QoS Policies for Link Bundles

QoS policies for link bundles

A link bundle QoS policy is a Modular QoS feature that

- applies unified QoS policies across a group of aggregated Ethernet or VLAN interfaces
- ensures consistent bandwidth management, load balancing, and scheduling for all member links, and
- simplifies configuration by allowing a single QoS policy to govern multiple physical links treated as one logical interface.

Load balancing on link bundles

Load balancing on link bundles distributes traffic across member links by using Layer 3 routing information. This mechanism includes per-destination load balancing. It ensures that all packets from the same source–destination pair use the same link, preserving packet order and preventing re-sequencing delays when multiple links are available. Per-destination load balancing is only supported on the router, which can distribute packets over any of the links in the bundle.

Layer 3 load balancing for link bundles operates on Ethernet Flow Points (EFPs) associated with the bundle interface. It is based on the IPv4 source and destination addresses in the packet. When IPv4 information is unavailable, the system automatically reverts to the default load balancing that uses the source and destination MAC (SA/DA) fields.

Usage guidelines for QoS policies on link bundles

Applying and inheriting QoS policies

- QoS policies are configured on link bundles in the same manner as they are on individual interfaces.

- When you apply a QoS policy on a bundle (in either ingress or egress direction), the queuing policy is automatically applied to each member interface. The reference bandwidth used for calculating shaper or bandwidth values is derived from the physical member interface, not the aggregate bundle bandwidth.
- The QoS policy applied to the bundle is inherited by all member links. The reference bandwidth used for shaping and scheduling is based on the bandwidth of each physical member interface, not the aggregate bandwidth of the bundle.
- The effective aggregate shaped rate depends on the number of active bundle members and how traffic is distributed across those members.
- If a QoS policy is not applied to a bundle interface, both the ingress and egress traffic use the default queue of the per link member port.
- The shape rate specified in a bundle policy map is not an aggregate rate for all bundle members. The configured shaper rate is applied per bundle member, and the effective aggregate shaped rate depends on traffic distribution across active member links.

  For example, if a 10 Mbps shaping policy is applied to a bundle with two member links and all traffic flows through one member, the total shape rate remains 10 Mbps. However, if traffic is balanced evenly across both members, the combined shaping rate becomes 20 Mbps.
- Add bundle members before you apply the QoS policy to the bundle. This action ensures that the QoS policy applies across the bundle and on the new member.

  If you add a new member after the policy is applied, and the new member's configuration is incompatible, the policy may fail to propagate and cause configuration errors.
- For absolute-value shapers on bundle interfaces, the configured shaper rate is replicated on each bundle member interface. The configured value is not a single aggregate bundle-wide shaper rate.
- If you add a lower-bandwidth member to a bundle that already has an absolute-value shaper policy, the commit can succeed, but member policy replication can remain incomplete for that member.

  To recover, remove the lower-bandwidth member, reduce the shaper rate so that it is less than or equal to the bandwidth of that member, and add the member again.

  If in-place policy modification fails after bundle member changes, remove the egress service policy from the bundle interface, update the shaper rate, and apply the egress service policy again.

Managing bundle changes

- If you add a lower-bandwidth member to a bundle that already has an absolute-value shaper policy, the commit can succeed, but member policy replication can remain incomplete for that member.

  To recover, remove the lower-bandwidth member, reduce the shaper rate so that it is less than or equal to the bandwidth of that member, and add the member again.
- If in-place policy modification fails after bundle member changes, remove the egress service policy from the bundle interface, update the shaper rate, and apply the egress service policy again.

Troubleshooting queue-limit errors

- When you encounter a queue-limit error, remove the existing QoS policy from the bundle interface. The removal process may take up to ten minutes to take effect, depending on the time elapsed since the error occurred. To monitor progress, check the **Retries performed** information in the syslog message.

- If a queue-limit error occurs, the syslog lists only one interface, even if multiple interfaces are affected. Review **operations** in the syslog to determine the number of affected interfaces.

```
Router#Aug 17 09:56:37.417 PDT: BM-DISTRIB[1156]:
%L2-BM-4-ERR_OP_RETRY_THRESHOLD :
Exceeded threshold for the number of retries for 4 operations on members.
First (and error) FortyGigE0/0/0/34 ('DPA_QOSEA' detected the 'warning'
condition 'Invalid Queue Limit.
Value should be in range 614400-390070272 bytes or equivalent value in usec or
msec based on the service rate of the queue'),
retries performed for 10 hrs 16 mins

Router#Aug 17 10:06:40.388 PDT: BM-DISTRIB[1156]:
%L2-BM-4-ERR_OP_RETRY_THRESHOLD :
Exceeded threshold for the number of retries for 4 operations on members.
First (and error) FortyGigE0/0/0/34 ('DPA_QOSEA' detected the 'warning'
condition 'Invalid Queue Limit.
Value should be in range 614400-390070272 bytes or equivalent value in
usec or msec based on the service rate of the queue'),
retries performed for 10 hrs 26 mins
```

Managing statistics and bundle changes

When you remove a member interface from a bundle, the total bundle statistics change because the counters associated with the detached link are lost. Always capture statistics before removing member links to ensure accurate reporting.

- Classification resources (TCAM or MAP) and counters are allocated per bundle interface.
- Queuing resources (VOQs) and counters are allocated per bundle member interface.

Configuring policing and queuing parameters

- You can apply a policing policy on either the bundle's main interface or a bundle subinterface. You cannot apply it to both at the same time.
- Starting with Cisco IOS XR Release Release 26.3.1, you can configure committed information rate (CIR) and peak information rate (PIR) values using absolute units on supported bundle main interfaces and bundle subinterfaces. On unsupported releases or routers, configure CIR and PIR values using percentage units.
- For absolute-value bundle policers, the configured CIR and PIR values are applied per logical hardware unit. On A100, K100, and P100-based Cisco 8000 routers, the logical hardware unit is a slice. On Q200-based Cisco 8000 routers, the logical hardware unit is an interface group.

  If bundle members are spread across multiple slices or interface groups, each participating logical hardware unit can enforce the configured CIR and PIR independently. Do not interpret the configured value as a single aggregate bundle-wide policer in that case.
- You can apply the queuing policy on the bundle's main interface or on its subinterface. You cannot apply it to both interfaces simultaneously.
- You can apply classification and marking policies on the bundle main interface or on its subinterface. You cannot apply these policies to both interfaces at the same time.

Absolute-value policing on bundle interfaces

Absolute-value policing on a bundle interface is a bundle QoS traffic-policing feature that

- uses the existing **police rate** command to configure fixed committed information rate (CIR) and peak information rate (PIR) values on supported bundle main interfaces and bundle subinterfaces for Layer 2 and Layer 3 use cases,
- applies the configured policer rate per logical hardware unit, and
- preserves the existing percentage-based policing behavior for bundle interfaces.

Use absolute-value policing when you want an ingress policer rate to match a fixed customer or service bandwidth profile without converting that rate into a percentage of interface bandwidth.

Before Cisco IOS XR Release 26.3.1, QoS policers on bundle interfaces supported percentage-based values, while policers on physical interfaces supported absolute values. This feature closes that bundle QoS gap for supported platforms.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Absolute-Value Policing on Bundle Interfaces	Release 26.3.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Centralized Systems (8600 [ASIC: Q200], 8400 [ASIC: K100]); Modular Systems (8800 [LC ASIC: Q200, P100, K100]). You can control ingress bandwidth with fixed policing rates on supported bundle interfaces. This helps enforce predictable bandwidth limits for customers, services, or AI workloads using explicit bandwidth values instead of interface-relative percentages.

The configured absolute policer rate is programmed by logical hardware unit. Use this behavior when you plan and verify the effective policing rate on a bundle interface.

Table 2: Logical hardware units for absolute-value policing

Cisco 8000 ASIC	Logical Hardware Unit	Policer Behavior
A100, K100, and P100	Slice	The configured policer rate is shared by all bundle members within the same slice. If bundle members are spread across slices, the configured rate is programmed on each participating slice.
Q200	Interface group	The configured policer rate is shared by all bundle members within the same interface group. If bundle members are spread across interface groups, the configured rate is programmed on each participating interface group.

Absolute-value policing and percentage-based policing use the same bundle QoS policy workflow, but they express the policer rate differently.

Table 3: Absolute-value policing and percentage-based policing

Absolute-Value Policing	Percentage-Based Policing
Uses fixed units, such as **police rate 30 mbps peak-rate 50 mbps**.	Uses a relative value that is based on interface bandwidth.
Helps align ingress rate enforcement with fixed customer or service bandwidth profiles.	Helps enforce ingress rates as a percentage of available interface bandwidth.
Is useful when you already know the CIR and PIR in bps, kbps, Mbps, or Gbps.	Is useful when the target rate must remain relative to the interface or member bandwidth.

The **show qos interface** output can differ based on whether bundle members are in the same logical hardware unit or in different logical hardware units. Use the output to verify the programmed policer rates for the bundle interface and its members.

Bundle policer across logical hardware units

A bundle interface has three active 10-Gbps member interfaces. Two members are in one slice or interface group, and one member is in another slice or interface group. You configure an ingress class with a 100-Mbps absolute policer rate on the bundle interface.

For A100, K100, and P100 platforms, the 100-Mbps policer is programmed per slice. For Q200 platforms, the 100-Mbps policer is programmed per interface group. In this example, traffic on the two members in the same logical hardware unit shares one 100-Mbps policer, while traffic on the member in the other logical hardware unit is policed by another 100-Mbps policer.

What absolute-value policing on bundle interfaces is not

Bundle aggregate mode is not supported on bundle interfaces. This restriction applies regardless of whether the policer uses absolute values or percentage-based values.

Absolute-value policing on a bundle interface does not always enforce one bundle-wide policer rate. If bundle members are spread across multiple slices or interface groups, each participating logical hardware unit can enforce the configured rate independently. The effective policing behavior across the bundle can therefore be higher than the single configured value.

Absolute-value policing is not to be confused with absolute-value shaping on bundle interfaces. Policing enforces an ingress traffic contract by transmitting, marking, or dropping traffic based on the configured rate. Shaping controls egress traffic by scheduling packets for transmission at a configured rate.

Benefits of absolute-value policing on bundle interfaces

Use this reference to identify the benefits, supported scope, and key behavior details for absolute-value policing on bundle interfaces.

Absolute-value policing on bundle interfaces provides these benefits:

- Fixed ingress bandwidth control: You can configure policer rates using absolute units instead of converting the target rate to a percentage.
- Existing MQC workflow: You use the existing **police rate** command for supported bundle QoS deployments.
- Physical-to-bundle policy alignment: You can use the same absolute rate design approach that is already available on physical interfaces in supported bundle QoS deployments.
- Logical-unit aware enforcement: The configured policer rate is applied per slice on A100, K100, and P100-based Cisco 8000 routers and per interface group on Q200-based Cisco 8000 routers.
- Predictable service limits: You can align ingress policing with fixed customer, service, or AI workload bandwidth profiles.
- Percentage-based compatibility: Existing percentage-based policing behavior remains supported.

Table 4: Reference details for absolute-value policing on bundle interfaces

Reference item	Supported value or behavior	Applies to	Notes
Interface scope	Bundle main interfaces and bundle subinterfaces	Layer 2 and Layer 3 bundle QoS	Applies to supported bundle QoS deployments.
Direction	Ingress	Traffic policing	Policing enforces ingress traffic contracts by transmitting, marking, or dropping traffic based on the configured rate.
Policer programming model	Per logical hardware unit	Bundle members within the same slice or interface group	The configured absolute value is not always one aggregate bundle-wide policer rate.
A100, K100, and P100-based Cisco 8000 router behavior	Shared policer per slice	Bundle members within a slice	All bundle members within the same slice share the configured policer rate. If members are spread across slices, each participating slice can enforce the configured rate.
Q200-based Cisco 8000 router behavior	Shared policer per interface group	Bundle members within an interface group	All bundle members within the same interface group share the configured policer rate. If members are spread across interface groups, each participating interface group can enforce the configured rate.
ETM and non-ETM mode	Supported	Supported bundle QoS deployments	Support applies to ETM and non-ETM modes for this feature.
Existing percentage behavior	Supported	Percentage-based policer configurations	Existing percentage-based policing behavior is unchanged.

Reference item	Supported value or behavior	Applies to	Notes
Scale	No scale change	Existing policer scale	The feature uses existing QoS scale expectations.
Member bandwidth	Policer rate must not exceed member interface bandwidth	Absolute policer rates on bundle interfaces	A policer rate that is greater than the member interface bandwidth fails.
Hierarchical policer rate relationship	Parent policer rate must be greater than the sum of child policer rates	Hierarchical policer policies	Applies when hierarchical policers are configured.

Policer behavior by bundle member placement

The effective policing behavior depends on whether bundle members are in the same logical hardware unit or spread across multiple logical hardware units.

Table 5: Comparison of bundle member placement for absolute-value policing

Attribute	Members in the Same Logical Hardware Unit	Members in Different Logical Hardware Units
Programmed rate	The configured policer rate is shared by the members in that slice or interface group.	The configured policer rate is programmed separately on each participating slice or interface group.
Effective aggregate behavior	The effective behavior maps more closely to the configured rate for that logical hardware unit.	The effective aggregate behavior can be higher than the single configured value because each logical hardware unit can enforce that value independently.
Verification output	The **show qos interface Bundle-Ether** *interface-id* **input** output shows the programmed policer values for that logical hardware unit.	The **show qos interface Bundle-Ether** *interface-id* **input** output can show separate bundle main entries by node or logical hardware unit.

Note Do not interpret an absolute policer value on a bundle interface as one aggregate bundle-wide policer when bundle members are spread across multiple slices or interface groups.

Guidelines for absolute-value policing on bundle interfaces

Use absolute-value policing for fixed ingress bandwidth limits

Use absolute-value policing when a supported bundle interface must enforce ingress traffic to a fixed CIR or PIR.

- Configure the rate with the existing **police rate** command.
- Use percentage-based policing when your router does not support absolute policer values on bundle interfaces or when the service design requires interface-relative rates.

This guideline applies to absolute-value policers on bundle main interfaces and bundle subinterfaces on supported Cisco 8000 Series routers. The supported bundle scope includes Layer 2 and Layer 3 QoS deployments, and ETM and non-ETM modes.

Before this release, policers on bundle interfaces supported percentage-based values, while policers on physical interfaces supported absolute values. This feature extends fixed-rate policing to supported bundle QoS deployments.

Using absolute values helps align ingress policing with customer, service, or AI workload bandwidth profiles that require explicit bandwidth limits.

For releases or routers outside this feature scope, configure percentage-based policing on bundle interfaces.

Use the supported bundle QoS scope for bundle policers

Absolute-value policing on bundle interfaces is supported for the documented bundle QoS scope. Bundle aggregate mode is not supported on bundle interfaces.

This guideline applies when you configure policers on supported bundle main interfaces and bundle subinterfaces, whether the policer uses absolute values or percentage-based values.

Interpret the configured value as a logical-unit policer rate

Treat the configured absolute policer value as a rate applied per logical hardware unit, not always as one aggregate bundle-wide policer.

- On A100, K100, and P100-based Cisco 8000 routers, the configured policer value is applied per slice.
- On Q200-based Cisco 8000 routers, the configured policer value is applied per interface group.

This guideline applies to every absolute-value policer configured on a supported bundle interface.

For policing on bundle interfaces, the router shares the configured policer rate across all bundle members within the same logical hardware unit.

If all bundle members are in the same slice or interface group, the configured rate maps more closely to the intended aggregate behavior for that logical hardware unit. If bundle members are spread across slices or interface groups, each participating logical hardware unit can enforce the configured rate independently, so the effective aggregate behavior can be higher than the single configured value.

Plan bundle member placement with the required policing behavior in mind, and verify the programmed policer values with the **show qos interface Bundle-Ether** *interface-id* **input** command.

Keep policer values within member bandwidth

Ensure that the configured absolute policer rates are less than or equal to the bandwidth of the bundle member interfaces.

- Check the bandwidth of existing bundle members before attaching an absolute-value policer policy.
- Check the bandwidth of planned bundle members before adding them to a bundle that uses an absolute-value policer policy.
- Account for the CIR and PIR when you configure **police rate** with **peak-rate**.

This guideline applies to bundles with one member speed and to bundles that include or may later include members with different speeds.

The policer is programmed on the logical hardware unit that includes the bundle members, so the configured policer rate must fit within the member interface bandwidth.

A policer rate that is greater than the member interface bandwidth fails.

Reduce the policer rate so that the configured value is supported by the bundle member bandwidth, or use bundle members with bandwidth that supports the configured rate.

Set parent policer rates above child policer rates

For hierarchical policer policies, configure the parent policer rate to be greater than the sum of the policer rates in the child classes.

- Calculate child policer rates before you configure the parent policy.
- Update the parent policer rate when you add child classes or increase child policer rates.

This guideline applies when you configure hierarchical policers with absolute values on supported bundle interfaces.

The parent policer must be able to contain the aggregate of the child policer rates.

An invalid parent and child policer rate relationship can cause the hierarchical policer configuration to fail.

Increase the parent policer rate or reduce the child policer rates so that the parent rate is greater than the sum of the child rates.

Apply the policing policy in the ingress direction

Apply absolute-value policing policies to supported bundle interfaces in the ingress direction.

This guideline applies when you attach a policy map that contains **police rate** *value* or **police rate** *value* **peak-rate** *value* to a bundle interface.

Applying the policy in the ingress direction aligns the configuration with the supported feature behavior and verification output.

Verify the attached input policy and programmed policing values with **show qos interface Bundle-Ether** *interface-id* **input** and **show policy-map type qos interface Bundle-Ether** *interface-id* **input pmap-name** *policy-map-name*.

Plan with existing QoS scale and percentage-based behavior

Plan absolute-value policing with the existing QoS scale expectations for bundle policers.

- Use the existing percentage-based policing behavior when the service design requires relative policing rates.
- Use existing bundle QoS planning guidance for inheritance, hierarchical policies, and other supported policer variations.

This guideline applies when you add absolute-value policing to an existing bundle QoS design or migrate a physical-interface policer design to a bundle interface.

The feature changes absolute-value support on bundle policers. It does not introduce a scale change, and existing percentage-based policing remains supported.

Keeping the existing scale and percentage-based behavior in view helps you add fixed-rate policing without redesigning unrelated QoS behavior.

If the design depends on relative rates or an unsupported release or router, keep the bundle policer percentage-based.

Configure an absolute-value policer on a bundle interface

Use this task to configure fixed ingress policing rates for traffic classes on a bundle interface.

- The configured policer value is applied per logical hardware unit.
- The configured value is not always a single aggregate bundle-wide policer rate.

This example matches IPv4 traffic by precedence value and applies an absolute committed rate and peak rate in Mbps. Replace the class name, policy-map name, bundle interface, match criteria, and policer rates with values for your network.

Percentage-based policing remains supported on bundle interfaces, but this example focuses solely on absolute-value policing.

Before you begin

- Use this task only for supported bundle QoS deployments. Bundle aggregate mode is not supported on bundle interfaces.
- Verify that your router supports absolute-value policing on bundle interfaces.
- Ensure that the bundle interface and its member interfaces are configured.
- Ensure that each absolute policer rate is less than or equal to the bandwidth of the bundle member interfaces.
- If you configure hierarchical policers, ensure that the parent policer rate is greater than the sum of the policer rates in the child classes.
- Identify the traffic classes and fixed ingress rates that you want to police.

Follow these steps to configure an absolute-value policer on a bundle interface.

Procedure

Step 1 Enter global configuration mode.

Example:

```
Router#configure
```

Step 2 Create a class map to identify the traffic class that requires policing.

Example:

```
Router(config)#class-map match-all bundle_prec3
Router(config-cmap)#match precedence ipv4 3
Router(config-cmap)#exit
```

This example uses IPv4 precedence as the match criterion. Use the match criteria that apply to your QoS design.

The class map identifies the traffic that the policy map polices.

Step 3 Create a policy map and configure an absolute policer rate for the traffic class.

Example:

```
Router(config)#policy-map bundle_abs_police
Router(config-pmap)#class bundle_prec3
Router(config-pmap-c)#police rate 30 mbps peak-rate 50 mbps
Router(config-pmap-c)#exit
Router(config-pmap)#class class-default
Router(config-pmap-c)#exit
Router(config-pmap)#end-policy-map
```

Use the **police rate** command with the absolute unit that matches the rate you want to configure. This example uses Mbps consistently for readability.

The policy map defines a fixed ingress CIR and PIR for the matched traffic class. With the default policer actions shown in the verification output, conforming traffic is transmitted, while exceed and violate traffic are dropped.

Step 4 Apply the policy map to the bundle interface in the ingress direction.

Example:

```
Router(config)#interface Bundle-Ether10
Router(config-if)#service-policy input bundle_abs_police
```

The input direction is the ingress direction. For bundle policers, the router applies the configured absolute policer value per logical hardware unit.

Step 5 Commit the configuration.

Example:

```
Router(config-if)#commit
Router(config-if)#end
```

The router attaches the ingress QoS policy to the bundle interface.

Step 6 Verify that the QoS policy is programmed on the bundle interface.

Example:

```
Router#show qos interface Bundle-Ether10 input
NOTE: Configured values are displayed within parentheses
Node 0/0/CPU0, Interface Bundle-Ether10 Ifh 0x78000084 (BundleMain) -- input policy
NPU Id:                        0
Total number of classes:       2
Interface Bandwidth:           100000000 kbps
Policy Name:                   bundle_abs_police
Accounting Type:               Layer1 (Include Layer 1 encapsulation and above)
------------------------------------------------------------------------------
Level1 Class                             =   bundle_prec3
Policer committed rate                   =   30000 kbps (30 mbits/sec)
Policer peak rate                        =   50000 kbps (50 mbits/sec)
Policer conform burst                    =   1024000 bytes (default)
Policer exceed burst                     =   2048000 bytes (default)
Policer conform action                   =   Just TX
Policer exceed action                    =   DROP PKT
Policer violate action                   =   DROP PKT

Level1 Class                             =   class-default

Interface HundredGigE0/0/0/1 Ifh 0x218 (Member) -- input policy
NPU Id:                        0
```

```
Total number of classes:       2
Interface Bandwidth:           100000000 kbps
Policy Name:                   bundle_abs_police
Accounting Type:               Layer1 (Include Layer 1 encapsulation and above)
------------------------------------------------------------------------------
Level1 Class                             =   bundle_prec3

Level1 Class                             =   class-default
```

The output excerpt confirms that the policy is applied in the input direction and that the configured CIR and PIR are programmed. The **Policer committed rate** and **Policer peak rate** fields show the programmed CIR and PIR. The configured values appear within parentheses.

The output confirms that absolute-value policing is implemented for the configured traffic class.

Step 7 Verify the policy-map statistics on the bundle interface.

Example:

```
Router#show policy-map type qos interface Bundle-Ether10 input pmap-name bundle_abs_police

Bundle-Ether10 input: bundle_abs_police

Class bundle_prec3
  Classification statistics          (packets/bytes)     (rate - kbps)
    Matched             :              900000/345600000              55000
    Transmitted         :              500000/192000000              30000
    Total Dropped       :              400000/153600000              25000
  Policing statistics                (packets/bytes)     (rate - kbps)
    Policed(conform)    :              500000/192000000              30000
    Policed(exceed)     :              200000/76800000               12000
    Policed(violate)    :              200000/76800000               13000
    Policed and dropped :              400000/153600000

Class class-default
  Classification statistics          (packets/bytes)     (rate - kbps)
    Matched             :                     0/0                        0
    Transmitted         :                     0/0                        0
    Total Dropped       :                     0/0                        0
```

The output excerpt shows that traffic is matching the configured class and that the policer is enforcing the configured traffic contract. Conform traffic is transmitted, while exceed and violate traffic are dropped based on the default policer actions.

Step 8 Verify the configured policy definition.

Example:

```
Router#show policy-map pmap-name bundle_abs_police detail
class-map match-all bundle_prec3
 match precedence ipv4 3
 end-class-map
!
class-map match-any class-default
 end-class-map
!
policy-map bundle_abs_police
 class bundle_prec3
  police rate 30 mbps peak-rate 50 mbps
  !
 !
 class class-default
  !
 !
```

```
 end-policy-map
!
```

The policy definition confirms that the configured **police rate** values match the values that you configured in the policy map.

The bundle interface has an ingress QoS policy that polices matched traffic to fixed absolute rates.

What to do next

After you add, remove, or replace bundle members, verify the programmed policer values with the **show qos interface Bundle-Ether10 input** command.

- If all bundle members are in the same slice or interface group, the configured rate maps more closely to the intended aggregate behavior for that logical hardware unit.
- If bundle members are spread across slices or interface groups, each participating logical hardware unit can enforce the configured rate independently.

Absolute-value shaping on bundle interfaces

Absolute-value shaping on a bundle interface is a bundle QoS traffic shaping feature that

- uses the existing **shape average** command to configure a fixed shaper rate on supported bundle main interfaces and bundle subinterfaces for Layer 2 and Layer 3 use cases,
- replicates the configured shaper rate on each bundle member interface, and
- preserves the existing percentage-based shaping behavior for bundle interfaces.

Use absolute-value shaping when customers, services, or AI workloads require predictable egress bandwidth allocations without relying on interface-relative percentages.

Feature history and shaping behavior

Table 6: Feature History Table

Feature Name	Release Information	Feature Description
Absolute-value shaping on bundle interfaces	Release 26.3.1	Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Centralized Systems (8600 [ASIC: Q200], 8400 [ASIC: K100]); Modular Systems (8800 [LC ASIC: Q200, P100, K100]). You can control egress bandwidth with fixed shaping rates on supported bundle interfaces. This helps enforce predictable bandwidth allocations for customers, services, or AI workloads without relying on interface-relative percentages.

Absolute-value shaping and percentage-based shaping use the same bundle-member replication model, but they express the shaper rate differently.

Table 7: Absolute-value shaping and percentage-based shaping

Absolute-Value Shaping	Percentage-Based Shaping
Uses fixed units, such as `shape average 500 mbps`.	Uses a relative value, such as `shape average percent 4`.
Helps align egress shaping with fixed customer, service, or transport bandwidth profiles.	Helps shape traffic as a percentage of available interface bandwidth.
Useful when you already know the target service rate in bps, kbps, mbps, or gbps.	Useful when the target rate must remain relative to the interface or member bandwidth.

Bundle member shaping

A bundle interface has three active 10-Gbps member interfaces. You configure an egress class with `shape average 500 mbps` on the bundle interface.

The configured 500 Mbps shaper rate is replicated on each bundle member interface. If traffic for that class is distributed across all three members, each member can shape that class to the configured 500 Mbps rate. The configured rate is not a single 500 Mbps aggregate shaper for the entire bundle.

What absolute-value shaping on bundle interfaces is not

Absolute-value shaping on a bundle interface is not a single aggregate bundle-wide shaper. If you need one configured value to represent the total shaped bandwidth across all active bundle members, do not interpret this feature as providing that behavior. The effective aggregate shaped rate depends on the number of active bundle members and how traffic is distributed across those members.

Absolute-value shaping is not to be confused with absolute-value policing on bundle interfaces. Shaping controls egress traffic by scheduling packets for transmission at a configured rate. Policing enforces a traffic contract by marking or dropping traffic that exceeds the configured rate.

Benefits of absolute-value shaping on bundle interfaces

Absolute-value shaping on bundle interfaces enables you to configure fixed egress shaping rates for supported bundle interfaces without converting those rates into percentages. This feature is particularly useful when a service, customer, or transport design requires a specific configured rate in bps, kbps, Mbps, or Gbps.

Absolute-value shaping on bundle interfaces provides these benefits:

- Fixed service-rate expression: You can configure shaping values using absolute units (bps, kbps, Mbps, Gbps) instead of converting the target rate to a percentage.
- Existing MQC workflow: You use the existing **shape average** command. No new aggregate bundle mode CLI is introduced.
- Per-member egress rate control: The configured shaper value is replicated on each bundle member interface, so each member shapes traffic to the configured value.

- Percentage-based compatibility: Existing percentage-based shaping behavior remains supported alongside absolute-value shaping.
- Predictable bandwidth allocations: You can align bundle egress shaping with fixed customer, service, or AI workload bandwidth profiles.

Guidelines for absolute-value shaping on bundle interfaces

Use absolute-value shaping for fixed egress rates

Use absolute-value shaping when customers, services, or AI workloads require fixed egress bandwidth allocations on a supported bundle interface.

- Configure the rate with the existing **shape average** command.
- If your router does not support absolute shaper values on bundle interfaces, use percentage-based shaping instead.

This guideline applies to absolute-value shapers on bundle main interfaces and bundle subinterfaces on supported Cisco 8000 Series routers. The supported bundle scope includes Layer 2 and Layer 3 QoS deployments, and ETM and non-ETM modes.

Before this release, shapers on bundle interfaces supported percentage-based values, while shapers on physical interfaces supported absolute values. This feature extends fixed-rate shaping to bundle QoS without introducing a new aggregate bundle mode CLI.

Using absolute values helps align bundle egress shaping with customer, service, or transport bandwidth profiles.

For releases or routers outside this feature scope, configure percentage-based shaping on bundle interfaces.

Interpret the configured value as a per-member shaper rate

Treat the configured absolute shaper value as a per-member egress rate.

- The router replicates the configured shaper value to each bundle member interface.
- Plan the aggregate behavior based on active bundle members and traffic distribution across those members.

This guideline applies to every absolute-value shaper configured on a supported bundle interface.

A configured value such as **500 mbps** is therefore applied to each member interface that receives the replicated policy.

When you plan the service rate as a per-member value, the bundle design reflects the way the shaper is programmed and avoids incorrect aggregate-rate assumptions.

If the service outcome requires an aggregate bundle target, calculate the per-member shaper value for the expected active-member set and verify the programmed values with the **show qos interface** command.

Keep the shaper value within member bandwidth

Ensure that the configured absolute shaper value is less than or equal to the bandwidth of each bundle member interface.

- Check the bandwidth of existing bundle members before attaching an absolute-value shaper policy.

- Check the bandwidth of planned bundle members before adding them to a bundle that already has an absolute-value shaper policy.
- Check the **show bundle** output after you add a lower-bandwidth member to a bundle that already has an absolute-value shaper policy.

This guideline applies to bundles with one member speed and to bundles that include or may later include members with different speeds.

The shaper is programmed per member interface, so the absolute shaper rate must fit within the bandwidth of each member interface.

If the absolute shaper rate is greater than the member reference bandwidth, attaching the policy can fail with this message: `MQC: Shape bandwidth rate is greater than the parent reference bandwidth.`

If you add a lower-bandwidth member to a bundle that already has a valid absolute-value shaper policy, the commit can succeed, but member policy replication for that member can remain incomplete. The **show bundle** output can display this warning: `'qos-lib' detected the 'warning' condition 'MQC: Shape bandwidth rate is greater than the parent reference bandwidth'.`

After you remove a lower-bandwidth member, increasing the shaper value in an attached policy can fail with this message, even when the remaining member bandwidth can support the higher rate: `MQC: Shape bandwidth rate is greater than the parent reference bandwidth: InPlace Modify Error:.`

If the warning appears after you add a lower-bandwidth member, remove the member, reduce the shaper rate so that the configured value is less than or equal to the bandwidth of the lower-bandwidth member, and add the member again.

If in-place policy modification fails after bundle member changes, remove the egress service policy from the bundle interface, update the shaper rate to a value supported by the current bundle member bandwidth, and apply the egress service policy again.

Apply the shaping policy in the egress direction

Apply absolute-value shaping policies to supported bundle interfaces in the egress direction.

- Use this feature for egress shaping on bundle interfaces.
- Use bundle main interfaces or bundle subinterfaces within the documented Layer 2 and Layer 3 bundle QoS scope.

This guideline applies when you attach a policy map that contains **shape average <absolute-value with unit>** to a bundle interface in the egress direction.

Applying the policy in the egress direction aligns the configuration with the supported feature behavior and verification output.

Verify the attached output policy and programmed shaping values with **show qos interface** and **show policy-map type qos interface**.

Plan with existing QoS scale and percentage-based behavior

Plan absolute-value shaping with the existing QoS scale expectations for bundle shapers.

- Use the existing percentage-based shaping behavior when the service design requires relative shaping rates.
- Use existing bundle QoS planning guidance for hierarchical policies, inheritance, and other supported shaper variations.

This guideline applies when you add absolute-value shaping to an existing bundle QoS design or migrate a physical-interface shaping design to a bundle interface.

The feature changes absolute-value support on bundle shapers. It does not introduce a scale change, and existing percentage-based shaping remains supported.

Keeping the existing scale and percentage-based behavior in view helps you add fixed-rate shaping without redesigning unrelated QoS behavior.

If the design depends on relative rates or an unsupported release or router, keep the bundle shaper percentage-based.

Configure an absolute-value shaper on a bundle interface

Use this task to configure fixed egress shaping rates for traffic classes on a bundle interface.

- The configured shaper value is replicated on each bundle member interface.
- The configured value is not one aggregate bundle-wide shaper rate.

This example matches traffic by **traffic-class** values and applies absolute shaper values in units of Gbps, Mbps, and kbps. Replace the class names, policy-map name, bundle interface, traffic classes, and shaper rates with values for your network.

Percentage-based shaping remains supported on bundle interfaces. This example uses only absolute shaper values to focus on absolute-value shaping.

Before you begin

- Verify that your router supports absolute-value shaping on bundle interfaces.
- Ensure that the bundle interface and its member interfaces are configured.
- Ensure that each absolute shaper value is less than or equal to the bandwidth of each bundle member interface.
- Identify the traffic classes and fixed egress rates that you want to shape.

Follow these steps to configure an absolute-value shaper on a bundle interface.

Procedure

Step 1 Enter global configuration mode.

Example:

```
Router# configure
```

Step 2 Create class maps to identify the traffic classes that require shaping.

Example:

```
Router(config)#class-map match-any bundle_tc1
Router(config-cmap)#match traffic-class 1
Router(config-cmap)#exit
Router(config)#class-map match-any bundle_tc2
Router(config-cmap)#match traffic-class 2
Router(config-cmap)#exit
```

```
Router(config)#class-map match-any bundle_tc3
Router(config-cmap)#match traffic-class 3
Router(config-cmap)#exit
```

This example uses **traffic-class** values as the match criteria. Use the match criteria that apply to your QoS design.

The class maps identify the traffic classes to which the shaping policy applies.

Step 3 Create a policy map and configure absolute shaper rates for the traffic classes.

Example:

```
Router(config)#policy-map bundle_abs_shape
Router(config-pmap)#class bundle_tc1
Router(config-pmap-c)#shape average 500 mbps
Router(config-pmap-c)#exit
Router(config-pmap)#class bundle_tc2
Router(config-pmap-c)#shape average 1000 mbps
Router(config-pmap-c)#exit
Router(config-pmap)#class bundle_tc3
Router(config-pmap-c)#shape average 1500 mbps
Router(config-pmap-c)#exit
Router(config-pmap)#class class-default
Router(config-pmap-c)#exit
Router(config-pmap)#end-policy-map
```

Use the **shape average** command with the absolute unit that matches the rate you want to configure. This example uses Mbps consistently for readability.

The policy map defines absolute egress shaping rates for the matched traffic classes.

Step 4 Apply the policy map to the bundle interface in the egress direction.

Example:

```
Router(config)#interface Bundle-Ether30
Router(config-if)#service-policy output bundle_abs_shape
```

The output direction is the egress direction. The router replicates the egress shaping policy to each bundle member interface.

Step 5 Commit the configuration.

Example:

```
Router(config-if)#commit
Router(config-if)#end
```

The router applies the QoS policy to the bundle interface.

Step 6 Verify that the QoS policy is programmed on the bundle interface and on the member interface.

Example:

```
Router#show qos interface Bundle-Ether30 output
NOTE: Configured values are displayed within parentheses
Node 0/7/CPU0, Interface Bundle-Ether30 Ifh 0x7800011c (BundleMain) -- output policy
NPU Id:                        1
Total number of classes:       4
Interface Bandwidth:           100000000 kbps
Policy Name:                   bundle_abs_shape
Accounting Type:               Layer1 (Include Layer 1 encapsulation and above)
------------------------------------------------------------------------------
Level1 Class                             =   bundle_tc1

Level1 Class                             =   bundle_tc2
```

```
Level1 Class                             =   bundle_tc3

Level1 Class                             =   class-default

Interface HundredGigE0/7/0/11 Ifh 0x1c000240 (Member) -- output policy
NPU Id:                        1
Total number of classes:       4
Interface Bandwidth:           100000000 kbps
Policy Name:                   bundle_abs_shape
VOQ Base:                      41600
Accounting Type:               Layer1 (Include Layer 1 encapsulation and above)
VOQ Mode:                      8
Shared Counter Mode:           1
------------------------------------------------------------------------------
Level1 Class                             =   bundle_tc1
Egressq Queue ID                         =   41601 (LP queue)
Queue Max. BW.                           =   500000 kbps (500 mbits/sec)
Guaranteed service rate                  =   500000 kbps

Level1 Class                             =   bundle_tc2
Egressq Queue ID                         =   41602 (LP queue)
Queue Max. BW.                           =   1000000 kbps (1000 mbits/sec)
Guaranteed service rate                  =   1000000 kbps

Level1 Class                             =   bundle_tc3
Egressq Queue ID                         =   41603 (LP queue)
Queue Max. BW.                           =   1500000 kbps (1500 mbits/sec)
Guaranteed service rate                  =   1500000 kbps
```

The output confirms that the output QoS policy is programmed on the bundle member interface and that each configured traffic class has its configured shaping rate. The configured values appear within parentheses.

The output confirms that the shaping policy has been replicated to the member interface.

Step 7 Verify the policy-map statistics on the bundle interface.

Example:

```
Router#show policy-map type qos interface Bundle-Ether30 output pmap-name bundle_abs_shape

Bundle-Ether30 output: bundle_abs_shape

Class bundle_tc1
  Classification statistics          (packets/bytes)     (rate - kbps)
    Matched             :             1200000/921600000             600000
    Transmitted         :             1000000/768000000             500000
    Total Dropped       :              200000/153600000             100000
  Queueing statistics
    Queue ID                           : None (Bundle)
    Taildropped(packets/bytes)         : 200000/153600000

Class bundle_tc2
  Classification statistics          (packets/bytes)     (rate - kbps)
    Matched             :             2400000/1843200000           1200000
    Transmitted         :             2000000/1536000000           1000000
    Total Dropped       :              400000/307200000             200000
  Queueing statistics
    Queue ID                           : None (Bundle)
    Taildropped(packets/bytes)         : 400000/307200000

Class bundle_tc3
  Classification statistics          (packets/bytes)     (rate - kbps)
    Matched             :             3600000/2764800000           1800000
    Transmitted         :             3000000/2304000000           1500000
```

```
    Total Dropped       :               600000/460800000            300000
  Queueing statistics
    Queue ID                             : None (Bundle)
    Taildropped(packets/bytes)           : 600000/460800000
```

The output shows that packets are matching the configured classes and that traffic is being transmitted according to the configured shaping policy. Use these counters to verify that the policy is active and handling traffic as expected.

Step 8 Verify the configured policy definition.

Example:

```
Router#show policy-map pmap-name bundle_abs_shape detail
class-map match-any bundle_tc1
 match traffic-class 1
 end-class-map
!
class-map match-any bundle_tc2
 match traffic-class 2
 end-class-map
!
class-map match-any bundle_tc3
 match traffic-class 3
 end-class-map
!
policy-map bundle_abs_shape
 class bundle_tc1
  shape average 500 mbps
 !
 class bundle_tc2
  shape average 1000 mbps
 !
 class bundle_tc3
  shape average 1500 mbps
 !
 class class-default
 !
 end-policy-map
!
```

The policy definition confirms that the configured **shape average** values match the values that you configured in the policy map.

The bundle interface has an egress QoS policy that shapes matched traffic classes to fixed absolute rates. The configured shaper rate is applied per bundle member interface.

What to do next

After you add or replace bundle members, verify the bundle state using the **show bundle Bundle-Ether 30** command.

- If the command output displays the warning `MQC: Shape bandwidth rate is greater than the parent reference bandwidth`, remove the lower-bandwidth member, reduce the shaper rate so that it fits within the member bandwidth, and add the member again.

Restrictions for QoS policies on link bundles

Interface and link bundling

- Link bundling is supported only on Ethernet interfaces. Other interface types cannot participate in a bundle.
- Each bundle must contain only physical links. You cannot add subinterfaces or virtual interfaces to a link bundle.
- All links within a bundle must run the same bundling protocol—either IEEE 802.3ad (LACP) or EtherChannel (non-LACP). Mixing protocols within a single bundle is not supported.

QoS shaping and accounting

- From Cisco IOS XR Release 26.3.1, absolute shaper values are supported on bundle interfaces on supported routers. On unsupported releases or routers, use percentage-based shaper values.
- The absolute shaper value must be less than or equal to the bandwidth of each bundle member interface.
- Ethernet link bundles do not support MAC-level accounting. To collect traffic statistics, use higher-layer QoS counters instead.

Scalability limits

- Each bundle can include a maximum of 64 Ethernet links.
- The maximum number of link bundles supported is 128.

Configure QoS policies on link bundles

Before you begin

When you attach a QoS policy to a bundle interface, the policy is inherited by member interfaces as applicable. The enforcement behavior depends on the QoS action, direction, and supported platform. Review the bundle-specific usage guidelines for shaping, queuing, and policing behavior before you configure the policy. Complete these prerequisites before you begin this procedure.

- Prepare your class and policy definitions.

  For more information on how to configure class maps and policy maps, see *Classify Packets to Identify Specific Traffic* in the *Modular QoS Configuration Guide for Cisco 8000 Series Router*.

- Create the link bundle.

Procedure

Step 1 Configure an Ethernet link bundle and attach the traffic policy to it.

Example:

```
Router(config)#interface Bundle-Ether 12000
Router(config-if)#mtu 9100
Router(config-if)#service-policy input ingress
Router(config-if)#service-policy output egress
Router(config-if)#ipv4 address 100.12.0.0 255.255.255.254
Router(config-if)#bundle maximum-active links 64
Router(config-if)#commit
```

This running configuration example shows how a queuing policy is applied on an Ethernet link bundle. The policy is applied to all interfaces that are members of the Ethernet link bundle.

```
/* Policy-map */
policy-map ingress
class inet4-classifier-af1
set qos-group 1
!
class inet4-classifier-af2
set qos-group 2
!
class inet4-classifier-af3
set qos-group 3
!
class inet4-classifier-af4
set qos-group 4
!
class inet4-classifier-be1
set qos-group 5
!
class inet4-classifier-nc1
set qos-group 6
!
class class-default
!
end-policy-map
!
/* Ether Bundle */
interface Bundle-Ether12000
mtu 9100
service-policy input ingress
service-policy output egress
ipv4 address 100.12.0.0 255.255.255.254
load-interval 30
!
```

Step 2 Verify that the bundle status is up.

Example:

```
router#show bundle bundle-ether 1200
Wed Dec 16 19:55:49.974 PST
Bundle-Ether12000
Status: Up
Local links <active/standby/configured>: 35 / 0 / 35
Local bandwidth <effective/available>: 3500000000 (3500000000) kbps
MAC address (source): ea3b.745f.c4b0 (Chassis pool)
Inter-chassis link: No
Minimum active links / bandwidth: 1 / 1 kbps
Maximum active links: 64
Wait while timer: 2000 ms
Load balancing: Default
LACP: Operational
Flap suppression timer: Off
Cisco extensions: Disabled
Non-revertive: Disabled
mLACP: Not configured
```

```
IPv4 BFD: Not configured
Port Device State Port ID B/W, kbps
------------------- --------------- ----------- -------------- ----------
Hu0/4/0/0 Local Active 0x8000, 0x0009 100000000
Link is Active
Hu0/4/0/1 Local Active 0x8000, 0x000a 100000000
Link is Active
- - -
- - -
Hu0/4/0/35 Local Active 0x8000, 0x002b 100000000
Link is Active
```

Step 3 Verify the bundle statistics.

Example:

```
router#show policy-map interface bundle-ether 12000
Bundle-Ether12000 input: ingress
Class inet4-classifier-af1
Classification statistics (packets/bytes) (rate - kbps)
Matched : 4647401962/21236124455654 26403040
Transmitted : 4647401962/21236124455654 26403040
Total Dropped : 0/0 0
Class inet4-classifier-af2
Classification statistics (packets/bytes) (rate - kbps)
Matched : 4502980177/20576584333939 25571493
Transmitted : 4502980177/20576584333939 25571493
Total Dropped : 0/0 0
Class inet4-classifier-af3
Classification statistics (packets/bytes) (rate - kbps)
Matched : 4647404125/21236213667880 26389086
Transmitted : 4647404125/21236213667880 26389086
Total Dropped : 0/0 0
Class inet4-classifier-af4
Classification statistics (packets/bytes) (rate - kbps)
Matched : 9291188840/42456120548683 52771168
Transmitted : 9291188840/42456120548683 52771168
Total Dropped : 0/0 0
Class inet4-classifier-be1
Classification statistics (packets/bytes) (rate - kbps)
Matched : 4647413429/21235847852686 26393414
Transmitted : 4647413429/21235847852686 26393414
Total Dropped : 0/0 0
Class inet4-classifier-nc1
Classification statistics (packets/bytes) (rate - kbps)
Matched : 9294887621/42473100149807 52778258
Transmitted : 9294887621/42473100149807 52778258
Total Dropped : 0/0 0
Class class-default
Classification statistics (packets/bytes) (rate - kbps)
Matched : 0/0 0
Transmitted : 0/0 0
Total Dropped : 0/0 0
Bundle-Ether12000 output: egress
Class c1
Classification statistics (packets/bytes) (rate - kbps)
Matched : 16665494532/75878118942463 8760591
Transmitted : 16655834643/75834136022017 8760591
Total Dropped : 9659889/43982920446 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 9659889/43982920446
Class c2
Classification statistics (packets/bytes) (rate - kbps)
Matched : 16665421959/75877849543188 8718687
```

```
Transmitted : 16665421959/75877849543188 8718687
Total Dropped : 0/0 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 0/0
Class c3
Classification statistics (packets/bytes) (rate - kbps)
Matched : 16665247833/75877509455458 8703470
Transmitted : 16665187414/75877234624197 8703470
Total Dropped : 60419/274831261 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 60419/274831261
Class c4
Classification statistics (packets/bytes) (rate - kbps)
Matched : 33330896131/151755393012945 17470745
Transmitted : 33330745421/151754709368565 17470745
Total Dropped : 150710/683644380 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 150710/683644380
Class c5
Classification statistics (packets/bytes) (rate - kbps)
Matched : 16878910340/76849791869834 8833394
Transmitted : 16878849464/76849514633309 8833394
Total Dropped : 60876/277236525 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 60876/277236525
Class c6
Classification statistics (packets/bytes) (rate - kbps)
Matched : 33330898844/151756094112925 17456785
Transmitted : 33330752668/151755427708382 17456785
Total Dropped : 146176/666404543 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 146176/666404543
Class c7
Classification statistics (packets/bytes) (rate - kbps)
Matched : 244106/79922040 74
Transmitted : 244106/79922040 74
Total Dropped : 0/0 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 0/0
Class class-default
Classification statistics (packets/bytes) (rate - kbps)
Matched : 267075066180/1215993441123215 139917482
Transmitted : 267075066180/1215993441123215 139917482
Total Dropped : 0/0 0
Queueing statistics
Queue ID : None (Bundle)
Taildropped(packets/bytes) : 0/0
```