



CHAPTER 23

Configuring Gigabit EtherChannel Features

On a Cisco 10000 Series router, a Gigabit EtherChannel (GEC) is a specialized interface type comprising aggregated Gigabit Ethernet links. A GEC bundle is synonymous with port channel and can have a minimum of one or a maximum of 8 active links. The bandwidth of the GEC interface is the aggregate of all the physical member links comprising the GEC bundle.



Note

Cisco IOS Release 12.2(31)SB supports a maximum of 4 member links per GEC bundle. In Cisco IOS Release 12.2(15)BX, the maximum number of links per GEC bundle has been increased from 4 to 8.

The Gigabit EtherChannel can be deployed in two ways on the Cisco 10000 Series router:

- *Core facing* or network facing deployment is an uplink EtherChannel that connects the Cisco 10000 Series router to the service provider. This setup has multiple physical links bundled per GEC interface and allows:
 - Load balancing across all the active interfaces.
 - Combination of different Gigabit Ethernet (GE) ports (both shared port adaptors and line cards)
- *Access facing* or subscriber facing deployment connects the Cisco 10000 Series router to the subscriber edge. This setup typically has only one active member link on a GEC bundle interface. The remaining links in the GEC bundle serve as passive links. Traffic is sent only through the active member link, while the passive link is used as a backup when the active member link fails. This arrangement provides link redundancy with no loss of Point-to-Point Protocol over Ethernet (PPPoE) sessions during link failover.
 - Load balancing is not applicable when there is only one active link in a GEC bundle.
 - Queuing action is allowed only on the GEC bundle and not on member links.



Note

A GEC bundle can include a combination of active and passive links. In an M:N mode, 'M' denotes active links and 'N' denotes passive links. In a 1:N mode only one link is active per GEC bundle and 'N' denotes passive links. Passive links operate as backup links but do not transfer any network traffic.

This chapter describes Gigabit EtherChannel (GEC) enhancements implemented on the Cisco 10000 Series routers and includes the following topics:

- [Feature History for Gigabit EtherChannel, page 23-2](#)
- [Prerequisites for Gigabit EtherChannel Configuration, page 23-3](#)
- [Restrictions for Gigabit EtherChannel Configuration, page 23-3](#)

■ Feature History for Gigabit EtherChannel

- Configuring QoS Service Policies on GEC Interfaces, page 23-3
- Configuring Policy Based Routing Support on a GEC Bundle, page 23-7
- Configuring IEEE 802.1Q and QinQ Support on GEC Bundle, page 23-7
- Configuring MVPN Support on GEC Bundle, page 23-9
- Configuring PPPoX Support on a GEC Bundle, page 23-9
- Configuring High Availability Support on GEC Bundle, page 23-11
- Configuring 8 Member Links per GEC Bundle, page 23-11
- Configuring VLAN-Based Load Balancing, page 23-12

For more information on Gigabit EtherChannels (GEC) see, the *Link Aggregation Control Protocol (LACP) (802.3ad) for Gigabit Interfaces* feature guide at:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sb/newft/122sb31/10gigeth.htm>

Feature History for Gigabit EtherChannel

Cisco IOS Release	Description	Required PRE
12.2(31)SB	This feature was supported only on native line cards. • QoS policies supported only on GEC member links.	PRE2 and PRE3
12.2(33)SB	This feature is supported on native line cards ¹ and the SPA Interface Processor (SIP) and Shared Port Adapters (SPA) ² on the Cisco 10000 Series router.	PRE3 and PRE4
12.2(33)SB	The following Gigabit EtherChannel enhancements were added on the Cisco 10000 Series router: • QoS Service Policies on GEC Bundle ³ • PPPoE hitless ⁴ switchover support with Link Aggregation Control Protocol (LACP) (802.3ad) port channel • PBR Support for GEC Bundle • IEEE 802.1Q and QinQ Support on GEC Bundle • Multicast VPN Support on GEC Bundle • PPPoX Support for GEC Bundle (includes PPPoEoE, PPPoEoQinQ, PPPoVLAN) • High availability for SSO, NSF, ISSU, and NSR • 8 Member Links per GEC bundle	PRE2, PRE3, and PRE4
12.2(33)XNE	VLAN-based Load Balancing feature for GEC interface was introduced to Cisco 10000 series router.	PRE3 and PRE4

1. 1 Port GE – Half-height and 1 Port GE – Full-height.

2. 2 Port GE – Half-height and 5 Port GE – Half-height.

3. Queuing actions still require QoS application on each member link.

4. A hitless switchover implies that no PPPoX sessions are dropped when an active member link fails and the passive or backup link takes over as the active link.

Prerequisites for Gigabit EtherChannel Configuration

The following are the prerequisites for configuring GEC bundles:

- Create a GEC bundle interface before adding GE links to the GEC bundle using the **channel-group** command.
- Add GE links to the GEC bundle and configure all the links identically.

Restrictions for Gigabit EtherChannel Configuration

The following are general restrictions applicable to GEC bundles:

- Bidirectional Forwarding Detection Protocol (BFD) is not supported on GEC bundles.
 - A dot1Q/QinQ subinterface created on a GEC bundle requires one VCCI on all the member links. The GEC bundle also uses one VCCI.
 - Intelligent Service Gateway (ISG) IP sessions are not supported on GEC bundles.
 - Gigabit EtherChannels are not supported on a 10 GE shared port adaptor (SPA).
 - On core facing deployments, queuing actions are only supported when they are applied on each individual member link.
 - In 1:N deployment mode:
 - Combination of any SPA based Gigabit Ethernet ports with either a full height GE line card or half height GE line card is not supported.
 - A maximum of 4 Gigabit Ethernet SPAs form a GEC bundle. Each SPA interface must have the same bay number and port number, assuming the representation is GigabitEthernetSlotNumber/BayNumber/Port-Number.
- For example, in the case of SPAs, Gi1/2/1 can be bundled with Gi5/2/1, but Gi1/2/1 cannot be bundled with Gi1/0/1.
- Member link counters are not updated.
- We do not recommend the deletion of the GEC bundle main interface before removing the member links, using the **no channel-group** command. Similarly, before you delete the GEC bundle main interface, remove the subinterfaces.

Configuring QoS Service Policies on GEC Interfaces

The QoS support feature on Gigabit EtherChannel allows service policies to be applied on GEC traffic. Support for QoS can be applied differently to GEC bundles or main interfaces of member links and GEC bundle subinterfaces.

- Input QoS can directly be applied on GEC bundle subinterfaces or main interfaces. Input QoS can be applied on member links using the vlan-group QoS feature. For details, see the “[Configuration Example for Using the VLAN Group Feature to Apply QoS on Member Links](#)” section on page 23-5 and the “[Configuration Example for Applying QoS on GEC Bundle Subinterfaces](#)” section on page 23-6.

Configuring QoS Service Policies on GEC Interfaces

- Output QoS can directly be applied on GEC bundle subinterfaces similar to the GEC main interfaces. Alternatively, output QoS can be applied on member links using the `vlan-group` QoS feature. The service policy with **match-vlan class-maps** is applied on the member link main interface.

The application of QoS depends on the deployment mode of the GEC bundle interface as described in the [Table 23-1](#):

Table 23-1 Service Policies Applied on GEC Bundles

QoS for GEC	M:N Deployment	1:N Deployment
Input QoS for GEC	<p>Service policy can be applied for both GEC bundle interface and member links. If applied on a:</p> <ul style="list-style-type: none"> • GEC bundle, the aggregate ingress traffic on GEC bundle is subject to this service-policy. • Member link, ingress traffic on a particular member link is subject to the service-policy applied on the member link. 	Service policy can be applied only on the GEC bundle interface.
Output QoS for GEC	<p>Service-policy <i>without</i> queuing actions can be applied either on GEC bundle interface or on the member links. If applied on a:</p> <ul style="list-style-type: none"> • GEC bundle, the aggregate egress traffic on GEC bundle is subject to this service-policy. • Member link, egress traffic on a particular member link is subject to the service-policy applied on the member link. <p>Service-policies <i>with</i> queuing actions can be applied only on member links.</p>	Service policies <i>with or without</i> queuing actions can be applied only on the GEC bundle interface.
Input QoS for GEC subinterface	<p>Input QoS applied on GEC bundle interface and on member main-interfaces. If the service-policy is applied on a:</p> <ul style="list-style-type: none"> • GEC bundle subinterface, the aggregate ingress traffic on the GEC bundle subinterface is subject to this service-policy. • GEC member main-interface using <code>vlan-group</code> feature, the ingress traffic on that member link with the <code>vlan-ids</code> specified in the <code>vlan-group</code> service-policy is subject to the corresponding actions as specified in the service-policy. 	Input QoS can only be applied on a bundle subinterface. The aggregate bundle subinterface traffic is subject to this service-policy.
Output QoS for GEC subinterface	<p>Service-policies <i>without</i> queuing actions can be applied either on the GEC bundle subinterface or on the member main-interface. If the service-policy is applied on a:</p> <ul style="list-style-type: none"> • GEC bundle subinterface, the aggregate egress traffic on that GEC bundle subinterface is subject to this service-policy. • GEC member main-interface using <code>vlan-group</code> feature, the egress traffic on that member link with the <code>vlan-ids</code> specified in the <code>vlan-group</code> service-policy is subject to the corresponding actions as specified in the service-policy. <p>Service policies <i>with</i> queuing actions can only be applied on member links. The egress traffic on that member link with the <code>vlan-ids</code> specified in the <code>vlan-group</code> service-policy is subject to the corresponding actions as specified in the service-policy.</p>	Output QoS <i>with or without</i> queuing actions can only be applied on a GEC bundle subinterface.

Restrictions for QoS Service Policies on GEC Bundles

The following restrictions are applicable to QoS service policies applied on GEC bundle interfaces and subinterfaces:

- Both ingress and egress service-policy without any queuing actions can only be applied on member links for M:N deployment, and are restricted for 1:N deployment.
- Egress service-policy with queuing action can only be applied on:
 - Member interfaces for an M:N GEC deployment
 - Bundle interface for a 1:N GEC deployment.
- Restriction for application of QoS on VLAN groups:
 - A VLAN group restricts the application of QoS to a maximum of 255 individual dot1Q (VLAN) subinterfaces. 255 denotes the number of user classes that can be used in VLAN group parent policy.
 - Matching a VLAN group user class on QinQ subinterfaces is not supported using VLAN group policy. Each VLAN group user class at parent hierarchy can only match a set of dot1Q subinterfaces.
 - VLAN group QoS policy at member link is used only when dot1Q subinterfaces are defined on the GEC bundles and not when QinQ subinterfaces are defined on the GEC bundle.
- Input Quality of Service (QoS) on member links is not supported for QinQ subinterfaces.
- The classification criteria of **match input-interface port-channel** is not supported. Instead, packets are classified by matching them with member-links.

Configuration Examples

This section provides the following configuration examples:

- [Configuration Example for Using the VLAN Group Feature to Apply QoS on Member Links, page 23-5](#)
- [Configuration Example for Applying QoS on GEC Bundle Subinterfaces, page 23-6](#)

Configuration Example for Using the VLAN Group Feature to Apply QoS on Member Links

Step 1 Consider a GEC bundle interface with two member links Gig3/0/0 and Gig4/0/0.

Step 2 Assume subinterfaces exist on the GEC bundle interface having the following configurations:

```
Port-channel 1.1
  Encapsulation dot1q 2
Port-channel 1.2
  Encapsulation dot1q 3
Port-channel 1.3
  Encapsulation dot1q 4
```

Assume that the following configurations need to be performed on each member link

- Police ingress traffic for subinterface port-channel 1.1 at 100 mbps
- Police ingress traffic for subinterface port-channel 1.2 at 150 mbps
- Shape egress traffic for subinterface port-channel 1.2 at 50 mbps

- Shape egress traffic for subinterfaces port-channel 1.1 and port-channel 1.3 together at 150 mbps

Step 3 Create **match-vlan class-maps** as follows:

```
Class-map match-any vlan_2
  Match vlan 2
Class-map match-any vlan_3
  Match vlan 3
Class-map match-any vlan_4
  Match vlan 4
Class-map match-any vlan_2_4
  Match vlan 2 4
```

Step 4 Create **policy-maps** as follows:

```
Policy-map mega_ingress
  Class vlan_2
    Police 100 mbps
  Class vlan_3
    Police 150 mbps
Policy-map mega_egress
  Class vlan_3
    Shape 50 mpbs
  Class vlan_2_4
    Shape 150 mbps
```

Step 5 Apply this policy on the GEC member links.

```
Interface Gig3/0/0
  Service-policy input mega_ingress
  Service-policy output mega_egress
Interface Gig4/0/0
  Service-policy input mega_ingress
  Service-policy output mega_egress
```

Configuration Example for Applying QoS on GEC Bundle Subinterfaces

[Example 23-1](#) shows how QoS is applied on GEC bundle subinterfaces:

Example 23-1 Applying QoS on GEC Bundle Subinterfaces

```
Class-map match-any dscp_20_30
  Match dscp 20 30
Class-map match-any dscp_40
  Match dscp 40

Policy-map police_dscp
  Class dscp_20_30
    Police 50 3000 3000 conform-action transmit exceed-action drop
    Set ip dscp af22
  Class dscp_40
    Police 10 3000 3000 conform-action transmit exceed-action drop

Policy-map customer_A
  Class class-default
    Police 100 mbps
    service-policy police_dscp

Policy-map customer_B
  Class class-default
    Police 150 mbps
```

```

Service-policy police_dscp

Interface Port-channel 1.1
  Service-policy input customer_A

Interface Port-channel 1.2
  Service-policy input customer_B

```

Configuring Policy Based Routing Support on a GEC Bundle

Cisco Policy Based Routing (PBR) provides a flexible mechanism for network administrators to customize the operation of the routing table and the flow of traffic within their networks.

Load balancing is performed on packets that pass through PBR. If a PBR clause is applied to outbound packets, and the clause results in the selection of an EtherChannel egress interface, then a new hash is generated based on the new address information. This hash is used to select an appropriate member link as the packet's final destination.

Policy based routing is supported on Gigabit EtherChannel. You can configure PBR directly on a GEC bundle interface.

Restriction for Configuring PBR Support on a GEC Bundle

- Use of the **set interface** command is restricted in the set clause for PBR, on GEC bundle interfaces. Only the IP address for the next hop can be specified.

Configuring IEEE 802.1Q and QinQ Support on GEC Bundle

Support for both dot1Q and QinQ subinterfaces is available for GEC bundle interfaces. Configuring subinterface on a GEC bundle interface is similar to a normal Gigabit Ethernet interface configuration. When a subinterface is configured on a GEC bundle interface, the GEC *Toaster* client creates subinterface instances in the Toaster for all the GEC member links. The subinterface instances created internally on GEC member links are hidden and not available to the user for applying configurations.



Note

The *toaster* is designed to process IP packets at very high rates using existing forwarding algorithms, though it may also be programmed to perform other tasks and protocols.

Prerequisites for Configuring IEEE 802.1Q and QinQ Support

- Create a GEC bundle main interface before creating subinterfaces.

Restrictions for Configuring IEEE 802.1Q and QinQ Support on GEC Bundle

- A dot1Q/QinQ subinterface created on GEC bundle requires a VCCI on all the member-links. For example, a GEC bundle interface with 8 member-links uses 9 (1+8) VCIs for each dot1Q/QinQ subinterface created on the GEC bundle.

- Ingress packet accounting for QinQ subinterfaces is carried out at the bundle level. Accounting of these ingress packets per member link is not supported.

Configuration Tasks for IEEE 802.1Q and QinQ on Subinterfaces

To create a GEC bundle subinterface and configure a dot1Q/QinQ encapsulation, enter the following commands beginning from the global configuration mode in the following table:

	Command	Purpose
Step 1	router(config)# interface port-channel <i>number</i>	Creates a GEC bundle.
Step 2	router(config)# interface port-channel <i>subinterface</i>	Creates a GEC bundle subinterface and enters the subinterface mode.
Step 3	router(config-subif)# encapsulation dot1Q <i>vlan-id</i>	Enables IEEE 802.1Q encapsulation on a specified subinterface in a VLAN.
Step 4	router(config-subif)# encapsulation dot1Q <i>vlan-id second-dot1q inner vlan-id</i>	Enables 802.1Q encapsulation on a specified subinterface in an inner VLAN.
Step 5	router# show running-config interface port-channel <i>subinterface</i>	Displays the current configuration for the GEC bundle subinterface.
Step 6	router# show interface port-channel <i>subinterface</i>	Displays status, traffic data, and configuration information about the subinterface you specify.

Configuration Examples

[Example 23-2](#) and [Example 23-3](#) show the encapsulation configuration details:

Example 23-2 show interface Command for the GEC Bundle Subinterface

```
router# show interface port-channel 1.1
Port-channel1.1 is up, line protocol is up
  Hardware is GEChannel, address is 0004.9b3e.101a (bia 0004.9b3e.1000)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID  20.
  ARP type: ARPA, ARP Timeout 04:00:00
  Last clearing of "show interface" counters never
```

Example 23-3 show running-config Command for the GEC Bundle Subinterface

```
router# show running-config interface port-channel 1.1
Building configuration...

Current configuration : 134 bytes
!
interface Port-channel1.1
  encapsulation dot1Q 20 second-dot1q 200
  ip address 3.0.0.1 255.255.255.0
  end
```

Configuring MVPN Support on GEC Bundle

The Multicast VPN (MVPN) feature allows a service provider to configure and support multicast traffic within a Virtual Private Network (VPN) environment. MVPN also supports routing and forwarding of multicast packets for each individual VPN routing and forwarding (VRF) instance, and it also provides a mechanism to transport VPN multicast packets across the service provider backbone.

On the Cisco 10000 Series router, when we use GEC as a core facing link (from the provider edge to the provider) the MVPN packet sent on the GEC interface has the IP header encapsulated inside a GRE Header or the Tunnel Header. The hash function is calculated based on the tunnel header's source and destination IP address and the original IP header's (Inner IP header) source and destination address. Load balancing is used on outbound GEC packets to find out the member link on which the packet is sent.

For more information on MVPN, see the “IP Multicast VPN” section in the *Multicast VPN—IP Multicast Support for MPLS VPN* guide at:

http://www.cisco.com/en/US/docs/ios/12_2sb/feature/guide/sbb_mvpn.html#wp1040907

Configuration Tasks and Examples

For configuration information and examples, see the “How to Configure Multicast VPN—IP Multicast Support for MPLS VPNs” section in the *How to Configure Multicast VPN—IP Multicast Support for MPLS VPNs* at:

http://www.cisco.com/en/US/docs/ios/12_2sb/feature/guide/sbb_mvpn.html#wp1041284

Configuring PPPoX Support on a GEC Bundle

PPPoE, PPPoEoQinQ, PPPoEoVLAN sessions are supported only in 1:N GEC mode and are provisioned on the GEC bundle interface. The complete session traffic is directed towards the active member link. Hence, when the active link goes down, the session traffic is directed towards the passive member link, which then becomes the active link.

PPPoX sessions on GEC work similar to that of a normal Gigabit Ethernet interface and QoS policy inheritance is similar to that of a normal Gigabit Ethernet interface.

Restrictions for Configuring PPPoX Support for GEC Bundle

- Support for PPPoX sessions is allowed only in a 1:N mode, where there is only one active GEC link.
- At any point of time the bandwidth of the 1:N GEC bundle will be 1 Gbps.



Note Multiple passive links can be added, but only one active link is supported for PPPoE.

For more information on PPPoEoQinQ support for subinterfaces, see PPPoE - QinQ Support feature guide at:

http://www.cisco.com/en/US/docs/ios/12_3t/12_3t7/feature/guide/gt_qinq.html

Configuration Tasks

To enable PPPoE session creation on a GEC bundle, enter the following commands:

	Command	Purpose
Step 1	router(config)# interface port-channel <i>number</i>	Creates a GEC bundle.
Step 2	router(config)# lacp max-bundle <i>1-8</i>	Sets the maximum number of active links per GEC bundle. For PPPoE sessions maximum number of active links is one.
Step 3	router(config)# lacp fast-switchover	Retains PPPoX sessions incase of member link switchover.
Step 4	router(config)# interface port-channel <i>subinterface</i>	Creates a GEC bundle subinterface and enters the subinterface mode.
Step 5	router(config-subif)# encapsulation dot1Q <i>vlan-id</i>	Enables IEEE 802.1Q encapsulation of traffic on a specified subinterface in a VLAN. Specify the VLAN identifier.
Step 6	router(config-subif)# pppoe enable group global	Enables PPPoE session on the GEC bundle subinterface. global is the default group used when a group name is not specified.
Step 7	router(config-subif)# end	Exits to the global configuration mode.

For more information on PPPoE over Ethernet, see the *Cisco 10000 Series Router Software Configuration Guide* at:

<http://www.cisco.com/en/US/docs/routers/10000/10008/configuration/guides/broadband/vlan.html>

Configuration Examples

Example 23-4 shows how to enable a PPPoE session on a GEC bundle:

Example 23-4 Enabling a PPPoE Session

```
interface Port-channel32
no ip address
no negotiation auto
lacp max-bundle 1
lacp fast-switchover
!
interface Port-channel32.1
encapsulation dot1Q 10
pppoe enable group bba_group_1
!
interface Port-channel32.2
encapsulation dot1Q 20
pppoe enable group bba_group_1
!
```

Configuring High Availability Support on GEC Bundle

The following high availability features are supported on GEC bundle interfaces, on the Cisco 10000 Series router.

- Stateful Switchover (SSO)
- In Service Software Upgrade (ISSU))
- Nonstop Forwarding (NSF)
- Nonstop Routing (NSR)

The EtherChannel and the IEEE 802.3ad LACP protocol are SSO and ISSU aware. This feature makes the GEC bundle interface available after a PRE switchover, in the event of a catastrophic failure.

For more information on NSF, see *Cisco Nonstop Forwarding* whitepaper at:

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsnsf20s.html

For information on SSO see, *Stateful Switchover* feature guide at:

http://www.cisco.com/en/us/docs/ios/12_2s/feature/guide/fssso20s.html

For more information on ISSU, see *Cisco IOS In Service Software Upgrade and Enhanced Fast Software Upgrade Process* feature guide, at:

http://www.cisco.com/en/US/docs/ios/12_2sb/feature/guide/sb_issu.html

Configuring 8 Member Links per GEC Bundle

A maximum of 8 configured member links per GEC bundle and 64 port channels are supported on the Cisco 10000 Series router. The number of member links per GEC bundle has been increased from 4 to 8 in the Cisco IOS Release 12.2(15)BX.

Configuration Tasks

The following table lists the configuration commands used to configure the maximum and minimum links on a GEC bundle on a Cisco 10000 Series router.

Command	Purpose
router(config-if)# lacp max-bundle 1-8	Assigns the maximum number of links per bundle. The default value is 8.
router(config-if)# lacp min-link 1-8	Sets the minimum number of active links required before declaring the port channel interface down. The default value is 1.
router(config-if)# lacp fast-switchover	Reduces the switchover time from 2 seconds to 10ms. By default lacp fast-switchover configuration is not enabled.

For more information on configuring member links, see the *Link Aggregation Control Protocol (LACP) (802.3ad) for Gigabit Interfaces* feature guide at:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sb/newft/122sb31/10gigeth.htm>

Configuring VLAN-Based Load Balancing

For more information on how to aggregate multiple Ethernet links into one logical channel, see *IEEE 802.3ad Link Bundling* feature guide at:
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122sb/newft/122sb31/sbcelacp.htm#wp1053782>

Configuring VLAN-Based Load Balancing

In Cisco IOS Release 12.2(33)XNE, there is support for VLAN-based load balancing for the GEC interface on the Cisco 10000 series routers. The user can enable manual VLAN load balancing and select the member-links on which a particular VLAN traffic is to be forwarded.

The VLAN load balancing feature can map a VLAN sub-interface to a member-link called a primary member-link. The egress traffic for the VLAN sub-interface is then transmitted through that primary member-link. The feature also allows the user to specify a standby member-link called a secondary member-link for a VLAN sub-interface. The secondary member-link is used if the primary member-link goes down.

Table 23-2 shows the active and standby links for different primary and secondary states.

Table 23-2 Active and standby links for different primary and secondary states

Primary	Secondary	Active	Standby
Up	Up	Primary	Secondary
Up	Down	Primary	Secondary
Down	Up	Secondary	Primary
Down	Down	Primary	Secondary

All packets forwarded over a VLAN sub-interface are considered to be part of the same flow that is mapped to one bucket. Each bucket is associated with both primary and secondary member-links. The bucket points to the active interface in the pair, either primary or secondary. Multiple VLAN flows can be mapped to the same bucket if their primary and secondary member-links mapping is the same.

Restrictions for VLAN-Based Load Balancing

- Only static mapping of VLAN sub-interfaces is supported; internal, dynamic load balancing is not supported.
- The primary member-link must be configured.
- When service policy is applied to port channel's main or subinterface, changing the load balancing mode from VLAN to flow is not supported.
- On a GEC bundle, executing the VLAN-based load balancing feature, along with the VLAN group QoS feature, requires that all VLAN sub-interfaces in a particular VLAN group have the same primary and secondary member-links configured.
- When a hierarchical queuing policy is applied on a port channel VLAN sub-interface, the hierarchical queues related to the policymap are created on both primary and secondary member-links associated with the VLAN sub-interface. Egress traffic for this VLAN sub-interface is enqueued on to the corresponding queues, at the primary member-link. In an event of the primary member-link going down, the traffic for the VLAN sub-interface is redirected to the queues on the secondary member-link.

- When LACP max-bundle is used with VLAN-based load balancing, user needs to verify whether both primary and secondary member-links for a VLAN sub-interface are not selected as standby links by the LACP protocol.
- Either of the Flow-based Load-balancing feature or VLAN-based Load Balancing feature for a GEC bundle is supported. The Flow and VLAN-based Load Balancing can co-exist on the same router, if they are selected for different GEC bundles.
- VLAN-based load balancing for QinQ sub-interfaces is not supported.
- Per VLAN bandwidth allocation cannot be more than 1 Gbps.
- Per VLAN group bandwidth allocation cannot be more than 1 Gbps.
- BRR can work only on a per member-link basis.
- PPPoX and IP/DHCP session are not supported.
- QoS on member-links is not supported.
- 1:N mode is not supported, only M:N mode is supported.
- The Link-switchover time depends on the carrier-delay configured on the member links and for VLAN-based load balancing the switchover time is expected to be same as flow-based load balancing.
- Because the HQF hierarchy for a VLAN subinterface on the GEC bundle is programmed on both primary and secondary member-links, the HQF-related resources, such as class queues, and the logical BLTs usage, would double and therefore, reduce the scalability of the QoS.
- If the member-link corresponding to the VLAN is not over-subscribed, there is no impact on the traffic of other VLANs, when new VLANs are added on the functional port channel.
- With the VLAN group QoS feature, the class queues are created on all the member links. The scalability of QoS reduces as the class queues are allocated from the same pool.

Configuration Tasks

To configure the VLAN-based Load Balancing feature on a GEC bundle, enter the following commands:

	Command	Purpose
Step 1	router(config)# interface port-channel number	Creates a GEC bundle.
Step 2	router(config-if)# load-balancing vlan	Enables the VLAN-based Load Balancing feature on the GEC bundle.
Step 3	router(config-if)# lacp max-bundle 1-8	Assigns the maximum number of links per bundle. The default value is 8.
Step 4	router(config-if)# lacp min-link 1-8	Sets the minimum number of active links required before declaring the port channel interface down. The default value is 1.
Step 5	router(config)# interface interface	Selects an GigabitEthernet interface to configure and enters config-if mode.
Step 6	router(config-if)# channel-group 1-64 [mode {active passive}]	Configures the member link so that it can be added to the port channel interface.
Step 7	router(config)# interface port-channel subinterface	Creates a GEC bundle subinterface and enters the subinterface mode.

Configuring VLAN-Based Load Balancing

	Command	Purpose
Step 8	<code>router(config-subif)# encapsulation dot1Q vlan-id primary member-link secondary member-link</code>	Enables IEEE 802.1Q encapsulation of traffic on a specified subinterface in a VLAN. Specify the VLAN identifier, primary and secondary links. Note The primary and secondary links must be a part of the port channel so that traffic can be forwarded to these links. If any of primary or secondary links are not part of the port channel at any state, the VLAN sub-interface gets in a suspended state and traffic is not forwarded.
Step 9	<code>router(config-subif)# end</code>	Exits to the global configuration mode.



Note The primary and secondary links must be configured. If primary and secondary links are not specified during the configuration of the VLAN-based Load Balancing feature, by default, the primary and secondary links are selected. When default links are chosen, equal VLAN distribution across links does not occur. If only a primary link is specified, by default, a secondary link that is different than the specified primary is selected.

The system allows the user to see the type of port channel load-balancing method that is enabled, using the **show ether-channel load-balancing** command. The output of the **show vlans** command displays information such as the mapping of VLANs, the member link (primary or secondary) that is currently used by the traffic for a given VLAN, and the VLANs for which a given member-link is active.



Note The **port-channel load-balancing vlan-manual** command applies the VLAN-manual load-balancing method globally to all GEC interfaces.

Configuration Example

This section provides the following configuration examples:

- Configuration Example of VLAN-Based Load Balancing, page 23-14
- Configuration Example for Applying VLAN QoS on GEC Bundle Subinterfaces, page 23-15
- Configuration Example for Using the VLAN Group Feature to Apply QoS, page 23-16



Note When service policy is applied to port channel's main or subinterface, changing the load balancing mode from VLAN to flow is not supported.

Configuration Example of VLAN-Based Load Balancing

Example 23-5 shows how to configure the VLAN-based Load Balancing feature on a GEC subinterface:

Example 23-5 Configuring The VLAN-Based Load Balancing Feature

```

configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
interface port-channel 1
load-balancing vlan
lacp max-bundle 2
exit
!
interface gigabitethernet2/1/0
no ip address
channel-group 1 mode active
exit
!
interface gigabitethernet8/0/0
no ip address
channel-group 1 mode active
exit
!
interface port-channel 1.1
encapsulation dot1q 2 primary gigabitethernet2/1/0 secondary gigabitethernet8/0/0
ip address 3.0.0.1 255.255.255.0
no sh

```

Configuration Example for Applying VLAN QoS on GEC Bundle Subinterfaces

[Example 23-6](#) shows how VLAN QoS is applied on GEC bundle subinterfaces:

Example 23-6 Applying VLAN QoS on GEC Bundle Subinterfaces

```

Class-map match-any dscp_20_30
    Match dscp 20 30
Class-map match-any dscp_40
    Match dscp 40

Policy-map police_dscp
    Class dscp_20_30
        Police 50 3000 3000 conform-action transmit exceed-action drop
        Set ip dscp af22
    Class dscp_40
        Police 10 3000 3000 conform-action transmit exceed-action drop

Policy-map customer_A
    Class class-default
        Police 100 mpbs
        service-policy police_dscp

Policy-map customer_B
    Class class-default
        Police 150 mbps
        Service-policy police_dscp

Interface Port-channel 1.1
    Service-policy input customer_A
    encapsulation dot1q 1 primary gigabitethernet2/1/0 secondary gigabitethernet8/0/0

Interface Port-channel 1.2
    Service-policy input customer_B

```

■ Configuring VLAN-Based Load Balancing

```
encapsulation dot1q 1 primary gigabitethernet2/1/0 secondary gigabitethernet8/0/0
```

Configuration Example for Using the VLAN Group Feature to Apply QoS

Assume that the following configurations need to be performed on a port channel bundle

- Police ingress traffic for VLAN 2 at 100 mbps
- Police ingress traffic for VLAN 3 at 150 mbps
- Shape egress traffic for VLAN 3 at 50 mbps
- Shape egress traffic for VLAN 2 & 4 together at 150 mbps

Step 1 Create a port-channel bundle as follows:

```
interface port-channel 1
load-balancing vlan
no sh
```

Step 2 Create VLAN subinterfaces as follows:

```
interface port-channel 1.1
encapsulation dot1q 2 primary gig2/0/0 secondary gig3/0/0
ip add 3.0.0.1 255.255.255.0
no sh

interface port-channel 1.2
encapsulation dot1q 3 primary gig2/0/0 secondary gig3/0/0
ip add 3.1.0.1 255.255.255.0
no sh

interface port-channel 1.3
encapsulation dot1q 4 primary gig2/0/0 secondary gig3/0/0
ip add 3.2.0.1 255.255.255.0
no sh
```

Step 3 Create match-any vlan class-maps as follows:

```
Class-map match-any vlan_2
Match vlan 2
Class-map match-any vlan_3
Match vlan 3
Class-map match-any vlan_4
Match vlan 4
Class-map match-any vlan_2_4
Match vlan 2 4
```

Step 4 Create policy-maps as follows:

```
Policy-map mega_ingress
Class vlan_2
Police 100 mbps
Class vlan_3
Police 150 mbps
Policy-map mega_egress
Class vlan_3
Shape 50 mpbs
Class vlan_2_4
Shape 150 mbps
```

Step 5 Apply the policy on the port-channel bundle

```
Interface port-channel 1
Service-policy input mega_ingress
Service-policy output mega_egress
```

■ Configuring VLAN-Based Load Balancing