



Configuring Modular QoS on Link Bundles

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.

Line Card, SIP, and SPA Support

Feature	ASR 9000 Ethernet Line Cards	SIP 700 for the ASR 9000
QoS on Link Bundles	Yes	Yes

Feature History for Configuring Modular QoS on Link Bundles on Cisco ASR 9000 Series Routers

Release	Modification
Release 3.9.0	The QoS on Link Bundles feature was introduced on ASR 9000 Ethernet Line Cards.
Release 6.0.1	The aggregate bundle QoS feature was introduced on ASR 9000 Ethernet Line Cards.

- [Link Bundling Overview, on page 1](#)
- [Load Balancing, on page 2](#)
- [QoS and Link Bundling, on page 3](#)
- [QoS for POS link bundling, on page 3](#)
- [Aggregate Bundle QoS Mode, on page 4](#)
- [Bundle Aggregate Policer, on page 7](#)
- [Additional References, on page 10](#)

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.

A link bundle is simply a group of ports that are bundled together and act as a single link. The advantages of link bundles are as follows:

- Multiple links can span several line cards 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 flow on the available links if one of the links within a bundle fails. Bandwidth can be added without interrupting packet flow.

All the individual links within a single bundle must be of the same type and the same speed.

Cisco IOS XR software supports these methods of forming bundles of Ethernet 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—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.

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.

Layer 3 Load Balancing on Link Bundles

By default, load balancing on Layer 2 link bundles is done based on the MAC SA/DA fields in the packet header. Layer 3 load balancing for link bundles is done on Ethernet Flow Points (EFPs) and is based on the IPv4 source and destination addresses in the packet. When Layer 3 service-specific load balancing is configured, all egressing bundles are load balanced based on the IPv4 source and destination addresses. When packets do not have IPv4 addresses, default load-balancing is used.

QoS and Link Bundling

All Quality of Service (QoS) features, currently supported on physical interfaces and subinterfaces, are also supported on all Link Bundle interfaces and subinterfaces. QoS is configured on Link Bundles in the same way that it is configured on individual interfaces. However, the following points should be noted:

- When a QoS policy is applied on a bundle (ingress or egress directions), the policy is applied at each member interface. Any queues and policers in the policy map (ingress or egress directions) will be replicated on each bundle member.
- If a QoS policy is not applied to a bundle interface or bundle VLAN, both the ingress and egress traffic will use the per link members port default queue.
- Link bundle members may appear across multiple Network Processing Units and linecards. The shape rate specified in the bundle policymap is not an aggregate for all bundle members. The shape rate applied to the bundle will depend on the load balancing of the links. For example, if a policy map with a shape rate of 10 Mbps is applied to a bundle with two member links, and if the traffic is always load-balanced to the same member link, then an overall rate of 10 Mbps will apply to the bundle. However, if the traffic is load-balanced evenly between the two links, the overall shape rate for the bundle will be 20 Mbps.

[Example 1](#) shows how a traffic policy is applied on an Ethernet link bundle, in the ingress direction. The policy is applied to all interfaces that are members of the Ethernet link bundle.

Example 1 Applying a Traffic Policy to an Ethernet Link Bundle

```
interface Bundle-Ether bundle-id
  service-policy input policy-1
end
```

QoS for POS link bundling

For POS link bundles, percentage-based bandwidth is supported for policers and output queues. Time-based queue limit is supported for output queues.

Input QoS Policy setup

For input QoS, queuing is not supported and thus bandwidth is used for policer only. As a member link is added or removed from a bundle with input QoS configured, the aggregate bundle bandwidth for that affected line card will change. One input QoS policy instance is assigned for each SIP 700 line card that is part of the POS link bundle.

Output QoS Policy setup

When a member link is added to a bundle with output QoS configured, the policy-map of the bundle is applied to the member link.

Example 2 shows the output QoS policy supported on POS link bundles.

Example 2 : Output QoS policy supported on POS link bundles

```

policy-map out-sample
  class voice
    priority level 1
    police rate percent 10
  class premium
    bandwidth percent 30
    queue-limit 100 ms
  class class-default
    queue-limit 100 ms

```

Aggregate Bundle QoS Mode

Aggregated Bundle QoS allows the shape, bandwidth, police rates, and burst values to be distributed between the active members of a bundle where a QoS policy-map is applied. For instance, consider that the traffic is load-balanced among the members of the bundle. In aggregate mode, the bundle ethernet traffic is shaped to 10 Mbps to match the configuration of QoS policy.

When the policy is applied on a member of the bundle, a ratio can be calculated based on the total bandwidth of the bundle to that of the bandwidth of a member in the bundle. For example, if the bandwidth of the bundle is 20 Gbps, and the bandwidth of a member in the bundle is 10 Gbps, then the ratio will be 2:1.

A change in the bundle (with a member down, added, removed or activated) or mode results in the automatic recalculation of QoS rate.

The user QoS policy is invalid when applied to the bundle interface under the following scenarios:

- A 10 Gbps interface and 40 Gbps interface are part of a bundle, and the 40 Gbps interface is inactive. Currently, QoS policy is also programmed on non-active members. When programming the 40Gbps bundle member, the bundle bandwidth is 10 Gbps, but the member bandwidth is 40 Gbps. The ratio of bundle bandwidth to member bandwidth does not work for this member.
- Consider a shape of 15 Gbps. This action is valid on bundles with multiple 10G active members, but invalid when only one member is active and that member is in QoS inconsistent state. To view inconsistency details for the QoS policy, run the **show qos inconsistency** command in EXEC mode. This scenario is also applicable during the reload of a router where only few interfaces in line cards (LC) becomes available before the rest of the interfaces in all LCs.
- A failure in the hardware when programming a rate change during the bundle bandwidth change or when a new member is added to the bundle.
- If an interface has the QoS policy configured to an absolute value, you cannot change the aggregated bundle mode from enabled to disabled. You must modify the policy or remove it before attempting to disable the aggregate bundle mode.
- An invalid policy combination, with the absolute values of port shaper less than the policy shape rate is accepted without an error in console or log file.
- You apply QoS policy on bundle interfaces for BNG subscribers. In this case, there is no aggregation across bundle members, and the policy is applied individually. For the policy to work correctly, modify the rate according to the number of members in the bundle.
- You apply QoS policy on PWHE and BVI interfaces. Both these interfaces don't support bundle aggregate mode.

Load Balancing in Aggregate Bundle QoS

Load balancing requires a large number of flows in order to distribute the traffic among the members of the bundle. Ensure that load is balanced evenly among the members of the bundle before using the aggregate bundle QoS mode. If the under-lying traffic is only a few tunnels (GRE, TE-TUNNELS), it may be possible that the load balancing is not distributing the traffic evenly and may cause problems.

For example, consider bandwidth of a bundle is 20 Gbps, and the bandwidth of a member in the bundle is 10 Gbps. If the traffic is not load balanced, the aggregate traffic output may not reach 10 Mbps even when more than 10 Mbps is sent to the bundle-ether interface.

QoS Policy in Aggregate bundle mode

The following table shows the behavior of aggregate QoS policy mode to various actions:

Action	Behavior
Policing and Shaping / Bandwidth	<p>Percentage: No change. The percentage is calculated based on <code>max-rate / interface bandwidth</code></p> <p>PPS / Absolute rate: Divide the rate based on bandwidth</p> <p>Burst-size: If <code>time-units</code>, no change. Convert <code>time-units</code> based on <code>service-rate</code></p> <p>If configured in absolute value, divide the absolute value by bandwidth ratio</p>
Wred / Queue-Limit Threshold	<p>Time-Units: No Change. Use the <code>service-rate</code> to convert</p> <p>Absolute value: Divide the absolute value based on bandwidth</p>

Enabling Aggregate Bundle QoS

To enable the aggregate bundle QoS, perform these steps:

SUMMARY STEPS

1. `configure`
2. `hw-module all qos-modeaggregate-bundle-mode`
3. `class class-name`
4. `end` or `commit`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>configure</code> Example: <code>RP/0/RSP0/CPU0:router# configure</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	hw-module all qos-modeaggregate-bundle-mode Example: RP/0/RSP0/CPU0:router(config)# hw-module all qos-mode aggregate-bundle-mode	Enters policy map configuration mode. When aggregated bundle mode changes, QoS polices on bundle interfaces and sub-interfaces are modified automatically. A reload of the line card is not required.
Step 3	class class-name Example: RP/0/RSP0/CPU0:router(config-pmap)# class class-default	Enters policy map class configuration mode. Specifies the name of the class whose policy you want to create or change.
Step 4	end or commit Example: RP/0/RSP0/CPU0:router(config-pmap-c-police)# end or RP/0/RSP0/CPU0:router(config-pmap-c-police)# commit	

Policing using Aggregate Bundle QoS

```

policy-map grand-parent
  class class-default
    service-policy parent
    police rate 200 mbps burst 1 kbytes
  end-policy-map

policy-map parent
  class class-default
    service-policy child
    police rate 300 mbps
  end-policy-map

policy-map child
  class 3play-voip
    police rate 10 mbps burst 10 kbytes peak-rate 20 mbps peak-burst 20 kbytes
    conform-color red-cos
    conform-action set precedence 1
    exceed-action set precedence 2
    violate-action drop

class 3play-video
  police rate 15 mbps burst 10 kbytes peak-rate 30 mbps peak-burst 20 kbytes
  conform-color yellow-cos
  conform-action set precedence 1
  exceed-action set precedence 2
  violate-action drop
!
!
class 3play-premium
  police rate 25 mbps burst 10 kbytes peak-rate 35 mbps peak-burst 20 kbytes
  conform-color green-cos

```

```
conform-action set precedence 1
exceed-action set precedence 2
violate-action drop
!
!
class class-default
  police rate 6 mbps
!
!
end-policy-map
!
```

Bundle Aggregate Policer

If a policy map having policer is applied on a bundle interface, the policy is programmed individually on the bundle members. This behaviour means that each bundle member has its own token bucket which results in inconsistencies. For example, if the policer is 1 Gbps and there are three bundle members, the presence of three token buckets allow for 3 Gbps traffic.

Configuring the **hw-module all qos-mode bundle-agg policer** command shares a policer token bucket between members of the bundle interface in the same network processor. In the preceding example, the three bundle members share 1 Gbps policer in the same network processor.

You must reload the line card after configuring this command. Because this operation is in the global configuration mode, all policers on the line card are configured in the aggregate mode.

In bundle aggregate policer mode, a policy replicates across every network processor. In the normal mode, the policy replicates across every member.

Restrictions

Ensure that you read these points before you configure the bundle aggregate policer.

- For policer configured with percent option, the policer percentage is calculated on the bundle member's bandwidth which gets added to the bundle first, and the policer is shared between all the bundle members in the same network processor. This scenario is applicable for mixed-speed bundles as well.



Note From Cisco IOS XR Release 7.5.1 onwards, configuring a percent policer value on a class allocates separate token buckets for each bundle member, which was the default behavior when **hw-module all qos-mode bundle-agg policer** wasn't enabled. See [Separate Token Buckets for Percentage Policers](#), on page 9 for details.

Policer percentage isn't calculated based on the overall aggregate bandwidth of bundle members in the same network processor.

- To remove a bundle member, ensure that you first remove all members from that network processor and then add back the required members to the processor. If you remove only one member or don't remove all members from the network processor, the policer may not work properly.

Enabling Bundle Aggregate Policer

To enable the bundle aggregate policer, perform the following steps:

SUMMARY STEPS

1. **configure**
2. **hw-module all qos-mode bundle-agg-policer**
3. **end** or **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure Example: RP/0/RSP0/CPU0:router# configure	Enters global configuration mode.
Step 2	hw-module all qos-mode bundle-agg-policer Example: RP/0/RSP0/CPU0:router(config)# hw-module all qos-mode bundle-agg-policer	You must reload the line card.
Step 3	end or commit Example: RP/0/RSP0/CPU0:router(config)# end Or RP/0/RSP0/CPU0:router(config)# commit	

Separate Token Buckets for Percentage Policers

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Separate Token Buckets for Percentage Policers	Release 7.5.1	<p>From this release, configuring a percent policer value on a class allocates separate token buckets for each bundle member. In other words, the policer percentage value is calculated on each single bundle member's bandwidth after configuring hw-module all qos-mode bundle-agg policer.</p> <p>This functionality allows you to use both aggregated absolute policers and nonaggregated percent policers. Plus, the percent policers span the entire bundle bandwidth (from 0 through 100%), allowing for better and efficient control over traffic flow.</p> <p>In earlier releases, the policer percentage for the entire bundle was calculated based on the bundle member's bandwidth that you added first. Such a calculation dragged down the traffic rate for the entire bundle.</p>

Overview and Benefits

In earlier releases, configuring the **hw-module all qos-mode bundle-agg policer** command shared a policer token bucket between members of the bundle interface in the same network processor. However, all policers on the line card were configured in the aggregate mode. For percent policers, the percentage was calculated based on the bandwidth of the first bundle member and not on the aggregated bandwidth. Such a calculation dragged down the traffic rate for the entire bundle.

You could have scenarios where you require configuring policers with absolute aggregate values and nonaggregated percent values. For example:

- Configuring ingress policers on your customer subinterfaces, where the policers must have absolute values and aggregated on bundles, because you don't control ingress load-balancing.
- Configuring egress policers on main interfaces, where they can be percent values and not be aggregated because you control egress load-balancing.

From Cisco IOS XR Release 7.5.1, running the **hw-module all qos-mode bundle-agg policer** command allocates separate token buckets for each bundle member for percent policers, enabling their calculation across the entire bandwidth of the bundle.

When you configure **hw-module all qos-mode bundle-agg policer** and reload the line card:

- For classes that have absolute policer values configured: token bucket is shared between bundle members on the same network processor.
- For classes that have percent policer values configured: separate token bucket is allocated to each bundle member irrespective of the network processor they're on. This allocation is the default behavior when you don't enable **hw-module all qos-mode bundle-agg policer**.

This enhancement provides the following benefits:

- The ability for percent policers to span the whole bundle bandwidth (from 0% through 100%), thus allowing better and efficient control over traffic flow.
- The flexibility to configure policers with absolute values (aggregated) and percent values (nonaggregated). As an example, the following sample policy has class default with policer percent and other two classes with policer absolute rate values.

```

policy-map p_bundle
  class c_policer_2r3c
    police rate 10 mbps peak-rate 20 mbps
    conform-action transmit
    exceed-action set cos 4
    violate-action drop
  !
  !
  class c_policer_ratel
    police rate 5 gbps
  !
  !
  class class-default
    police rate percent 10
  !
  !
end-policy-map
!
```

Guidelines

- The feature is supported on ingress and egress service policies.
- It is supported on satellite bundle interfaces.
- This feature supports in-place modification of a service policy.
- [User Based Rate-Limiting \(UBRL\)](#) isn't supported with this functionality.

Additional References

These sections provide references related to implementing QoS on Link Bundles.

Related Documents

Related Topic	Document Title
Initial system bootup and configuration	<i>Cisco ASR 9000 Series Aggregation Services Router Getting Started Guide</i>
Link Bundling	“Configuring Link Bundling on the Cisco ASR 9000 Series Aggregation Services Router” module of <i>Cisco ASR 9000 Series Aggregation Services Router and Hardware Component Configuration Guide</i>
QoS commands	<i>Cisco ASR 9000 Series Aggregation Services Router Module of Service Command Reference</i>
User groups and task IDs	“Configuring AAA Services on Cisco ASR 9000 Series Router” module of <i>Cisco ASR 9000 Series Aggregation Services Router 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 the appropriate MIBs 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/cisco/web/support/index.html