Fast Ethernet port 1 and subinterface 2.1 on the Catalyst 4840G SLB switch are in bridge group 1. They are part of VLAN 50, which uses ISL encapsulation.
- Fast Ethernet port 3 and subinterface 2.2 are in bridge group 2. They are part of VLAN 100, which uses ISL encapsulation.
- Fast Ethernet port 1 is configured as an ISL trunk.

Figure 4-2 Example of an ISL VLAN Bridging Configuration

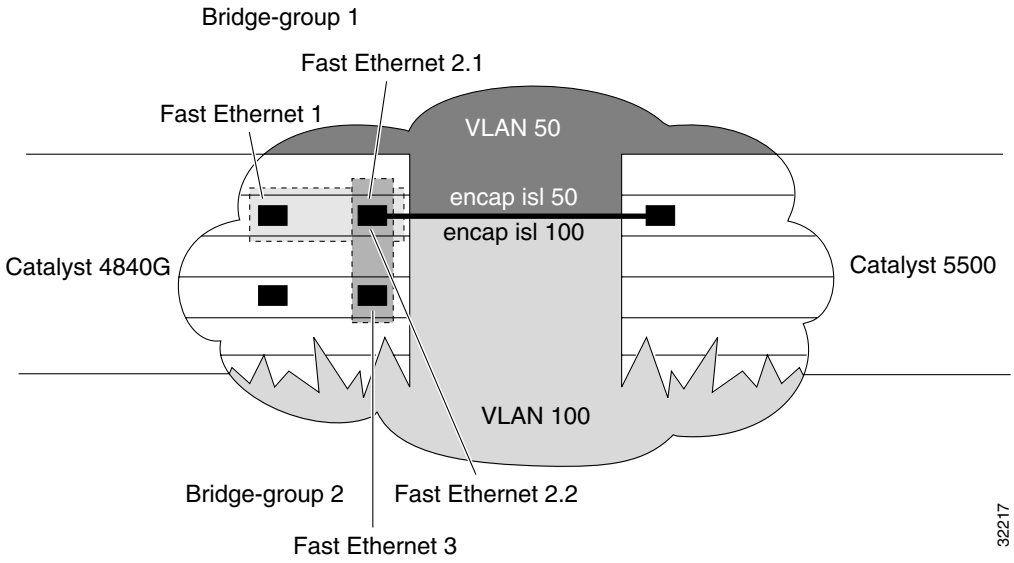

To configure the Layer 3 VLANs shown in Figure 4-2, perform this task:

	Command	Purpose
Step 1	SLB-Switch(config)# **interface fastethernet** *interface_number.subinterface* SLB-Switch(config-subif)#	Enter subinterface configuration mode.
Step 2	SLB-Switch(config-subif)# **encapsulation isl** *vlan-id*	Specify ISL encapsulation for the Ethernet frames sent from this subinterface with a header that maintains the specified VLAN ID between network nodes.
Step 3	SLB-Switch(config-subif)# **bridge-group** *bridge-group*	Assign the subinterface a bridge group number. **Note** When you are configuring VLAN routing, skip this step.
Step 4	SLB-Switch(config-subif)# **interface fastethernet** *interface_number* SLB-Switch(config-if)#	Enter interface configuration mode to configure the Fast Ethernet main interface.
Step 5	SLB-Switch(config-if)# **bridge-group** *bridge-group*	Assign the main interface to the bridge group.
Step 6	SLB-Switch(config-if)# **exit** SLB-Switch(config)#	Return to global configuration mode.
Step 7	SLB-Switch(config)# **bridge** *bridge-group* **protocol ieee**	Specify that the bridge group will use the IEEE Ethernet Spanning Tree Protocol.

This example shows how to configure the interfaces for VLAN bridging with ISL encapsulation shown in Figure 4-2:

```
SLB-Switch(config)# interface fastethernet 1.1
SLB-Switch(config-subif)# encap isl 50
SLB-Switch(config-subif)# bridge-group 1
SLB-Switch(config-subif)# interface fastethernet 0
SLB-Switch(config-if)# bridge-group 1
SLB-Switch(config-if)# exit
SLB-Switch(config)# bridge 1 protocol ieee
SLB-Switch(config)# interface fastethernet 1.2
SLB-Switch(config-subif)# encap isl 100
SLB-Switch(config-subif)# bridge-group 2
SLB-Switch(config-subif)# interface fastethernet 3
SLB-Switch(config-subif)# bridge-group 2
SLB-Switch(config-subif)# exit
SLB-Switch(config)# bridge 2 protocol ieee
SLB-Switch(config)# exit
SLB-Switch# copy system:running-config nvram:startup-config
```

For a complete configuration example for VLANs with ISL encapsulation, see the "Example ISL VLAN and BVI with GEC Configuration" section on page 4-11.

To monitor the VLANs once they are configured, use the commands described in the "Monitoring VLAN Operation" section on page 4-10.

Configuring 802.1Q VLAN Encapsulation

The IEEE 802.1Q standard provides a method for secure bridging of data across a shared backbone. 802.1Q VLAN encapsulation uses an internal, or one level, packet-tagging scheme to multiplex VLANs across a single physical link, while maintaining strict adherence to the individual VLAN domains.

On an 802.1Q trunk port, all transmitted and received frames are tagged except for those on the one VLAN configured as the port VLAN identifier (PVID) or native VLAN for the port. Frames on the native VLAN are always transmitted untagged and are normally received untagged.

The VLAN configuration shown in Figure 4-3 depicts the following:

- Fast Ethernet ports 1 and subinterface 2.1 on the Catalyst 4840G SLB switch are in bridge group 1. They are part of native VLAN 1, which uses 802.1Q encapsulation.
- Fast Ethernet port 3 and subinterface 2.2 are in bridge group 2. They are part of VLAN 100, which uses 802.1Q encapsulation.
- Fast Ethernet port 1 is configured as an 802.1Q trunk.

Figure 4-3 Example of Bridging Between Native and Non-Native 802.1Q VLANs

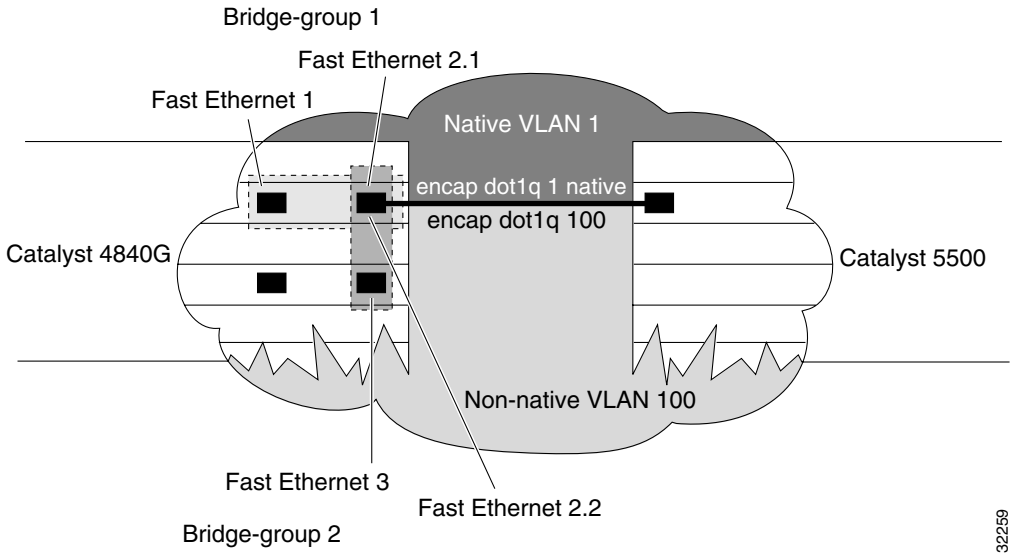

To configure the bridging between native VLAN 1 and non-native VLAN 100 depicted in Figure 4-3, perform this task:

	Command	Purpose
Step 1	`SLB-Switch(config)# interface fastethernet` *interface_number.subinterface*	Enter subinterface configuration mode.
Step 2	`SLB-Switch(config-subif)# encap dot1q` *vlan-id* `native`	Specify 802.1Q encapsulation for Ethernet frames sent from the subinterface with a header that maintains the specified native VLAN ID between network nodes.

	Command	Purpose
Step 3	`SLB-Switch(config-subif)# bridge-group` *bridge-group*	Assign the subinterface a bridge group number. **Note** When you are configuring VLAN routing, skip this step.
Step 4	`SLB-Switch(config-subif)# interface fastethernet` *interface_number*	Enter interface configuration mode to configure the Fast Ethernet main interface.
Step 5	`SLB-Switch(config-if)# bridge-group` *bridge-group*	Assign the main interface to the bridge group.
Step 6	`SLB-Switch(config-if)# exit`	Return to global configuration mode.
Step 7	`SLB-Switch(config)# bridge` *bridge-group* `protocol ieee`	Specify that the bridge group will use the IEEE Spanning Tree Protocol.

This example shows how to configure the bridging between native and non-native 802.1Q VLANS shown in Figure 4-3:

```
SLB-Switch(config)# interface fastethernet 2.1
SLB-Switch(config-subif)# encap dot1q 1 native
SLB-Switch(config-subif)# bridge-group 1
SLB-Switch(config-subif)# interface fastethernet 1
SLB-Switch(config-if)# bridge-group 1
SLB-Switch(config-if)# exit
SLB-Switch(config)# bridge 1 protocol ieee
SLB-Switch(config)# interface fastethernet 2.2
SLB-Switch(config-subif)# encap dot1q 100
SLB-Switch(config-subif)# bridge-group 2
SLB-Switch(config-subif)# interface fastethernet 3
SLB-Switch(config-subif)# bridge-group 2
SLB-Switch(config-subif)# exit
SLB-Switch(config)# bridge 2 protocol ieee
SLB-Switch(config)# exit
SLB-Switch# copy system:running-config nvram:startup-config
```

Monitoring VLAN Operation

After the VLANs are configured on the SLB switch, you can monitor their operation by performing one of these tasks:

Command	Purpose
show vlan *vlan-id*	Display information on all configured VLANs or on a specific VLAN (by VLAN ID number).
clear vlan *vlan-id*	Clear the counters for all VLANs, when the VLAN ID is not specified.
debug vlan packet	Display contents of the packets sent to and exiting from the route processor.

Example ISL VLAN and BVI with GEC Configuration

This example configuration focuses on both ISL and VLANs, as well as integrated routing and bridging (IRB) using a bridge-group virtual interface (BVI) over Gigabit EtherChannel. The Cisco proprietary ISL allows any Fast Ethernet port to be configured as a trunk. The Spanning Tree Protocol detects and breaks loops on all the VLANs carried across the trunk.

```
!
ip subnet-zero
no ip domain-lookup
ip name-server 171.69.2.132
ip name-server 198.92.30.32
bridge irb
!
interface FastEthernet1
    no ip address
    no ip directed-broadcast
    no keepalive
!
interface FastEthernet1.128
    ip address 172.68.16.10 255.255.255.0
    ip helper-address 172.68.16.15
    no ip redirects
    no ip directed-broadcast
    encapsulation isl 128
!
interface FastEthernet1.199
    ip address 172.68.17.15 255.255.255.0
    ip helper-address 172.68.16.16
    ip helper-address 172.68.16.17
    ip helper-address 172.68.16.18
    no ip redirects
    no ip directed-broadcast
    encapsulation isl 199
!
interface FastEthernet1.201
    ip address 172.68.18.10 255.255.255.0
    ip helper-address 172.68.16.16
    ip helper-address 172.68.16.17
    ip helper-address 172.68.16.18
    no ip redirects
    no ip directed-broadcast
    encapsulation isl 201
!
interface FastEthernet2
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet3
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet4
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
```

```
interface FastEthernet5
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet6
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet7
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet8
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet9
    ip address 172.68.19.10 255.255.255.0
    ip helper-address 172.68.16.16
    ip helper-address 172.68.16.17
    ip helper-address 172.68.16.18
    no ip redirects
    no ip directed-broadcast
    ip sdr listen
    no keepalive
!
interface FastEthernet10
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
interface FastEthernet11
    no ip address
    no ip directed-broadcast
    no keepalive
    shutdown
!
.
(Information Deleted)
.
interface GigabitEthernet41
!
interface GigabitEthernet42
 ip address 172.68.1.1 255.255.255.0
 no ip directed-broadcast
!
interface BVI1
 ip address 171.201.1.2 255.255.255.0
 no ip directed-broadcast
 no ip route-cache cef
!
router eigrp 170
 network 171.200.0.0
 network 171.201.0.0
 network 172.68.0.0
```

```
 network 172.69.0.0
 no auto-summary
!
router bgp 180
 network 172.68.1.0
 network 172.69.1.0
 no auto-summary
!
ip classless
!
bridge 1 protocol ieee
bridge 1 route ip
!
ip http server
!
line con 0
line aux 0
line vty 0 4
    login
!
ntp clock-period 17181168
ntp update-calendar
ntp server 171.71.150.52
ntp server 171.69.4.143
ntp server 171.69.5.10
end
```