



Configuring Modular QoS on Link Bundles on Cisco IOS XR Software

A link bundle is a group of one or more ports that are aggregated together and treated as a single link. This module describes QoS on link bundles.

Feature History for Configuring Modular QoS Congestion Avoidance on Cisco IOS XR Software

Release	Modification
Release 3.6.0	The QoS on Link Bundles feature was introduced.

Contents

- [Link Bundling Overview, page 127](#)
- [Load Balancing, page 128](#)
- [QoS and Link Bundling, page 129](#)
- [Additional References, page 133](#)

Link Bundling Overview

The Link Bundling feature allows you to group multiple point-to-point links together into one logical link and provide higher bidirectional bandwidth, redundancy, and load balancing between two routers. A virtual interface is assigned to the bundled link. The component links can be dynamically added and deleted from the virtual interface.

The virtual interface is treated as a single interface on which one can configure an IP address and other software features used by the link bundle. Packets sent to the link bundle are forwarded to one of the links in the bundle.

The advantages of link bundles are as follows:

- Multiple links can span several line cards and SPAs to form a single interface. Thus, the failure of a single link does not cause a loss of connectivity.

- Bundled interfaces increase bandwidth availability, because traffic is forwarded over all available members of the bundle. Therefore, traffic can move onto another link if one of the links within a bundle fails. You can add or remove bandwidth without interrupting packet flow. For example, you can upgrade from an OC-48c PLIM modular services card to an OC-192 PLIM modular services card without interrupting traffic.

All links within a bundle must be of the same type. For example, a bundle can contain all Ethernet interfaces, or it can contain all POS interfaces, but it cannot contain Ethernet and POS interfaces at the same time.

Cisco IOS XR software supports the following methods of forming bundles of Ethernet and POS interfaces:

- IEEE 802.3ad—Standard technology that employs a Link Aggregation Control Protocol (LACP) to ensure that all the member links in a bundle are compatible. Links that are incompatible or have failed are automatically removed from a bundle.
- EtherChannel or POS Channel—Cisco proprietary technology that allows the user to configure links to join a bundle, but has no mechanisms to check whether the links in a bundle are compatible. (EtherChannel applies to Ethernet interfaces, and POS Channel applies to POS interfaces.)

The following types of link bundling are supported:

- EtherChannel is used to bundle multiple Gigabit Ethernet (GE) interfaces.
- Gigabit EtherChannel link bundles are at Layer 2 and use one MAC address and one IP address for all GigabitEthernet interfaces in the bundle.
- POS Channel is used to bundle multiple Packet-over-SONET (POS) interfaces.

Gigabit EtherChannel and POS Channel link bundling provide the following benefits:

- Flexible, incremental bandwidth
- Transparency to network applications
- Support for IP unicast and MPLS traffic
- Load balancing (equal cost) across all active links on the bundle
- Redundancy: if there is a failure of an individual GE or POS link, the traffic flow through the channel is evenly distributed across the available links.
- Interoperability with link bundling implementations in other Cisco and OEM routers and switches.
- Out-of-service support: a Gigabit EtherChannel or POS Channel is brought down if the minimum number of Gigabit Ethernet or POS links are not up.
- Bandwidth propagation support: bandwidth changes in a Gigabit EtherChannel or POS Channel can be (optionally) propagated to the upper-layer protocols until the amount of bandwidth required in the link bundle exceeds a specified threshold.

Load Balancing

Load balancing is supported on all links in the bundle. Load balancing function is a forwarding mechanism to distribute traffic over multiple links based on layer 3 routing information in the router. There are two types of load balancing schemes:

- Per-Destination Load Balancing
- Per-Packet Load Balancing

When a traffic stream arrives at the router, per-packet load balancing allows the traffic to be evenly distributed among multiple equal cost links. Per-packet schemes make routing decision based on round-robin techniques, regardless of the individual source-destination hosts.

Only Per-Destination Load Balancing is supported.

Per-destination load balancing allows the router to distribute packets over one of the links in the bundle to achieve load sharing. The scheme is realized through a hash calculating based on the source-destination address and user sessions.

When the per-destination load balancing is enabled, all packets for a certain source-destination pair will go through the same link, though there are multiple links available. In other words, per-destination load balancing can ensure that packets for a certain source-destination pair could arrive in order.

QoS and Link Bundling

Quality of Service (QoS) features, currently supported on Ethernet and Packet-over-SONET (POS) interfaces, are also supported on Link Bundle interfaces. This section describes the features and restrictions that apply to QoS Link Bundles.

Features

The following QoS features are supported on Link Bundles interfaces:

- IP precedence
- Differentiated Services Code Point (DSCP)
- MPLS Exp
- Class of Service (CoS)
- Discard
- Multicast QoS
- QoS group
- Shaping
- Hierarchical Policing
- Policing
- Bandwidth
- Bandwidth Remaining
- Policy Map Modification

The following QoS features are supported on Link Bundles only through the Modular QoS CLI (MQC).

- Marking (IP prec, MPLS exp, 802.1Q, DSCP)
- Priority
- Shaping
- Input/Output Committed Access Rate (CAR) policing
- Weighted Random Early Detection (WRED)
- Modified Deficit Round Robin (MDRR)

Restrictions

QoS is configured on Link Bundles primarily in the same way that it is configured on individual interfaces. However, there are some important restrictions that must be observed:

- QoS features are applied only on Link Bundle interfaces — never on individual members.
- For all calculations, the total bandwidth used on a Link Bundle line card is the sum of all bandwidths for all members of all bundles on the line card. Consequently, all rates for policing, shaping, and bandwidth must be configured in percentages instead of absolute values, and the WRED/Qlimit must be configured in units of time.
- In general, show commands aggregate statistics across all line cards on which a bundle interfaces have been replicated.
- The **show-policy map** command displays information for a bundle. As a result, queue-related information is not displayed in the output of this command.
- The **show qos** commands for output policies display data and queueing information in the Execution Agent (EA), for each member. For egress QoS policies, the policer is applied at bundle level, while all other QoS features (e.g. shaping) are applied at member level. As a result, the police information is displayed on a per bundle basis, while all other QoS information is displayed on a per member basis. For ingress QoS policies, all output from **show qos** commands is displayed on a per bundle basis.
- Queueing actions are applied on each link-bundle-member. This may cause unexpected QoS behaviour in some cases. For example, if 50% bandwidth is guaranteed on a bundle with two member links, using the **bandwidth** command, the effect is to guarantee 50% bandwidth on each member link. Since the traffic distribution on the members may not be 50% each, the bandwidth for each bundle may be different to what is expected.
- There is no QoS in slow path for Link Bundles.
- VLAN-based load sharing is not supported on EtherBundles, and QoS reflects this.
- **match cos** and **match vlan** commands are not supported on link bundles.

In addition, the following points should be noted, when configuring link bundles which contain both 2.5 Gbps IP Services Engine (E3) Line Card and 10 Gbps IP Services Engine (E5) Line Card links.

- 2.5 Gbps IP Services Engine (Engine 3) and 10 Gbps IP Services Engine (Engine 5) hardware granularities are different for shaper and policer. Hence, for the same configured rate, traffic flow via E3 and E5 line card links may differ.
- For E3 Line Card links, policy accounting is performed at Layer 3, by default. For E5 line card links, policy accounting is performed at Layer 2, by default.
- Full QoS for multicast traffic is supported on E5 Line Cards. The QoS policing and marking features are not supported for multicast traffic on E3 Line Cards.

Marking

When traffic matches a certain class, marking (coloring) can be used to set or change the value of these fields in packet headers:

- IP Precedence
- MPLS EXP

- IP DSCP
- QoS Group
- Discard Class
- Dot1q P

**Note**

Dot1q P marking occurs at L2 layer. All other fields are marked at L3 layer.

Multiple re-marking is possible on bundle interfaces using the following calculation on the same class:

```
"set mpls exp imposition <>" + "set qos-group" + "set discard-class"
```

IP Precedence Marking

The IP precedence value in the Type of Service (TOS) field of an input or output IP packet may be set or changed by a user configured value.

MPLS EXP Marking

The three MPLS EXP (experimental) bits in the shim header of an input or output MPLS packet header may be set or changed by a user configured value.

DSCP Marking

The six Differentiated Services Code Point (DSCP) bits in the IP header field of an input or output IP packet on an Ether channel may be set or changed by a user configured value.

QoS Group Marking

The QoS Group field of an input IP packet received on an Ether Bundle may be set or changed by a user configured value, using the **set qos-group** command.

Discard Class Marking

The Discard Class field of an input or output IP packet on an Ether Bundle may be set or changed by a user configured value, using the **set discard-class** command. For egress only for continue-bit case may be set.

Dot1q P Bit Marking

The 3 priority bits in the Dot1q P field of a tagged ethernet frame may be set or changed by a user configured value. Tagged Ethernet frames are normal Ethernet frames with additional 4 bytes of information inserted between the Source Address (S/A) field and the Type or Length (T/L) field.

Policing

Traffic policing is used to control the rate of traffic transmitted and received on an interface, based on certain matching criteria defined by the user. Traffic policing performs 2 basic functions:

- Controlling the reception and transmission rates for traffic with different Classes of Service (CoS), based on user-defined criteria.
- Marking reception and transmission packets by setting the following fields in packet headers:
 - IP Precedence
 - MPLS EXP

- IP DSCP
- QoS Group
- Discard Class
- Dot1q P

When traffic conforms to or exceeds a configured rate limit, policing performs the following configurable actions:

- transmit
- drop
- set precedence
- set dscp and set qos

Packets that conform to a rate limit are handled with the configured conform-action. Packets that do not conform to a rate limit are handled with the configured exceed-action.

Policing is supported on Link Bundles for ingress and egress traffic.

Single-Rate Color Policing

Single-rate 2-color policing is supported on Link Bundles.

Single-rate 3-color policing is supported on Link Bundles.

Dual-Rate Color Policing

Dual-rate 3-color policing is supported on Link Bundles, using the **peak-rate** keyword of the **police rate** command:

```
police rate peak-rate conform-action exceed-action
```

QoS Group as a Policing Action

You can set the QoS Group as a policing action, using the **police** command:

```
police conform-action set qos-group
```

```
police exceed-action set qos-group
```



Note

Only one **conform-action** or one **exceed-action** may be set per action type.

Hierarchical Policing

Both hierarchical ingress policing and enhanced hierarchical ingress policing is supported, where different policing actions may be configured at different levels; on the parent class level and on one or more of its child classes.

Additional References

The following sections provide references related to implementing QoS on Link Bundles.

Related Documents

Related Topic	Document Title
Initial system bootup and configuration	<i>Cisco IOS XR Getting Started Guide for the Cisco XR 12000 Series Router</i>
Link Bundling	“Configuring Link Bundling on Cisco IOS XR Software” module of <i>Cisco IOS XR Interface and Hardware Component Configuration Guide for the the Cisco XR 12000 Series Router</i>
Master command reference	<i>Cisco XR 12000 Series Router Master Command Listing</i>
QoS commands	<i>Cisco IOS XR Modular Quality of Service Command Reference for the Cisco XR 12000 Series Router</i>
User groups and task IDs	“Configuring AAA Services on Cisco IOS XR Software” module of <i>Cisco IOS XR System Security Configuration Guide</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport