Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast

First Published: February 27, 2007
Last Updated: June 20, 2007

The Per Interface Mroute State Limit feature provides the capability to limit the amount of multicast route (mroute) states on an interface for different access control list (ACL)-classified sets of multicast traffic. This feature can be used to prevent denial-of-service (DoS) attacks, or to provide a multicast Call Admission Control (CAC) mechanism in network environments where all the multicast flows roughly utilize the same amount of bandwidth.

The Bandwidth-Based CAC for IP Multicast feature enhances the Per Interface Mroute State Limit feature by implementing a way to count per interface mroute state limiters using cost multipliers. This feature can be used to provide bandwidth-based CAC on a per interface basis in network environments where the multicast flows utilize different amounts of bandwidth.

Finding Feature Information in This Module

Your Cisco IOS software release may not support all of the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To reach links to specific feature documentation in this module and to see a list of the releases in which each feature is supported, use the "Feature Information for Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast" section on page 19.

Finding Support Information for Platforms and Cisco IOS and Catalyst OS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS and Catalyst OS software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Contents

Information About Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast

Before you configure the Per Interface Mroute State Limit and Bandwidth-Based CAC for IP Multicast features, you should understand the following concepts:

Overview of Per Interface Mroute State Limit

The Per Interface Mroute State Limit feature provides the capability to limit the number of mroute states on an interface for different ACL-classified sets of multicast traffic. This feature can be used to prevent DoS attacks, or to provide a multicast CAC mechanism when all the multicast flows roughly utilize the same amount of bandwidth.

The Per Interface Mroute State Limit feature essentially is a complete superset of the IGMP State Limit feature (with the exception that it does not support a global limit). The Per Interface Mroute State Limit feature, moreover, is more flexible and powerful (albeit more complex) than the IGMP State Limit feature but is not intended to be a replacement for it because there are applications that suit both features.

Note The Per Interface Mroute State Limit feature can be used in conjunction with the IGMP State Limit feature. If both the Per Interface Mroute State Limit feature and IGMP State Limit feature are configured on an interface, the Cisco IOS software enforces both limits.

For more information about the IGMP State Limit feature, see the "Customizing IGMP" chapter in the *Cisco IOS IP Multicast Configuration Guide*, Release 12.4T.

The main differences between the Per Interface Mroute State Limit feature and the IGMP State Limit feature are as follows:

- The Per Interface Mroute State Limit feature allows multiple limits to be configured on an interface, whereas the IGMP State Limit feature allows only one limit to be configured on an interface. The Per Interface Mroute State Limit feature, thus, is more flexible than the IGMP State Limit feature in that it allows multiple limits to be configured for different sets of multicast traffic on an interface.
- The Per Interface Mroute State Limit feature can be used to limit both Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) joins, whereas the IGMP State Limit feature can only be used to limit IGMP joins. The IGMP State Limit feature, thus, is more limited in application in that it is best suited to be configured on an edge router to limit the number of groups that receivers can join on an outgoing interface. The Per Interface Mroute State Limit feature has a wider application in that it can be configured to limit IGMP joins on an outgoing interface, to limit PIM joins (for Any Source Multicast [ASM] groups or Source Specific Multicast [SSM] channels) on an outgoing interface connected to other routers, to limit sources behind an incoming interface from sending multicast traffic, or to limit sources directly connected to an incoming interface from sending multicast traffic.

Note Although the PIM Interface Mroute State Limit feature allows you to limit both IGMP and PIM joins, it does not provide the capability to limit PIM or IGMP joins separately because it does not take into account whether state is created as a result of an IGMP or PIM join. As such, the IGMP State Limit feature is more specific in application because it specifically limits IGMP joins.

- The Per Interface Mroute State Limit feature allows you to specify limits according to the direction of traffic; that is, it allows you to specify limits for outgoing interfaces, incoming interfaces, and for incoming interfaces having directly connected multicast sources. The IGMP State Limit feature, however, only can be used to limit outgoing interfaces. The Per Interface State Mroute State Limit feature, thus, is wider in scope in that it can be used to limit mroute states for both incoming and outgoing interfaces from both sources and receivers, whereas the IGMP State Limit feature is more narrow in scope in that it can only be used to limit mroute states for receivers on an LAN by limiting the number of IGMP joins on an outgoing interface.

Both the IGMP State Limit and Per Interface Mroute State Limit features provide a rudimentary multicast CAC mechanism that can be used to provision bandwidth utilization on an interface when all multicast flows roughly utilize the same amount of bandwidth. The Bandwidth-Based CAC for IP Multicast feature, however, offers a more flexible and powerful alternative for providing multicast CAC in network environments where IP multicast flows utilize different amounts of bandwidth.

Note For more information about the Bandwidth-Based CAC for IP Multicast feature, see the "Overview of Bandwidth-Based CAC for IP Multicast" section on page 5.

Per Interface Mroute State Limit Feature Design

The Per Interface Mroute State Limit feature is configured using the **ip multicast limit** command in interface configuration mode. An **ip multicast limit** command configured on an interface is called an *mroute state limiter*. An mroute state limiter is defined by direction, ACL, and maximum number of mroutes. Each mroute state limiter maintains a counter to ensure that the maximum number of mroutes is not exceeded.

The following forms of the **ip multicast limit** command are available to configure mroute state limiters on an interface:

- **ip multicast limit** *access-list max-entries*

  Limits mroute state creation for an ACL-classified set of traffic on an interface when the interface is an outgoing (egress) interface, and limits mroute outgoing interface list (olist) membership when the interface is an incoming (ingress) Reverse Path Forwarding (RPF) interface.

  This type of mroute state limiter limits mroute state creation—by accounting each time an mroute permitted by the ACL is created or deleted—and limits mroute olist membership—by accounting each time that an mroute olist member permitted by the ACL is added or removed.

  Entering this form of the command (that is, with no optional keywords) is equivalent to specifying the **ip multicast limit rpf** and **ip multicast limit out** forms of the command.

- **ip multicast limit connected** *access-list max-entries*

  Limits mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface that is directly connected to a multicast source by accounting each time that an mroute permitted by the ACL is created or deleted.

- **ip multicast limit out** *access-list max-entries*

  Limits mroute olist membership on an outgoing interface for an ACL-classified set of multicast traffic by accounting each time that an mroute olist member permitted by the ACL is added or removed.

- **ip multicast limit rpf** *access-list max-entries*

  Limits the number of mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface by accounting each time an mroute permitted by the ACL is created or deleted.

ACLs are used with the **ip multicast limit** command to define the IP multicast traffic to be limited on an interface. Standard ACLs can be used to define the (*, G) state to be limited on an interface. Extended ACLs can be used to define the (S, G) state to be limited on an interface. Extended ACLs also can be used to define the (*, G) state to be limited on an interface, by specifying 0.0.0.0 for the source address and source wildcard—referred to as (0, G)—in the permit or deny statements that compose the extended access list.

Mechanics of the Per Interface Mroute State Limit Feature

The mechanics of the Per Interface Mroute State Limit feature are as follows:

- Each time the state for an mroute is created or deleted and each time an olist member is added or removed, the Cisco IOS software searches for a corresponding mroute state limiter that matches the mroute.
- In the case of the creation and deletion of mroutes, the Cisco IOS software searches for an mroute state limiter configured on the incoming (RPF) interface that matches the mroute to be created or deleted. In the case of olist member addition or removal, the Cisco IOS software searches for an mroute state limiter configured on the outgoing interface that matches the mroute to be added or removed.
- The Cisco IOS software performs a top-down search from the list of configured limiters on the interface. Only limiters that match the direction of traffic are considered. The first mroute state limiter that matches is used for limiting (sometimes referred to as *accounting*). A match is found when the ACL permits the mroute state.

- When a match is found, the counter of the mroute state limiter is updated (increased or decreased). If no mroute state limiter is found that matches an mroute, no accounting is performed for the mroute (because there is no counter to update).
- The amount to update the counter with is called the *cost* (sometimes referred to as the *cost multiplier*). The default cost is 1.

Note An mroute state limiter always allows the deletion of an mroute or the removal of an interface from the olist. In those cases, the respective mroute state limiter decreases the counter by the value of the cost multiplier. In addition, RPF changes to an existing mroute are always allowed (in order to not affect existing traffic). However, an mroute state limiter *only* allows the creation of an mroute or the addition of an mroute olist member if adding the cost does not exceed the maximum number of mroutes permitted.

Tips for Configuring Mroute State Limiters

- To ensure that all mroutes are accounted, you can configure an mroute state limiter whose ACL contains a **permit any** statement and set the maximum for the *max-entries* argument to 0. Configuring an mroute state limiter in this manner effectively denies all fall through states, which may be a way to prevent a multicast DoS attack in and out of the interface.
- When creating an ACL, remember that, by default, the end of the ACL contains an implicit **deny any** statement for everything if it did not find a match before reaching the end.
- An explicit deny statement for a specific mroute in an ACL can be used to specify the state that will not match the ACL (thus, preventing the ACL from being accounted). If an mroute matches a deny statement, the search immediately continues to the next configured mroute state limiter. Configuring an explicit deny statement in an ACL can be more efficient than forcing the mroute to fall through an ACL (by means of the implicit **deny any** statement at the end of the ACL).

Overview of Bandwidth-Based CAC for IP Multicast

The Bandwidth-Based CAC for IP Multicast feature enhances the Per Interface Mroute State Limit feature by implementing a way to count per interface mroute state limiters using cost multipliers. This feature can be used to provide bandwidth-based multicast CAC on a per interface basis in network environments where the multicast flows utilize different amounts of bandwidth.

Bandwidth-Based CAC for IP Multicast Feature Design

Bandwidth-Based CAC policies are configured using the **ip multicast limit cost** command in global configuration mode. The syntax of the **ip multicast limit cost** command is as follows:

ip multicast limit cost *access-list cost-multiplier*

ACLs are used with this command to define the IP multicast traffic for which to apply a cost. Standard ACLs can be used to define the (*, G) state. Extended ACLs can be used to define the (S, G) state. Extended ACLs also can be used to define the (*, G) state, by specifying 0.0.0.0 for the source address and source wildcard—referred to as (0, G)—in the permit or deny statements that compose the extended access list.

The Bandwidth-Based CAC for IP Multicast feature enhances the Per Interface Mroute State Limit feature by providing the capability to define costs (globally or per multicast VPN routing and forwarding [MVRF] instance) to be applied to mroutes that are being limited. The *cost-multiplier* argument is used to specify the cost to apply to mroutes that match the ACL specified for the *access-list* argument.

Mechanics of the Bandwidth-Based CAC for IP Multicast Feature

The mechanics of the Bandwidth-Based CAC for IP Multicast feature are as follows:

- Once an mroute matches an ACL configured for an mroute state limiter, the Cisco IOS software performs a top-down search from the global or per MVRF list of configured bandwidth-based CAC policies to determine if a cost should be applied to the mroute.
- A cost is applied to the first bandwidth-based CAC policy that matches the mroute. A match is found when the ACL applied to the bandwidth-based CAC policy permits the mroute state.
- The counter for the mroute state limiter either adds or subtracts the cost configured for the *cost-multiplier* argument. If no costs are configured or if the mroute does not match any of the configured bandwidth-based CAC polices, the default cost of 1 is used.

Tips for Configuring Bandwidth-Based CAC Policies for IP Multicast

- To ensure that a particular cost applies to all mroutes being limited, you can configure a bandwidth-based CAC policy whose ACL contains a **permit any** statement. Configuring a bandwidth-based CAC policy in this manner effectively ensures that the default cost is not applied to any mroutes being limited.
- Configuring a bandwidth-based CAC policy with a cost of 0 for the *cost-multiplier* argument can be used to skip the accounting of certain mroutes (for example, to prevent Auto-RP groups or a specific multicast channel from being accounted).
- An explicit deny statement for a specific mroute in an ACL can be used to specify the state that will not match the ACL (thus, preventing the ACL from being accounted). If an mroute matches a deny statement, the search immediately continues to the next configured bandwidth-based CAC policy. Configuring an explicit deny statement in an ACL can be more efficient than forcing the mroute to fall through an ACL (by means of the implicit **deny any** statement at the end of the ACL).

How to Configure Per Interface Mroute State Limiters and Bandwidth-Based CAC Policies for IP Multicast

This section contains the following procedures:

- Configuring Per Interface Mroute State Limiters, page 7 (required)
- Configuring Per Interface Mroute State Limiters with Bandwidth-Based CAC Policies for IP Multicast, page 9 (required)
- Monitoring Per Interface Mroute State Limiters and Bandwidth-Based CAC Policies for IP Multicast, page 11 (optional)

Configuring Per Interface Mroute State Limiters

Perform this task to configure per interface mroute state limiters. Configuring mroute state limiters can be used to prevent DoS attacks, or to provide a multicast Call Admission Control (CAC) mechanism to control bandwidth, when all the multicast flows roughly utilize the same amount of bandwidth.

Prerequisites

All ACLs you intend to use for configuring limiters should be configured prior to beginning this configuration task; otherwise, the limiters are ignored. For information about how to configure ACLs, see the "Creating an IP Access List and Applying It to an Interface" module.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ip multicast limit** [**connected** | **out** | **rpf**] *access-list max-entries*
5. Repeat Step 4, if you want to configure additional mroute state limiters on the interface.
6. Repeat Step 3 and 4, if you want to configure mroute state limiters on additional interfaces.
7. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Router(config)# interface FastEthernet 1`	Enters interface configuration mode for the specified interface type and number.

	Command or Action (continued)	Purpose (continued)
Step 4	**ip multicast limit** [**connected** | **out** | **rpf**] *access-list max-entries* **Example:** `Router(config-if)# ip multicast limit 15 100`	Configures mroute state limiters on an interface. • Specify the **ip multicast limit** command with no optional keywords to limit mroute state creation for an ACL-classified set of traffic on an interface when the interface is an outgoing (egress) interface, and to limit mroute olist membership when the interface is an incoming (ingress) Reverse Path Forwarding (RPF) interface. – This type of mroute state limiter limits mroute state creation by accounting each time an mroute permitted by the ACL is created or deleted and limits mroute olist membership by accounting each time that an mroute olist member permitted by the ACL is added or removed. – Entering this form of the command (that is, with no optional keywords) is equivalent to specifying the **ip multicast limit rpf** and **ip multicast limit out** forms of the command. • Use the **connected** keyword to configure an mroute state limiter that limits mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface that is directly connected to a multicast source by accounting each time that an mroute permitted by the ACL is created or deleted. • Use the **out** keyword to configure an mroute state limiter that limits mroute olist membership on an outgoing interface for an ACL-classified set of multicast traffic by accounting each time that an mroute olist member permitted by the ACL is added or removed. • Use the **rpf** keyword to configure an mroute state limiter that limits the number of mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface by accounting each time an mroute permitted by the ACL is created or deleted. • The range of values that can be entered for the *max-entries* argument is 0 to 2147483647.
Step 5	Repeat Step 4, if you want to configure additional mroute state limiters on the interface.	—
Step 6	Repeat Step 3 and 4, if you want to configure mroute state limiters on additional interfaces.	—
Step 7	**end** **Example:** `Router(config-if)# end`	Exits interface configuration mode, and enters privileged EXEC mode.

Configuring Per Interface Mroute State Limiters with Bandwidth-Based CAC Policies for IP Multicast

Perform this task to apply costs to mroutes that are being limited by configuring bandwidth-based CAC policies. This task can be performed to provide bandwidth-based multicast CAC on a per interface basis in network environments where the multicast flows utilize different amounts of bandwidth. Bandwidth-based CAC policies can be applied globally or per MVRF.

Prerequisites

All ACLs you intend to use for configuring limiters should be configured prior to beginning this configuration task; otherwise, the limiters are ignored. For information about how to configure ACLs, see the "Creating an IP Access List and Applying It to an Interface" module.

Note You can omit Steps 3 to 6, if you have already configured the mroute state limiters for which to apply costs.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ip multicast limit** [**connected** | **out** | **rpf**] *access-list max-entries*
5. Repeat Step 4, if you want to configure additional mroute state limiters on the interface.
6. Repeat Step 3 and 4, if you want to configure mroute state limiters on additional interfaces.
7. **exit**
8. **ip multicast** [**vrf** *vrf-name*] **limit cost** *access-list cost-multiplier*
9. Repeat Step 8, if you want to apply additional costs to mroutes.
10. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action (continued)	Purpose (continued)
Step 3	**interface** *type number* **Example:** `Router(config)# interface FastEthernet 1`	Enters interface configuration mode for the specified interface type and number.
Step 4	**ip multicast limit** [**connected** | **out** | **rpf**] *access-list max-entries* **Example:** `Router(config-if)# ip multicast limit acl-test 100`	Configures mroute state limiters on an interface. • Specify the **ip multicast limit** command with no optional keywords to limit mroute state creation for an ACL-classified set of traffic on an interface when the interface is an outgoing (egress) interface, and to limit mroute olist membership when the interface is an incoming (ingress) Reverse Path Forwarding (RPF) interface. – This type of mroute state limiter limits mroute state creation by accounting each time an mroute permitted by the ACL is created or deleted and limits mroute olist membership by accounting each time that an mroute olist member permitted by the ACL is added or removed. – Entering this form of the command (that is, with no optional keywords) is equivalent to specifying the **ip multicast limit rpf** and **ip multicast limit out** forms of the command. • Use the optional **connected** keyword to configure an mroute state limiter that limits mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface that is directly connected to a multicast source by accounting each time that an mroute permitted by the ACL is created or deleted. • Use the optional **out** keyword to configure an mroute state limiter that limits mroute olist membership on an outgoing interface for an ACL-classified set of multicast traffic by accounting each time that an mroute olist member permitted by the ACL is added or removed. • Use the optional **rpf** keyword to configure an mroute state limiter that limits the number of mroute states created for an ACL-classified set of multicast traffic on an incoming (RPF) interface by accounting each time an mroute permitted by the ACL is created or deleted. • The range of values that can be entered for the *max-entries* argument is 0 to 2147483647.
Step 5	Repeat Step 4, if you want to configure additional mroute state limiters on the interface.	—
Step 6	Repeat Step 3 and 4, if you want to configure mroute state limiters on additional interfaces.	—

	Command or Action (continued)	Purpose (continued)
Step 7	**exit** **Example:** `Router(config-if)# exit`	Exits interface configuration mode, and returns to global configuration mode.
Step 8	**ip multicast** [**vrf** *vrf-name*] **limit cost** *access-list cost-multiplier* **Example:** `Router(config)# ip multicast limit cost acl-MP2SD-channels 4000`	Applies costs to mroutes state limiters. • Use the optional **vrf** keyword to specify that the cost be applied only to mroutes associated with MVRF specified for the *vrf-name* argument. • The range of values that can be entered for the *cost-multiplier* argument is 0 to 2147483647.
Step 9	Repeat Step 8, if you want to apply additional costs to mroutes.	—
Step 10	**end** **Example:** `Router(config-if)# end`	Exits interface configuration mode, and enters privileged EXEC mode.

Monitoring Per Interface Mroute State Limiters and Bandwidth-Based CAC Policies for IP Multicast

Perform this optional task to monitor per interface mroute state limiters and bandwidth-based CAC policies for IP multicast.

SUMMARY STEPS

1. **debug ip mrouting limits** [*group-address*]
2. **show ip multicast limit** [*type number*]
3. **clear ip multicast limit** [*type number*]

DETAILED STEPS

Step 1 **debug ip mrouting limits** [*group-address*]

Use this command to display debugging information for mroute state limiters configured on interfaces.

Specify the optional *group-address* argument to restrict the output to display only mroute state limiter events related to a particular multicast group.

The following output is from the **debug ip mrouting limits** command. The output displays the following events:

- An mroute state being created and the corresponding mroute state limiter counter being increased by the default cost of 1 on incoming Ethernet interface 1/0.
- An mroute olist member being removed from the olist and the corresponding mroute limiter being decreased by the default cost of 1 on outgoing Ethernet interface 1/0.
- An mroute being denied by the mroute state limiter because the maximum number of mroute states has been reached.

- An mroute state being created and the corresponding mroute state limiter counter being increased by the cost of 2 on incoming Ethernet interface 1/0.
- An mroute olist member being removed from the olist and the corresponding mroute limiter being decreased by a cost of 2 on outgoing Ethernet interface 1/0.

```
Router# debug ip mrouting limits

MRL(0): incr-ed acl ‘rpf-list’ to (13 < max 32), [n:0,p:0], (main) Ethernet1/0,
(10.41.0.41, 225.30.200.60)
MRL(0): decr-ed acl ‘out-list’ to (10 < max 32), [n:0,p:0], (main) Ethernet1/0, (*,
225.40.202.60)
MRL(0): Add mroute (10.43.0.43, 225.30.200.60) denied for Ethernet0/2, acl std-list, (16 =
max 16)
MRL(0): incr-ed limit-acl `rpf-list' to (12 < max 32), cost-acl 'cost-list' cost 2,
[n:0,p:0], (main) Ethernet1/0, (10.41.0.41, 225.30.200.60)
MRL(0): decr-ed limit-acl `out-list' to (8 < max 32), cost-acl 'cost-list'' cost 2,
[n:0,p:0], (main) Ethernet1/0, (*, 225.40.202.60)
```

Step 2 **show ip multicast limit** [*type number*]

Use this command to display the counters related to mroute state limiters configured on an interface.

Specify the optional *type number* arguments to restrict the output to only display information about the mroute state limiters configured for the specified interface.

When mroute state limiters are configured on interfaces, each time the state for an mroute is created or deleted and each time an olist member is added or removed, the counters that are displayed in the output of the **show ip multicast limit** command are increased or decreased accordingly. The output for this command also accounts for each time an mroute is denied due to an mroute state limit being reached.

The following is sample output from **show ip multicast limit** command with the *type number* arguments. In this example, information about mroute state limiters configured on Ethernet interface 1/0 is displayed.

```
Router# show ip multicast limit ethernet 1/0
Interface Ethernet1/0
 Multicast Access Limits
out acl out-list (1 < max 32) exceeded 0
rpf acl rpf-list (6 < max 32) exceeded 0
con acl conn-list (0 < max 32) exceeded 0
```

Step 3 **clear ip multicast limit** [*type number*]

Use this command to reset exceeded counters for mroute state limiters.

Specify the optional *type number* arguments to clear the mroute state limit exceeded counters for mroute state limiters configured on the specified interface.

The output of the **show ip multicast limit** command includes an “exceeded” counter for each configured mroute state limit. Each time an mroute is denied because the maximum number of mroutes for an mroute state limiter is reached, the “exceeded” counter is increased by a value of 1. Use the **clear ip multicast limit** command to reset exceeded counters for mroute state limiters.

The following example shows how to reset exceeded counters for mroute state limiters configured on Fast Ethernet interface 1:

```
clear ip multicast limit interface FastEthernet 1
```

Configuration Examples for Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast

This section provides the following configuration examples:

Configuring Per Interface Mroute State Limiters: Example

The following example shows how to configure mroute state limiters on interfaces to provide multicast CAC in a network environment where all the multicast flows roughly utilize the same amount of bandwidth.

This example uses the topology illustrated in Figure 1.

Figure 1 Per Interface Mroute State Limit Example Topology

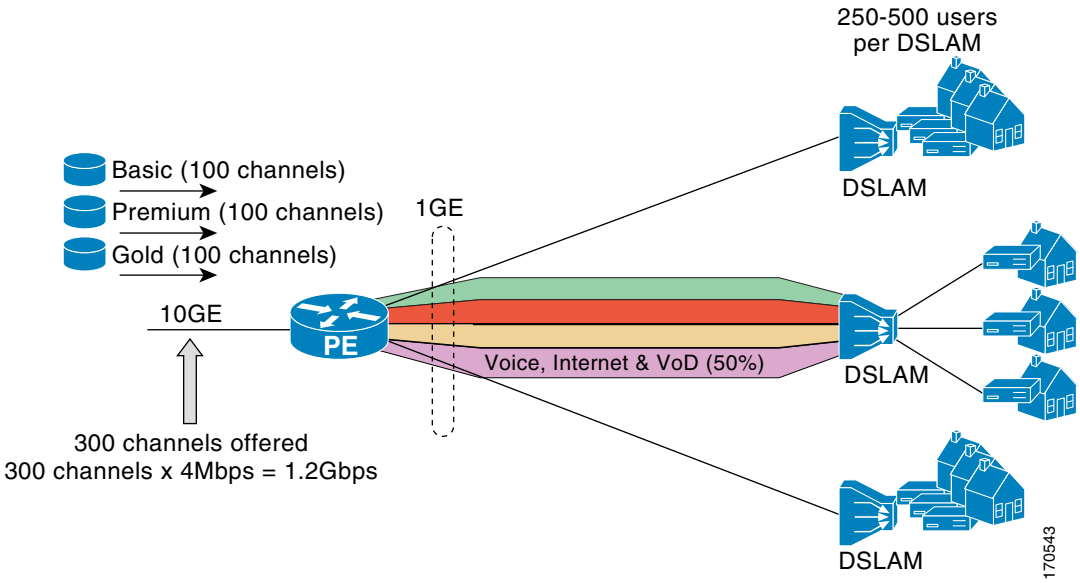

In this example, a service provider is offering 300 Standard Definition (SD) TV channels. The SD channels are being offered to customers in three service bundles (Basic, Premium, and Gold), which are available to customers on a subscription basis. Each bundle offers 100 channels to subscribers, and each channel utilizes approximately 4 Mbps of bandwidth.

The service provider must provision the Gigabit Ethernet interfaces on the provider edge (PE) router connected to Digital Subscriber Line Access Multiplexers (DSLAMs) as follows: 50% of the link's bandwidth (500 Mbps) must be available to subscribers of their Internet, voice, and video on demand (VoD) service offerings while the remaining 50% (500 Mbps) of the link's bandwidth must be available to subscribers of their SD channel bundle service offerings.

For the 500 Mbps of the link's bandwidth that must always be available to (but must never be exceeded by) the subscribers of the SD channel bundles, the interface must be further provisioned as follows:

- 60% of the bandwidth must be available to subscribers of the basic service (300 Mbps).
- 20% of the bandwidth must be available to subscribers of the premium service (100 Mbps).
- 20% of the bandwidth must be available to subscribers of the gold service (100 Mbps).

Because each SD channel utilizes the same amount of bandwidth (4 Mbps), the Per Interface Mroute State Limit feature can be used to provide the necessary CAC to provision the services being offered by the service provider. To determine the required CAC needed per interface, the number of channels for each bundle is divided by 4 (because each channel utilizes 4 Mbps of bandwidth). The required CAC needed per interface, therefore, is as follows:

- Basic Services: 300 / 4 = 75
- Premium Services: 100 / 4 = 25
- Gold Services: 100 / 4 = 25

Once the required CAC required per SD channel bundle is determined, the service provider uses the results to configure the mroute state limiters required to provision the Gigabit Ethernet interfaces on the PE router for the services being offered to subscribers behind the DSLAMs:

- For the Basic Services bundle, the service provider must limit the number of Basic Service SD channels that can be transmitted out a Gigabit Ethernet interface (at any given time) to 75. Configuring an mroute state limit of 75 for the SD channels offered in the Basic Service bundle provisions the interface for 300 Mbps of bandwidth (the 60% of the link's bandwidth that must always be available to [but never exceeded by] the subscribers of the Basic Services bundle).
- For the Premium Services bundle, the service provider must limit the number of Premium Service SD channels that can be transmitted out a Gigabit Ethernet interface (at any given time) to 25. Configuring an mroute state limit of 25 for the SD channels offered in the Premium Service bundle provisions the interface for 100 Mbps of bandwidth (the 20% of the link's bandwidth that must always be available to [but never exceeded by] the subscribers of the Premium Service bundle).
- For the Gold Services bundle, the service provider must limit the number of Gold Service SD channels that can be transmitted out a Gigabit Ethernet interface (at any given time) to 25. Configuring an mroute state limit of 25 for the SD channels offered in the Gold Service bundle provisions the interface for 100 Mbps of bandwidth (the 20% of the link's bandwidth that must always be available to [but never exceeded by] the subscribers of the Gold Service bundle).

The service provider then configures three ACLs to be applied to mroute state limiters. Each ACL defines the SD channels for each SD channel bundle to be limited on an interface:

- acl-basic—The ACL that defines the SD channels offered in the basic service.
- acl-premium—The ACL that defines the SD channels offered in the premium service.
- acl-gold—The ACL that defines the SD channels offered in the gold service.

These ACLs are then applied to mroute state limiters configured on the PE router's Gigabit Ethernet interfaces.

For this example, three mroute state limiters are configured on Gigabit Ethernet interface 0/0 to provide the multicast CAC needed to provision the interface for the SD channel bundles being offered to subscribers:

- An mroute state limit of 75 for the SD channels that match acl-basic.
- An mroute state limit of 25 for the SD channels that match acl-premium.
- An mroute state limit of 25 for the SD channels that match acl-gold.

The following configuration shows how the service provider uses mroute state limiters to provision Gigabit Ethernet interface 0/0 for the SD channel bundles and Internet, Voice, and VoD services being offered to subscribers:

```
interface GigabitEthernet0/0
description --- Interface towards the DSLAM ---
.
.
.
ip multicast limit out acl-basic 75
ip multicast limit out acl-premium 25
ip multicast limit out acl-gold 25
```

Configuring Per Interface Mroute State Limiters with Bandwidth-Based CAC Policies for IP Multicast: Example

The following example shows how to configure per interface mroute state limiters with bandwidth-based CAC policies to provide multicast CAC in a network environment where the multicast flows utilize the different amounts of bandwidth.

This example uses the topology illustrated in Figure 2.

Figure 2 Bandwidth-Based CAC for IP Multicast Example Topology

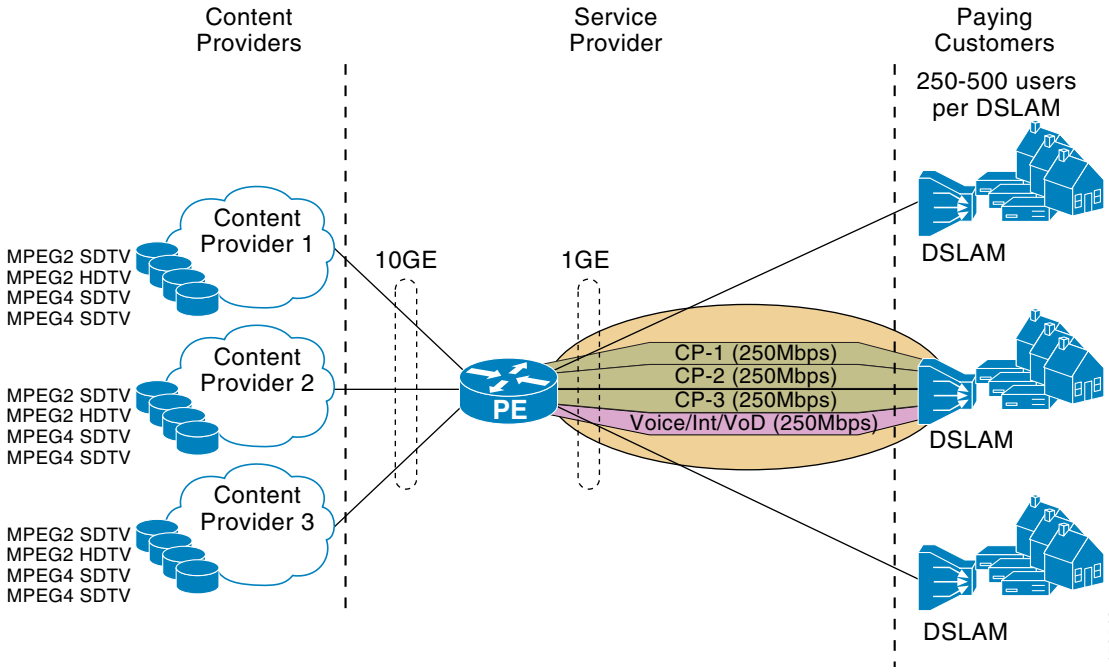

In this example, three content providers are providing TV services across a service provider core. The content providers are broadcasting TV channels that utilize different amounts of bandwidth:

- MPEG-2 SDTV channels—4 Mbps per channel.
- MPEG-2 HDTV channels—18 Mbps per channel.
- MPEG-4 SDTV channels—1.6 Mbps per channel.
- MPEG-4 HDTV channels—6 Mbps per channel.

The service provider needs to provision the fair sharing of bandwidth between these three content providers to its subscribers across Gigabit Ethernet interfaces. The service provider, thus, determines that it needs to provision each Gigabit Ethernet interface on the PE router connected to the DSLAMs as follows:

- 250 Mbps per content provider.
- 250 Mbps for Internet, voice, and VoD services.

The service provider then configures three ACLs:

- acl-CP1-channels—The ACL that defines the channels being offered by the content provider CP1.
- acl-CP2-channels—The ACL that defines the channels being offered by the content provider CP2.
- acl-CP3-channels—The ACL that defines the channels being offered by the content provider CP3.

Because the content providers are broadcasting TV channels that utilize different amounts of bandwidth, the service provider needs to determine the values that need to be configured for the mroute state limiters and bandwidth-based CAC policies to provide the fair sharing of bandwidth required between the content providers.

Prior to the introduction of the Bandwidth-Based CAC for IP Multicast feature, the mroute state limiters were based strictly on the number of flows. The introduction of cost multipliers by the Bandwidth-Based CAC for IP Multicast feature expands how mroute state limiters can be defined. Instead of defining the mroute state limiters based on the number of multicast flows, the service provider looks for a common unit of measure and decides to represent the mroute state limiters in Kbps. The service provider then configures three mroute states, one mroute state limiter per content provider. Because the link is a Gigabit, the service provider sets each limit to 250000 (because 250000 Kbps equals 250 Mbps, the number of bits that service provider needs to provision per content provider).

The service provider needs to further provision the fair sharing of bandwidth between the content providers, which can be achieved by configuring bandwidth-based CAC policies. The service provider decides to create four bandwidth-based CAC policies, one policy per channel based on bandwidth. For these policies, the service provider configures the following ACLs:

- acl-MP2SD-channels—Defines all the MPEG-2 SD channels offered by the three content providers.
- acl-MP2HD-channels—Defines all the MPEG-2 HD channels offered by the three content providers.
- acl-MP4SD-channels—Defines all the MPEG-4 SD channels offered by the three content providers.
- acl-MP4HD-channels—Defines all the MPEG-4 HD channels offered by the three content providers.

For each policy, a cost multiplier (represented in Kbps) is defined for each ACL that is based on the bandwidth of the channels defined in the ACL:

- 4000—Represents the 4 Mbps MPEG-2 SD channels.
- 18000—Represents the 18 Mbps MPEG-2 HD channels.
- 1600—Represents the 1.6 Mbps MPEG-4 SD channels.
- 6000—Represents the 6 Mbps MPEG-4 HD channels.

The following configuration example shows how the service provider used mroute state limiters with bandwidth-based CAC policies to provision Gigabit Ethernet interface 0/0 for the fair sharing of bandwidth required between the three content providers:

```
!
ip multicast limit cost acl-MP2SD-channels 4000
ip multicast limit cost acl-MP2HD-channels 18000
ip multicast limit cost acl-MP4SD-channels 1600
ip multicast limit cost acl-MP4HD-channels 6000
!
.
.
.
!
interface GigabitEthernet0/0
 ip multicast limit out acl-CP1-channels 250000
 ip multicast limit out acl-CP2-channels 250000
 ip multicast limit out acl-CP3-channels 250000
!
```

Additional References

The following sections provide references related to the Per Interface Mroute State Limit and Bandwidth-Based CAC for IP Multicast features.

Related Documents

Related Topic	Document Title
Multicast commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	*Cisco IOS IP Multicast Command Reference*

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
No new or modified MIBs are supported by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

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 Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/techsupport

Command Reference

The following commands are introduced or modified in the feature or features

- **clear ip multicast limit**
- **debug ip mrouting limits**
- **ip multicast limit**
- **ip multicast limit cost**
- **show ip multicast**

Feature Information for Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast

Table 1 lists the release history for this feature.

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Note Table 1 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 1 Feature Information for Per Interface Mroute State Limit with Bandwidth-Based CAC for IP Multicast

Feature Name	Releases	Feature Information
Per Interface Mroute State Limit	12.3(14)T 12.2(33)SRB	The Per Interface Mroute State Limit feature provides the capability to limit the number of mroute states on an interface for different ACL-classified sets of multicast traffic. This feature can be used to prevent DOS attacks, or to provide a multicast CAC mechanism in network environments where all the multicast flows roughly utilize the same amount of bandwidth. The following commands were introduced or modified by this feature: **clear ip multicast limit**, **debug ip mrouting limits**, **ip multicast limit**, **show ip multicast**.
Bandwidth-Based CAC for IP Multicast	12.2(33)SRB 12.4(15)T	The Bandwidth-Based CAC for IP Multicast enhances the Per Interface Mroute State Limit feature by implementing a way to count per interface mroute state limiters using cost multipliers. This feature can be used to provide bandwidth-based CAC on a per interface basis in network environments where the multicast flows utilize different amounts of bandwidth. The following command was introduced by this feature: **ip multicast limit cost**.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.