sequencing through show ip traffic-engineering configuration

sequencing

To configure the direction in which sequencing is enabled for data packets in a Layer 2 pseudowire, use the **sequencing** command in the appropriate configuration mode. To remove the sequencing configuration from the pseudowire class, use the **no** form of this command.

sequencing {**transmit**| **receive**| **both**| **resync** *number*}

no sequencing {**transmit**| **receive**| **both**| **resync** *number*}

Syntax Description

transmit	Updates the Sequence Number field in the headers of data packets sent over the pseudowire according to the data encapsulation method that is used.
receive	Keeps the value in the Sequence Number field in the headers of data packets received over the pseudowire. Out-of-order packets are dropped.
both	Enables both the **transmit** and **receive** options.
resync	Enables packet sequencing reset after the disposition router receives a specified number of out-of-order packets.
number	The number of out-of-order packets that cause reset of packet sequencing. The range is from 5 to 65535.

Command Default

Sequencing is disabled.

Command Modes

Interface configuration (config-if)

Pseudowire class configuration (config-pw-class)

Template configuration (config-template)

Command History

Release	Modification
12.0(23)S	This command was introduced for Layer 2 Tunnel Protocol Version 3 (L2TPv3).
12.3(2)T	This command was integrated into Cisco IOS Release 12.3(2)T.
12.0(29)S	This command was updated to support Any Transport over MPLS (AToM).
12.0(30)S	This command was modified. The **resync** keyword was added.

Release	Modification
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(27)SBC	This command was modified. L2TPv3 support for this command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(28)SB	This command was modified. AToM support for this command was integrated into Cisco IOS Release 12.2(28)SB.
Cisco IOS XE Release 3.7S	This command was integrated into a release prior to Cisco IOS XE Release 3.7S. This command was modified as part of the MPLS-based Layer 2 VPN (L2VPN) command modifications for cross-OS support and made available in interface configuration and template configuration modes in Cisco IOS XE Release 3.7S.
15.3(1)S	This command was integrated in Cisco IOS Release 15.3(1)S.

Usage Guidelines

When you enable sequencing using any available options, the sequence numbers are automatically sent and a request is sent to the remote provider edge (PE) peer for sequence numbers. Out-of-order packets that are received on the pseudowire are dropped only if you use the **sequencing receive** or **sequencing both** command.

If you enable sequencing for Layer 2 pseudowires on the Cisco 7500 series routers and use the **ip cef distributed** command, all traffic on the pseudowires is switched through the line cards.

Use the **resync** keyword when the disposition router receives many out-of-order packets. It allows the router to recover when too many out-of-order packets are dropped.

Examples

The following example shows how to enable sequencing in data packets in Layer 2 pseudowires that were created from the pseudowire class named ether-pw. The Sequence Number field is updated in tunneled packet headers for data packets that are both sent and received over the pseudowire:

```
Device(config)# pseudowire-class ether-pw
Device(config-pw-class)# encapsulation mpls
Device(config-pw-class)# sequencing both
```

The following example shows how to enable the disposition router to reset packet sequencing after it receives 1000 out-of-order packets:

```
Device(config)# pseudowire-class ether-pw
Device(config-pw-class)# encapsulation mpls
Device(config-pw-class)# sequencing both
Device(config-pw-class)# sequencing resync 1000
```

The following example shows how to enable the disposition router to reset packet sequencing after it receives 1000 out-of-order packets in interface configuration mode:

```
Device(config)# interface pseudowire 100
Device(config-if)# encapsulation mpls
Device(config-if)# sequencing both
Device(config-if)# sequencing resync 1000
```

The following example shows how to enable the disposition router to reset packet sequencing after it receives 1000 out-of-order packets in template configuration mode:

```
Device(config)# template type pseudowire template1
Device(config-template)# encapsulation mpls
Device(config-template)# sequencing both
Device(config-template)# sequencing resync 1000
```

Related Commands

Command	Description
encapsulation (pseudowire)	Specifies an encapsulation type for tunneling Layer 2 traffic over a pseudowire.
ip cef	Enables Cisco Express Forwarding on the Route Processor card.
pseudowire-class	Specifies the name of an L2TP pseudowire class and enters pseudowire class configuration mode.

set cos

To set the Layer 2 class of service (CoS) value of an outgoing packet, use the **setcos** command in policy-map class configuration mode. To remove a specific CoS value setting, use the **no** form of this command.

set cos {*cos-value*| *from-field* [**table** *table-map-name*]}

no set cos {*cos-value*| *from-field* [**table** *table-map-name*]}

Cisco CMTS and 10000 Series Router

set cos *cos-value*

Syntax Description

cos-value	Specific IEEE 802.1Q CoS value from 0 to 7.
from-field	Specific packet-marking category to be used to set the CoS value of the packet. If you are using a table map for mapping and converting packet-marking values, this establishes the "map from" packet-marking category. Packet-marking category keywords are as follows: • **precedence** • **dscp**
table	(Optional) Indicates that the values set in a specified table map will be used to set the CoS value.
table-map-name	(Optional) Name of the table map used to specify the CoS value. The table map name can be a maximum of 64 alphanumeric characters.

Command Default

No CoS value is set for the outgoing packet.

Command Modes

Policy-map class configuration

Command History

Release	Modification
12.1(5)T	This command was introduced.
12.2(13)T	This command was modified for Enhanced Packet Marking to allow a mapping table (table map)to be used to convert and propagate packet-marking values.

Release	Modification
12.0(16)BX	This command was implemented on the Cisco 10000 series router for the ESR-PRE2.
12.0(31)S	This command was integrated into Cisco IOS Release 12.0(31)S.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB and implemented on the Cisco 10000 series router.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.2(33)SCF	This command was integrated into Cisco IOS Release 12.2(33)SCF.
3.2SE	This command was integrated into Cisco IOS XE Release 3.2SE.

Usage Guidelines

CoS packet marking is supported only in the Cisco Express Forwarding switching path.

The **setcos** command should be used by a router if a user wants to mark a packet that is being sent to a switch. Switches can leverage Layer 2 header information, including a CoS value marking.

The **setcos** command can be used only in service policies that are attached in the output direction of an interface. Packets entering an interface cannot be set with a CoS value.

The **matchcos** and **setcos** commands can be used together to allow routers and switches to interoperate and provide quality of service (QoS) based on the CoS markings.

Layer 2 to Layer 3 mapping can be configured by matching on the CoS value because switches already can match and set CoS values. If a packet that needs to be marked to differentiate user-defined QoS services is leaving a router and entering a switch, the router should set the CoS value of the packet because the switch can process the Layer 2 header.

Using This Command with the Enhanced Packet Marking Feature

You can use this command as part of the Enhanced Packet Marking feature to specify the “from-field” packet-marking category to be used for mapping and setting the CoS value. The “from-field” packet-marking categories are as follows:

- Precedence
- Differentiated services code point (DSCP)

If you specify a “from-field” category but do not specify the **table** keyword and the applicable *table-map-nam* e argument, the default action will be to copy the value associated with the “from-field” category as the CoS value. For instance, if you configure the **setcosprecedence** command, the precedence value will be copied and used as the CoS value.

You can do the same for the DSCP marking category. That is, you can configure the **setcosdscp** command, and the DSCP value will be copied and used as the CoS value.

Note If you configure the **setcosdscp**command, only the *first three bits* (the class selector bits) of the DSCP field are used.

Examples

In the following example, the policy map called "cos-set" is created to assign different CoS values for different types of traffic. This example assumes that the class maps called "voice" and "video-data" have already been created.

```
Router(config)#
policy-map cos-set
Router(config-pmap)#
class voice
Router(config-pmap-c)#
set cos 1
Router(config-pmap-c)#
exit
Router(config-pmap)#
class video-data
Router(config-pmap-c)#
set cos 2
Router(config-pmap-c)#
end
```

Examples

In the following example, the policy map called "policy-cos" is created to use the values defined in a table map called "table-map1". The table map called "table-map1" was created earlier with the **table-map** (value mapping) command. For more information about the **table-map** (value mapping)command, see the **table-map**(value mapping) command page.

In this example, the setting of the CoS value is based on the precedence value defined in "table-map1":

```
Router(config)#
policy-map policy-cos
Router(config-pmap)#
class class-default
Router(config-pmap-c)#
set cos precedence table table-map1
Router(config-pmap-c)#
end
```

Examples

The following example shows how to set the class of service for the 802.1p domain:

```
Router(config)# policy-map cos7
Router(config-pmap)# class cos7
Router(config-pmap-c)# set cos 2
Router(config-pmap-c)# end
```

Note The **setcos** command is applied when you create a service policy in QoS policy-map configuration mode and attach the service policy to an interface or ATM virtual circuit (VC). For information on attaching a service policy, refer to the "Modular Quality of Service Command-Line Interface Overview" chapter of the *Cisco IOS Quality of Service Solutions Configuration Guide* .

Related Commands

Command	Description
match cos	Matches a packet on the basis of Layer 2 CoS marking.
policy-map	Creates or modifies a policy map that can be attached to one or more interfaces to specify a service policy.
service-policy	Attaches a policy map to an input interface or VC, or an output interface or VC, to be used as the service policy for that interface or VC.
set dscp	Marks a packet by setting the Layer 3 DSCP value in the ToS byte.
set precedence	Sets the precedence value in the packet header.
show policy-map	Displays the configuration of all classes for a specified service policy map or all classes for all existing policy maps.
show policy-map class	Displays the configuration for the specified class of the specified policy map.
show policy-map interface	Displays the configuration of all classes configured for all service policies on the specified interface or displays the classes for the service policy for a specific PVC on the interface.

set extcomm-list delete

To allow the deletion of extended community attributes based on an extended community list, use the **set extcomm-list delete** command in route-map configuration mode. To negate a previous **set extcomm-list detect** command, use the **no** form of this command.

set extcomm-list *extended-community-list-number* **delete**

no set extcomm-list *extended-community-list-number* **delete**

Syntax Description

extended-community-list-number	An extended community list number.

Command Default

Extended community attributes based on an extended community list cannot be deleted.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
12.0(26)S	This command was introduced.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.
Cisco IOS XE Release 3.8S	This command was integrated into Cisco IOS XE Release 3.8S.

Usage Guidelines

This command removes extended community attributes of an inbound or outbound Border Gateway Protocol (BGP) update using a route map to filter and determine the extended community attribute to be deleted and replaced. Depending upon whether the route map is applied to the inbound or outbound update for a neighbor, each extended community that passes the route map permit clause and matches the given extended community list will be removed and replaced from the extended community attribute being received from or sent to the BGP neighbor.

For information about how to use this command when translating a route target to a VPN distinguisher and vice versa, see the "BGP—VPN Distinguisher Attribute" module in the *IP Routing: BGP Configuration Guide*.

Examples

The following example shows how to replace a route target 100:3 on an incoming update with a route target of 100:4 using an inbound route map named extmap:

```
.
.
.
Device(config-af)# neighbor 10.10.10.10 route-map extmap in
.
.
.
Device(config)# ip extcommunity-list 1 permit rt 100:3
Device(config)# route-map extmap permit 10
Device(config-route-map)# match extcommunity 1
Device(config-route-map)# set extcomm-list 1 delete
Device(config-route-map)# set extcommunity rt 100:4 additive
```

The following example shows how to configure more than one replacement rule using the route-map configuration **continue** command. Prefixes with RT 100:2 are rewritten to RT 200:3 and prefixes with RT 100:4 are rewritten to RT 200:4. With the **continue** command, route-map evaluation proceeds even if a match is found in a previous sequence.

```
Device(config)# ip extcommunity-list 1 permit rt 100:3
Device(config)# ip extcommunity-list 2 permit rt 100:4
Device(config)# route-map extmap permit 10
Device(config-route-map)# match extcommunity 1
Device(config-route-map)# set extcomm-list 1 delete
Device(config-route-map)# set extcommunity rt 200:3 additive
Device(config-route-map)# continue 20
Device(config)# route-map extmap permit 20
Device(config-route-map)# match extcommunity 2
Device(config-route-map)# set extcomm-list 2 delete
Device(config-route-map)# set extcommunity rt 200:4 additive
Device(config-route-map)# exit
Device(config)# route-map extmap permit 30
```

Related Commands

Command	Description
ip community-list	Creates an extended community access list and controls access to it.
match extcommunity	Matches BGP extended community list attributes.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set extcommunity	Sets BGP extended community attributes.
set extcommunity vpn-distinguisher	Sets a VPN distinguisher attribute to routes.

set ipv6 default next-hop

To specify an IPv6 default next hop to which matching packets are forwarded, use the **set ipv6 default next-hop** command in route-map configuration mode. To delete the default next hop, use the **no** form of this command.

set ipv6 default [**vrf** *vrf-name*| **global**] **next-hop** *global-ipv6-address* [*global-ipv6-address...*]

no set ipv6 default [**vrf** *vrf-name*| **global**] **next-hop** *global-ipv6-address* [*global-ipv6-address...*]

Syntax Description

vrf *vrf-name*	(Optional) Specifies explicitly that the default next-hops are under the specific Virtual Routing and Forwarding (VRF) instance.
global	(Optional) Specifies explicitly that the default next-hops are under the global routing table.
global-ipv6-address	IPv6 global address of the next hop to which packets are output. The next-hop router must be an adjacent router. This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

Command Default

Packets are not forwarded to a default next hop.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
12.3(7)T	This command was introduced.
12.2(30)S	This command was integrated into Cisco IOS Release 12.2(30)S.
12.2(33)SXI4	This command was integrated into Cisco IOS Release 12.2(33)SXI4.
Cisco IOS XE Release 3.2S	This command was modified. It was integrated into Cisco IOS XE Release 3.2S.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Usage Guidelines

An ellipsis (...) in the command syntax indicates that your command input can include multiple values for the *global-ipv6-address* argument.

Use the **set ipv6 default next-hop** command in policy-based routing PBR for IPv6 to specify an IPv6 next-hop address to which a packet is policy routed when the router has no route in the IPv6 routing table or the packets match the default route. The IPv6 next-hop address must be adjacent to the router; that is, reachable by using a directly connected IPv6 route in the IPv6 routing table. The IPv6 next-hop address also must be a global IPv6 address. An IPv6 link-local address cannot be used because the use of an IPv6 link-local address requires interface context.

If the software has no explicit route for the destination in the packet, then the software routes the packet to the next hop as specified by the **set ipv6 default next-hop** command. The optional specified IPv6 addresses are tried in turn.

Use the **ipv6 policy route-map** command, the **route-map** command, and the **match** and **set route-map** commands to define the conditions for PBR packets. The **ipv6 policy route-map** command identifies a route map by name. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the match criteria, which are the conditions under which PBR occurs. The **set** commands specify the set actions, which are the particular routing actions to perform if the criteria enforced by the match commands are met.

The set clauses can be used in conjunction with one another. They are evaluated in the following order:

1. **set ipv6 next-hop**
2. **set interface**
3. **set ipv6 default next-hop**
4. **set default interface**

Note The **set ipv6 next-hop** and **set ipv6 default next-hop** are similar commands. The **set ipv6 next-hop** command is used to policy route packets for which the router has a route in the IPv6 routing table. The **set ipv6 default next-hop** command is used to policy route packets for which the router does not have a route in the IPv6 routing table (or the packets match the default route).

Examples

The following example shows how to set the next hop to which the packet is routed:

```
ipv6 access-list match-dst-1
  permit ipv6 any 2001:DB8:4:1::1/64 any
route-map pbr-v6-default
  match ipv6 address match-dst-1
  set ipv6 default next-hop 2001:DB8:4:4::1/64
```

Related Commands

Command	Description
ipv6 local policy route-map	Identifies a route map to use for local IPv6 PBR.
ipv6 policy route-map	Configures IPv6 policy-based routing (PBR) on an interface.

Command	Description
match ipv6 address	Specifies an IPv6 access list to use to match packets for PBR for IPv6.
match length	Bases policy routing on the Level 3 length of a packet.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or to enable policy routing.
set default interface	Indicates where to output packets that pass a match clause of a route map for policy routing and have no explicit route to the destination.
set interface	Indicates where to output packets that pass a match clause of a route map for policy routing.
set ipv6 next-hop (PBR)	Indicates where to output IPv6 packets that pass a match clause of a route map for policy routing.
set ipv6 precedence	Sets the precedence value in the IPv6 packet header.

set ipv6 next-hop (PBR)

To indicate where to output IPv6 packets that pass a match clause of a route map for policy-based routing (PBR), use the **set ipv6 next-hop** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set ipv6 [**vrf** *vrf-name*| **global**] **next-hop** *global-ipv6-address* [*global-ipv6-address...*]

no set ipv6 [**vrf** *vrf-name*| **global**] **next-hop** *global-ipv6-address* [*global-ipv6-address...*]

Syntax Description

vrf *vrf-name*	(Optional) Specifies explicitly that next-hops are under the specific Virtual Routing and Forwarding (VRF) instance.
global	(Optional) Specifies explicitly that next-hops are under the global routing table.
global-ipv6-address [*global-ipv6-address...*]	IPv6 global address of the next hop to which packets are output. The next-hop router must be an adjacent router. This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

Command Default

Packets are not forwarded to a default next hop.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
12.3(7)T	This command was introduced.
12.2(30)S	This command was integrated into Cisco IOS Release 12.2(30)S.
12.2(33)SXI4	This command was integrated into Cisco IOS Release 12.2(33)SXI4.
Cisco IOS XE Release 3.2S	This command was integrated into Cisco IOS XE Release 3.2S.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Usage Guidelines

The following **set** commands support inherit-VRF, inter-VRF, and VRF-to-global routing in an IPv6-specific implementation:

The **set ipv6 next-hop** command is similar to the **set ip next-hop** command, except that it is IPv6-specific.

An ellipsis (...) in the command syntax indicates that your command input can include multiple values for the *global-ipv6-address* argument. A global IPv6 address must be specified. An IPv6 link-local address cannot be used because the use of an IPv6 link-local address requires interface context.

The *global-ipv6-address* argument must specify an address that is installed in the IPv6 Routing Information Base (RIB) and is directly connected. A directly connected address is an address that is covered by an IPv6 prefix configured on an interface or an address covered by an IPv6 prefix specified on a directly connected static route.

Examples

The following example shows how to set the next hop to which the packet is routed:

```
ipv6 access-list match-dst-1
  permit ipv6 any 2001:DB8::1 any
route-map pbr-v6-default
  match ipv6 address match-dst-1
  set ipv6 next-hop 2001:DB8::F
```

Related Commands

Command	Description
ipv6 local policy route-map	Identifies a route map to use for local IPv6 PBR.
ipv6 policy route-map	Configures IPv6 PBR on an interface.
match ipv6 address	Specifies an IPv6 access list to use to match packets for PBR for IPv6.
match length	Bases policy routing on the Level 3 length of a packet.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set default interface	Indicates where to output packets that pass a match clause of a route map for policy routing and have no explicit route to the destination.
set interface	Indicates where to output packets that pass a match clause of a route map for policy routing.
set ipv6 default next-hop	Specifies an IPv6 default next hop to which matching packets are forwarded.
set ipv6 precedence	Sets the precedence value in the IPv6 packet header.

set mpls experimental

To set the Multiprotocol Label Switching (MPLS) experimental-bit value, use the **set mpls experimental** command in QoS policy-map configuration mode. To return to the default settings, use the **no** form of this command.

set mpls experimental {**imposition**| **topmost**} *experimental-value*

no set mpls experimental {**imposition**| **topmost**}

Syntax Description

imposition	Specifies the experimental-bit value on IP to Multiprotocol Label Switching (MPLS) or MPLS input in all newly imposed labels.
topmost	Specifies the experimental-bit value on the topmost label on the input or output flows.
experimental-value	Experimental-bit value; valid values are from 0 to 7.

Command Default

No experimental-bit value is set.

Command Modes

QoS policy-map configuration

Command History

Release	Modification
12.2(18)SXE	This command was introduced on the Supervisor Engine 720.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Usage Guidelines

This command is not supported on systems that are configured with a Supervisor Engine 2.

Examples

This example shows how to set the experimental-bit value on the topmost label on input or output:

```
Router(config)# policy-map policy1
Router(config-pmap)# class class1
Router(config-pmap-c)# set mpls experimental topmost 5
```

set mpls experimental imposition

To set the value of the Multiprotocol Label Switching (MPLS) experimental (EXP) field on all imposed label entries, use the **set mpls experimental imposition** command in QoS policy-map class configuration mode. To disable the setting, use the **no** form of this command.

set mpls experimental imposition {*mpls-exp-value*| *from-field* [**table** *table-map-name*]}

no set mpls experimental imposition {*mpls-exp-value*| *from-field* [**table** *table-map-name*]}

Cisco 10000 Series Router

set mpls experimental imposition *mpls-exp-value*

no set mpls experimental imposition *mpls-exp-value*

Syntax Description

mpls-exp-value	Specifies the value used to set MPLS EXP bits defined by the policy map. Valid values are numbers from 0 to 7.
from-field	Specific packet-marking category to be used to set the MPLS EXP imposition value. If you are using a table map for mapping and converting packet-marking values, this establishes the "map from" packet-marking category. Packet-marking category keywords are as follows: • **precedence** • **dscp**
table	(Optional) Used in conjunction with the *from-field*argument. Indicates that the values set in a specified table map will be used to set the MPLS EXP imposition value.
table-map-name	(Optional) Used in conjunction with the **table** keyword. Name of the table map used to specify the MPLS EXP imposition value. The name can be a maximum of 64 alphanumeric characters.

Command Default

No MPLS EXP value is set.

Command Modes

QoS policy-map class configuration

Command History

Release	Modification
12.2(13)T	This command was introduced; it replaces (renames) the **set mpls experimental** command, introduced in 12.1(5)T. The **set mpls experimental imposition** command was modified for the Enhanced Packet Marking feature. A mapping table (table map) can now be used to convert and propagate packet-marking values.
12.3(7)XII	This command was implemented on the ESR-PRE2.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The **set mpls experimental imposition** command is supported only on input interfaces. Use this command during label imposition. This command sets the MPLS EXP field on all imposed label entries.

Using This Command with the Enhanced Packet Marking Feature

If you are using this command as part of the Enhanced Packet Marking feature, you can use this command to specify the "from-field" packet-marking category to be used for mapping and setting the class of service (CoS) value. The "from-field" packet-marking categories are as follows:

- Precedence
- Differentiated services code point (DSCP)

If you specify a "from-field" category but do not specify the **table** keyword and the applicable *table-map-name* argument, the default action will be to copy the value associated with the "from-field" category as the MPLS EXP imposition value. For instance, if you configure the **set mpls experimental imposition precedence** command, the precedence value will be copied and used as the MPLS EXP imposition value.

If you configure the **set mpls experimental imposition dscp** command, the DSCP value will be copied and used as the MPLS EXP imposition value.

Note If you configure the **set mpls experimental imposition dscp** command, only the *first three bits* (the class selector bits) of the DSCP field are used.

Cisco 10000 Series Router

Cisco IOS software replaced the **set mpls experimental** command with the **set mpls experimental imposition** command. However, the Cisco 10000 series router continues to use the **set mpls experimental** command for ESR-PRE1. For ESR-PRE2, the command is **set mpls experimental imposition** .

Examples

The following example shows how to set the MPLS EXP value to 3 on all imposed label entries:

```
Router(config-pmap-c)# set mpls experimental imposition 3
```

The following example shows how to create the policy map named policy1 to use the packet-marking values defined in a table map named table-map1. The table map was created earlier with the **table-map** (value mapping) command. For more information about the **table-map** (value mapping) command, see the **table-map** (value mapping) command page. The MPLS EXP imposition value is set according to the DSCP value defined in table-map1.

```
Router(config)# policy-map policy1
Router(config-pmap)# class class-default
Router(config-pmap-c)# set mpls experimental imposition dscp table table-map1
Router(config-pmap-c)# exit
```

Related Commands

Command	Description
set dscp	Marks a packet by setting the Layer 3 DSCP value in the ToS byte.
set mpls experimental topmost	Sets the MPLS EXP field value in the topmost label on either an input or an output interface.
set precedence	Sets the precedence value in the packet header.
show table-map	Displays the configuration of a specified table map or all table maps.
table-map (value-mapping)	Creates and configures a mapping table for mapping and converting one packet-marking value to another.

set mpls experimental topmost

To set the Multiprotocol Label Switching (MPLS) experimental (EXP) field value in the topmost label on either an input or an output interface, use the **set mpls experimental topmost** command in QoS policy-map class configuration mode. To disable the setting, use the **no** form of this command.

set mpls experimental topmost {*mpls-exp-value*| **qos-group** [**table** *table-map-name*]}

no set mpls experimental topmost {*mpls-exp-value*| **qos-group** [**table** *table-map-name*]}

Syntax Description

mpls-exp-value	Specifies the value used to set MPLS experimental bits defined by the policy map. Valid values are numbers from 0 to 7.
qos-group	Specifies that the **qos-group** packet-marking category is used to set the MPLS EXP imposition value. If you are using a table map for mapping and converting packet-marking values, this establishes the "map from" packet-marking category.
table	(Optional) Used in conjunction with the **qos-group** keyword. Indicates that the values set in a specified table map will be used to set the MPLS EXP value.
table-map-name	(Optional) Used with the **table** keyword. Name of the table map used to specify the MPLS EXP value. The name can be a maximum of 64 alphanumeric characters.

Command Default

No MPLS EXP value is set.

Command Modes

QoS policy-map class configuration

Command History

Release	Modification
12.2(13)T	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.2(33)SCF	This command was integrated into Cisco IOS Release 12.2(33)SCF.

Release	Modification
15.4(1)S	This command was implemented on the Cisco ASR 901 series routers.

Usage Guidelines

This command sets the MPLS EXP value only in the topmost label. This command does not affect an IP packet. The MPLS field in the topmost label header is not changed.

Using This Command with the Enhanced Packet Marking Feature

If you are using this command as part of the Enhanced Packet Marking feature, you can use this command to specify the qos-group packet-marking category to be used for mapping and setting the differentiated services code point (DSCP) value.

If you specify the qos-group category but do not specify the **table** *table-map-name* keyword and argument, the default action will be to copy the value associated with the qos-group category as the MPLS EXP topmost value. For instance, if you configure the **set mpls experimental topmost qos-group** command, the QoS group value will be copied and used as the MPLS EXP topmost value.

The valid value range for the MPLS EXP topmost value is a number from 0 to 7. The valid value range for the QoS group is a number from 0 to 99. Therefore, when configuring the **set mpls experimental topmost qos-group** command, note the following points:

- If a QoS group value falls within both value ranges (for example, 6), the packet-marking value will be copied and the packets will be marked.
- If a QoS group value exceeds the MPLS EXP topmost range (for example, 10), the packet-marking value will not copied and the packet will not be marked. No action is taken.

Examples

The following example shows how to set the MPLS EXP value to 3 in the topmost label of an input or output interface:

```
Router(config-pmap)# set mpls experimental topmost 3
```

The following example shows how to create the policy map named policy1 to use the packet-marking values defined in a table map named table-map1. The table map was created earlier with the **table-map** (value mapping) command. For more information about the **table-map** (value mapping) command, see the **table-map** (value mapping) command page.

The following example shows how to set the MPLS EXP value according to the QoS group value defined in table-map1.

```
Router(config)# policy-map policy1
Router(config-pmap)# class class-default
Router(config-pmap-c)# set mpls experimental topmost qos-group table table-map1
Router(config-pmap-c)# exit
```

Related Commands

Command	Description
match mpls experimental topmost	Matches the MPLS EXP field value in the topmost label.

Command	Description
set mpls experimental imposition	Sets the value of the MPLS EXP field on all imposed label entries.
set qos-group	Sets a group ID that can be used later to classify packets.
show table-map	Displays the configuration of a specified table map or all table maps.
table-map (value mapping)	Creates and configures a mapping table for mapping and converting one packet-marking value to another.

set mpls-label

To enable a route to be distributed with a Multiprotocol Label Switching (MPLS) label if the route matches the conditions specified in the route map, use the **set mpls-label** command in route-map configuration mode. To disable this function, use the **no** form of this command.

set mpls-label

no set mpls-label

Syntax Description

This command has no arguments or keywords.

Command Default

No route with an MPLS label is distributed.

Command Modes

Route-map configuration

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	Support for IPv6 was added.
12.2(33)SB	This command was integrated into Cisco IOS Release 12.2(33)SB.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

This command can be used only with the **neighbor route-map out** command to manage outbound route maps for a Border Gateway Protocol (BGP) session.

Use the **route-map** global configuration command with **match** and **set route-map** commands to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the match criteria--the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the set actions--the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

Examples

The following example shows how to create a route map that enables the route to be distributed with a label if the IP address of the route matches an IP address in ACL1:

```
Router(config-router)# route-map incoming permit 10
Router(config-route-map)# match ip address 1
Router(config-route-map)# set mpls-label
```

Related Commands

Command	Description
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list.
match ipv6 address	Distributes IPv6 routes that have a prefix permitted by a prefix list or specifies an IPv6 access list to use to match packets for PBR for IPv6.
match mpls-label	Redistributes routes that contain MPLS labels and match the conditions specified in the route map.
neighbor route-map out	Manage outbound route maps for a BGP session.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.

set ospf router-id

To set a separate Open Shortest Path First (OSPF) router ID for each interface or subinterface on a provider edge (PE) router for each directly attached customer edge (CE) router, use the **set ospf router-id** command in route map configuration mode.

set ospf router-id

Syntax Description

This command has no arguments or keywords.

Command Default

OSPF router ID is not set.

Command Modes

Route map configuration

Command History

Release	Modification
12.0(7)T	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

To use this command, you must enable OSPF and create a routing process.

Examples

The following example shows how to match the PE router IP address 192.168.0.0 against the interface in access list 1 and set to the OSPF router ID:

```
router ospf 2 vrfvpn1-site1
 redistribute bgp 100 metric-type 1 subnets
 network 202.0.0.0 0.0.0.255 area 1
router bgp 100
 neighbor 172.19.89. 62 remote-as 100
 access-list 1 permit 192.168.0.0
 route-map vpn1-site1-map permit 10
 match ip address 1
 set ospf router-id
```

Related Commands

Command	Description
router ospf	Enables OSPF routing, which places the router in router configuration mode.

set vrf

To enable VPN routing and forwarding (VRF) instance selection within a route map for policy-based routing (PBR) VRF selection, use the **set vrf** command in route-map configuration mode. To disable VRF selection within a route map, use the **no** form of this command.

set vrf *vrf-name*

no set vrf *vrf-name*

Syntax Description

vrf-name	Name assigned to the VRF.

Command Default

VRF instance selection is not enabled within a route map for policy-based routing VRF selection.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
12.3(7)T	This command was introduced.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE Release 2.2	This command was integrated into Cisco IOS XE Release 2.2.
12.2(33)SXI	This command was integrated into Cisco IOS Release 12.2(33)SXI.
12.2(33)SXI4	This command was modified. Support for IPv6 was added.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Usage Guidelines

The **set vrf** route-map configuration command was introduced with the Multi-VRF Selection Using Policy-Based Routing feature to provide a PBR mechanism for VRF selection. This command enables VRF selection by policy routing packets through a route map. The route map is attached to the incoming interface. The match criteria are defined in an IP access list or in an IP prefix list. The match criteria can also be defined based on the packet length with the **match length** route map command. The VRF must be defined before you configure this command, and the **ip policy route-map** interface configuration command must be configured to enable policy routing under the interface or subinterface. If the VRF is not defined or if policy routing is not enabled, an error message will be displayed on the console when you attempt to configure the **set vrf** command.

Note The **set vrf** command is not supported in the hardware with the IP Services feature set. If this command is configured in IP Services, the packets are software switched. Hardware forwarding with this command in place requires packet circulation and is only supported in the Advanced IP Services feature set, which supports Multiprotocol Label Switching (MPLS).

In Cisco IOS Release 12.2(33)SXI4 on the Cisco Catalyst 6500, IPv6 PBR allows users to override normal destination IPv6 address-based routing and forwarding results. VRF allows multiple routing instances in Cisco software. The PBR feature is VRF-aware, meaning that it works under multiple routing instances, beyond the default or global routing table.

In PBR, the **set vrf** command decouples the VRF and interface association and allows the selection of a VRF based on the ACL-based classification using the existing PBR or route-map configurations. It provides a single router with multiple routing tables and the ability to select routes based on the ACL classification. The router classifies packets based on ACL, selects a routing table, looks up the destination address, and then routes the packet.

Note The functionality provided by the **set vrf** and **set ip global next-hop** commands can also be configured with the **set default interface**, **set interface**, **set ip default next-hop**, and **set ip next-hop** commands. However, the **set vrf** and **set ip global next-hop** commands take precedence over the **set default interface**, **set interface**,**set ip default next-hop**, and **set ip next-hop** commands. No error message is displayed indicating that VRF is already enabled if you attempt to configure the **set vrf** command with any of these four **set** commands.

Examples

The following example shows a route-map sequence that selects and sets a VRF based on the match criteria defined in three different access lists. (The access list configuration is not shown in this example.) If the route map falls through and a match does not occur, the packet will be dropped if the destination is local.

```
route-map PBR-VRF-Selection permit 10
match ip address 40
set vrf VRF1
!
route-map PBR-VRF-Selection permit 20
match ip address 50
set vrf VRF2
!
route-map PBR-VRF-Selection permit 30
match ip address 60
set vrf VRF3
```

Related Commands

Command	Description
access-list (IP standard)	Defines a standard IP access list.
debug ip policy	Displays the IP policy routing packet activity.
ip policy route-map	Identifies a route map to use for policy routing on an interface.

Command	Description
ip vrf	Configures a VRF routing table.
ip vrf receive	Inserts the IP address of an interface as a connected route entry in a VRF routing table.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, or performs policy routing on packets.
match length	Bases policy routing on the Level 3 length of a packet.
route-map	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set default interface	Indicates where to output packets that pass a match clause of a route map for policy routing and that have no explicit route to the destination.
set interface	Indicates where to forward packets that pass a match clause of a route map for policy routing.
set ip default next-hop	Indicates where to output packets that pass a match clause of a route map for policy routing and for which the Cisco software has no explicit route to a destination.
set ip next-hop	Indicates where to output packets that pass a match clause of a route map for policy routing.

show acircuit checkpoint

To display checkpointing information for each attachment circuit (AC), use the **show acircuit checkpoint** command in privileged EXEC mode.

show acircuit checkpoint

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.2(25)S	This command was introduced.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.2(33)SRC	This command was integrated into Cisco IOS Release 12.2(33)SRC.

Usage Guidelines

This command is used for interface-based attachment circuits. For Frame Relay and ATM circuits, use the following commands to show redundancy information:

- **debug atm ha-error**
- **debug atm ha-events**
- **debug atm ha-state**
- **debug atm l2transport**
- **debug frame-relay redundancy**

Examples

The following show acircuit checkpoint command displays information about the ACs that have been check-pointed. The output varies, depending on whether the command output is for the active or standby Route Processor (RP).

On the active RP, the command displays the following output:

```
Router# show acircuit checkpoint
AC HA Checkpoint info:
Last Bulk Sync: 1 ACs
 AC    IW    XC    Id  VCId   Switch    Segment  St  Chkpt
----  ----  ----  ---  ----  --------  --------  --  -----
```

```
HDLC  LIKE  ATOM    3   100      1000      1000   0    N
VLAN  LIKE  ATOM    2  1002      2001      2001   3    Y
```

On the standby RP, the command displays the following output::

```
Router# show acircuit checkpoint
AC HA Checkpoint info:
 AC    IW    XC    Id  VCId   Switch    Segment  St  F-SLP
----  ----  ----  ---  ----  --------  --------  --  -----
HDLC  LIKE  ATOM    3   100         0         0   0   001
VLAN  LIKE  ATOM    2  1002      2001      2001   2   000
```

The table below describes the significant fields shown in the display.

Table 1: show acircuit checkpoint Field Descriptions

Field	Description
Last Bulk Sync	The number of ACs that were sent to the backup RP during the last bulk synchronization between the active and backup RPs.
AC	The type of attachment circuit.
IW	The type of interworking, either like-to-like (AToM) or any-to-any (Interworking).
XC	The type of cross-connect. Only AToM ACs are checkpointed.
ID	This field varies, depending on the type of attachment circuit. For Ethernet VLANs, the ID is the VLAN ID. For PPP and High-Level Data Link Control (HDLC), the ID is the AC circuit ID.
VCID	The configured virtual circuit ID.
Switch	An ID used to correlate the control plane and data plane contexts for this virtual circuit (VC). This is an internal value that is not for customer use.
Segment	An ID used to correlate the control plane and data plane contexts for this VC. This is an internal value that is not for customer use.
St	The state of the attachment circuit. This is an internal value that is not for customer use.
Chkpt	Whether the information about the AC was checkpointed.
F-SLP	Flags that provide more information about the state of the AC circuit. These values are not for customer use.

Related Commands

Command	Description
show mpls l2transport vc	Displays AToM status information.
show mpls l2transport vc checkpoint	Displays the status of the checkpointing process for both the active and standby RPs.

show atm cell-packing

To display the average number of cells in packets sent from an ATM permanent virtual circuit (PVC) to a single Multiprotocol Label Switching (MPLS) pseudowire and the average number of cells in packets that are received from an MPLS pseudowire and sent to the respective ATM virtual circuits (VCs), use the **show atm cell-packing** command in privileged EXEC mode.

show atm cell-packing

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
Cisco IOS XE Release 3.7S	This command was introduced.

Usage Guidelines

To map one or more ATM PVCs to a single pseudowire, an N:1 PVC must be created on an ATM interface. The output of the **show atm cell-packing** command can be used to gauge the amount of cell packing in packets that originate from a device and are received by the device, for a specific pseudowire. Cisco IOS software calculates the average number of cells per packet in each direction.

Examples

The following is sample output from the **show atm cell-packing** command. The fields in the output are self-explanatory.

```
Device# show atm cell-packing

                                    average                     average
circuit                 local  nbr of cells     peer     nbr of cells       MCPT
type                    MNCP   rcvd in one pkt  MNCP     sent in one pkt    us)

ATM4/0/0.1  vc   1/41    20           1     20         1                  100
ATM4/0/0.1  vc   1/42    20           1     20         1                  100
```

Related Commands

Command	Description
cell-packing	Enables multiple cell packing.

show atm vc

To display all ATM permanent virtual circuits (PVCs), switched virtual circuits (SVCs), and traffic information, use the **show atm vc** command in privileged EXEC mode.

show atm vc [*vcd-number*| **range** *lower-limit-vcd upper-limit-vcd*] [**interface atm** *interface-number*] [**detail** [**prefix** {**vpi/vci**| **vcd**| **interface**| **vc_name**}]] [*connection-name*] [**signalling** [**freed-svcs**| [**cast-type** {**p2mp**| **p2p**}]]] **[detail]** [**interface atm** *interface-number*| **summary atm** *interface-number*]

Syntax Description

vcd-number	(Optional) Specifies a unique virtual circuit descriptor (VCD) number that identifies PVCs within one ATM interface.
range *lower-limit-vcd upper-limit-vcd*	(Optional) Specifies the range of VCs. Displays all the VC information for the specified range of VCDs. The *lower-limit-vcd* argument specifies the lower limit of the VCD range. The *upper-limit-vcd* argument specifies the upper limit of the VCD range.
interface atm *interface-number*	(Optional) Interface number or subinterface number of the PVC or SVC. Displays all PVCs and SVCs on the specified interface or subinterface. The *interface-number* uses one of the following formats, depending on the router platform you use: • For the ATM Interface Processor (AIP) on Cisco 7500 series routers; for the ATM port adapter, ATM-CES port adapter, and enhanced ATM port adapter on Cisco 7200 series routers; for the 1-port ATM-25 network module on Cisco 2600 and 3600 series routers: *slot* **/ 0 .** *subinterface-number* **multipoint** • For the ATM port adapter and enhanced ATM port adapter on Cisco 7500 series routers : *slot* **/** *port-adapter* **/ 0 .** *subinterface-number* **multipoint** • For the network processing module (NPM) on Cisco 4500 and Cisco 4700 routers : *number* **.** *subinterface-number* **multipoint** • For a description of these arguments, refer to the **interface atm** command.
detail	(Optional) Displays the detailed information about the VCs.

prefix	(Optional) Displays detailed information about the selected VC category. You must specify one of the following VC categories: • **vpi/vci** --Virtual path identifier and virtual channel identifier. • **vcd** --Virtual circuit descriptor. • **interface** --Interface in which the VCD is configured. • **vc_name** --Name of the PVC or SVC.
connection-name	(Optional) Connection name of the PVC or SVC.
signalling	(Optional) Displays the ATM interface signaling information for all the interfaces.
freed-svcs	(Optional) Displays the details of the last few freed SVCs.
cast-type	(Optional) SVC cast type. You must specify one of the following connections: • **p2mp** --Point to multipoint connection. • **p2p** --Point to point connection.
summary atm *interface-number*	(Optional) Displays a summary of VCs.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
10.0	This command was introduced.
11.1CA	This command was modified. Information about VCs on an ATM-CES port adapter was added to the command output.
12.0(5)T	This command was modified. Information about VCs on an extended Multiprotocol Label Switching (MPLS) ATM interface was added to the command output.
12.2(25)S	This command was modified. Information about packet drops and errors was added to the command output.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series routers.

Release	Modification
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB and the **signalling** keyword was added.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
Cisco IOS XE 2.3	This command was implemented on the Cisco ASR 1000 series routers.

Usage Guidelines

If no value is specified for the *vcd* argument, the command displays information for all PVCs and SVCs. The output is in summary form (one line per virtual circuit).

VCs on the extended MPLS ATM interfaces do not appear in the **show atm vc** command output. Instead, the **show xtagatm vc** command provides a similar output that shows information only on extended MPLS ATM VCs.

Note The SVCs and the **signalling** keyword are not supported on the Cisco ASR 1000 series routers.

Examples

The following is sample output from the **show atm vc** command when no value for the *vcd* argument is specified. The status field is either ACTIVE or IN (inactive).

```
Router# show atm vc
Interface     VCD   VPI    VCI Type  AAL/Encaps     Peak   Avg.  Burst Status
ATM2/0          1     0      5  PVC  AAL5-SAAL     155000 155000    93 ACTIVE
ATM2/0.4        3     0     32  SVC  AAL5-SNAP     155000 155000    93 ACTIVE
ATM2/0.65432   10    10     10  PVC  AAL5-SNAP     100000  40000    10 ACTIVE
ATM2/0         99     0     16  PVC  AAL5-ILMI     155000 155000    93 ACTIVE
ATM2/0.105    250    33     44  PVC  AAL5-SNAP     155000 155000    93 ACTIVE
ATM2/0.100    300    22     33  PVC  AAL5-SNAP     155000 155000    93 ACTIVE
ATM2/0.12345 2047   255 65535  PVC  AAL5-SNAP         56     28  2047 ACTIVE
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified for a circuit emulation service (CES) circuit:

```
Router# show atm vc 2
ATM6/0: VCD: 2, VPI: 10, VCI: 10
PeakRate: 2310, Average Rate: 2310, Burst Cells: 94
CES-AAL1, etype:0x0, Flags: 0x20138, VCmode: 0x0
OAM DISABLED
InARP DISABLED
OAM cells received: 0
OAM cells sent: 334272
Status: ACTIVE
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified, displaying statistics for that virtual circuit only:

```
Router# show atm vc 8
ATM4/0: VCD: 8, VPI: 8, VCI: 8
PeakRate: 155000, Average Rate: 155000, Burst Cells: 0
AAL5-LLC/SNAP, etype:0x0, Flags: 0x30, VCmode: 0xE000
OAM frequency: 0 second(s)
InARP frequency: 1 minute(s)
```

```
InPkts: 181061, OutPkts: 570499, InBytes: 757314267, OutBytes: 2137187609
InPRoc: 181011, OutPRoc: 10, Broadcasts: 570459
InFast: 39, OutFast: 36, InAS: 11, OutAS: 6
OAM cells received: 0
OAM cells sent: 0
Status: UP
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified, AAL3/4 is enabled, an ATM Switched Multimegabit Data Service (SMDS) subinterface has been defined, and a range of message identifier numbers (MIDs) has been assigned to the PVC:

```
Router# show atm vc 1
ATM4/0.1: VCD: 1, VPI: 0, VCI: 1
PeakRate: 0, Average Rate: 0, Burst Cells: 0
AAL3/4-SMDS, etype:0x1, Flags: 0x35, VCmode: 0xE200
MID start: 1, MID end: 16
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified and generation of Operation, Administration, and Maintenance (OAM) F5 loopback cells has been enabled:

```
Router# show atm vc 7
ATM4/0: VCD: 7, VPI: 7, VCI: 7
PeakRate: 0, Average Rate: 0, Burst Cells: 0
AAL5-LLC/SNAP, etype:0x0, Flags: 0x30, VCmode: 0xE000
OAM frequency: 10 second(s)
InARP DISABLED
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0, Broadcasts: 0
InFast:0, OutFast:0, InAS:0, OutAS:0
OAM cells received: 0
OAM cells sent: 1
Status: UP
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified, and there is an incoming multipoint virtual circuit:

```
Router# show atm vc 3
ATM2/0: VCD: 3, VPI: 0, VCI: 33
PeakRate: 0, Average Rate: 0, Burst Cells: 0
AAL5-MUX, etype:0x809B, Flags: 0x53, VCmode: 0xE000
OAM DISABLED
InARP DISABLED
InPkts: 6646, OutPkts: 0, InBytes: 153078, OutBytes: 0
InPRoc: 6646, OutPRoc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
interface =  ATM2/0, call remotely initiated, call reference = 18082
vcnum = 3, vpi = 0, vci = 33, state = Active
 aal5mux vc, multipoint call
Retry count: Current = 0, Max = 10
timer currently inactive, timer value = never
Root Atm Nsap address: DE.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified, and there is an outgoing multipoint virtual circuit:

```
Router# show atm vc 6
ATM2/0: VCD: 6, VPI: 0, VCI: 35
PeakRate: 0, Average Rate: 0, Burst Cells: 0
AAL5-MUX, etype:0x800, Flags: 0x53, VCmode: 0xE000
OAM DISABLED
InARP DISABLED
InPkts: 0, OutPkts: 818, InBytes: 0, OutBytes: 37628
InPRoc: 0, OutPRoc: 0, Broadcasts: 818
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
interface =  ATM2/0, call locally initiated, call reference = 3
vcnum = 6, vpi = 0, vci = 35, state = Active
 aal5mux vc, multipoint call
Retry count: Current = 0, Max = 10
```

```
timer currently inactive, timer value = never
Leaf Atm Nsap address: DE.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12
Leaf Atm Nsap address: CD.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.12
```

The following is sample output from the **show atm vc** command when a *vcd* value is specified and there is a PPP-over-ATM connection:

```
Router# show atm vc 1
ATM8/0.1: VCD: 1, VPI: 41, VCI: 41
PeakRate: 155000, Average Rate: 155000, Burst Cells: 96
AAL5-CISCOPPP, etype:0x9, Flags: 0xC38, VCmode: 0xE000
virtual-access: 1, virtual-template: 1
OAM DISABLED
InARP DISABLED
InPkts: 13, OutPkts: 10, InBytes: 198, OutBytes: 156
InPRoc: 13, OutPRoc: 10, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
OAM cells received: 0
OAM cells sent: 0
```

The following is sample output from the **show atm vc** command for IP multicast virtual circuits. The display shows the leaf count for multipoint VCs opened by the root. VCD 3 is a root of a multipoint VC with three leaf routers. VCD 4 is a leaf of some other router's multipoint VC. VCD 12 is a root of a multipoint VC with only one leaf router.

```
Router# show atm vc
            VCD/                                      Peak     Avg/Min     Burst
Interface   Name  VPI   VCI   Type     Encaps          Kbps      Kbps      Cells         Sts
0/0         1       0     5    PVC      SAAL       155000    155000       96           UP
0/0         2       0    16    PVC      ILMI       155000    155000       96           UP
0/0         3       0   124 MSVC-3      SNAP       155000    155000       96           UP
0/0         4       0   125   MSVC      SNAP       155000    155000       96           UP
0/0         5       0   126   MSVC      SNAP       155000    155000       96           UP
0/0         6       0   127   MSVC      SNAP       155000    155000       96           UP
0/0         9       0   130   MSVC      SNAP       155000    155000       96           UP
0/0         10      0   131    SVC      SNAP       155000    155000       96           UP
0/0         11      0   132 MSVC-3      SNAP       155000    155000       96           UP
0/0         12      0   133 MSVC-1      SNAP       155000    155000       96           UP
0/0         13      0   134    SVC      SNAP       155000    155000       96           UP
0/0         14      0   135 MSVC-2      SNAP       155000    155000       96           UP
0/0         15      0   136 MSVC-2      SNAP       155000    155000       96           UP
```

The following is sample output from the **show atm vc** command for an IP multicast virtual circuit. The display shows the owner of the VC and leaves of the multipoint VC. This VC was opened by IP multicast. The three leaf routers' ATM addresses are included in the display. The VC is associated with IP group address 10.1.1.1.

```
Router# show atm vc 11
ATM0/0: VCD: 11, VPI: 0, VCI: 132
PeakRate: 155000, Average Rate: 155000, Burst Cells: 96
AAL5-LLC/SNAP, etype:0x0, Flags: 0x650, VCmode: 0xE000
OAM DISABLED
InARP DISABLED
InPkts: 0, OutPkts: 12, InBytes: 0, OutBytes: 496
InPRoc: 0, OutPRoc: 0, Broadcasts: 12
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
OAM cells received: 0
OAM cells sent: 0
Status: ACTIVE, TTL: 2, VC owner: IP Multicast (10.1.1.1)
interface =  ATM0/0, call locally initiated, call reference = 2
vcnum = 11, vpi = 0, vci = 132, state = Active
 aal5snap vc, multipoint call
Retry count: Current = 0, Max = 10
timer currently inactive, timer value = 00:00:00
Leaf Atm Nsap address: 47.0091810000000002BA08E101.444444444444.02
Leaf Atm Nsap address: 47.0091810000000002BA08E101.333333333333.02
Leaf Atm Nsap address: 47.0091810000000002BA08E101.222222222222.02
```

The following is sample output from the **show atm vc** command where no VCD is specified and private VCs are present:

```
Router# show atm vc
AAL /         Peak   Avg.  Burst
Interface     VCD   VPI   VCI Type  Encapsulation  Kbps   Kbps  Cells Status
ATM1/0          1     0    40  PVC  AAL5-SNAP          0      0     0 ACTIVE
ATM1/0          2     0    41  PVC  AAL5-SNAP          0      0     0 ACTIVE
ATM1/0          3     0    42  PVC  AAL5-SNAP          0      0     0 ACTIVE
ATM1/0          4     0    43  PVC  AAL5-SNAP          0      0     0 ACTIVE
ATM1/0          5     0    44  PVC  AAL5-SNAP          0      0     0 ACTIVE
ATM1/0         15     1    32  PVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0         17     1    34  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0         26     1    43  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0         28     1    45  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0         29     1    46  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0         33     1    50  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
```

When you specify a VCD value and the VCD corresponds to that of a private VC on a control interface, the display output appears as follows:

```
Router# show atm vc 15
ATM1/0 33     1    50  TVC  AAL5-XTAGATM       0      0     0 ACTIVE
ATM1/0: VCD: 15, VPI: 1, VCI: 32, etype:0x8, AAL5 - XTAGATM, Flags: 0xD38
PeakRate: 0, Average Rate: 0, Burst Cells: 0, VCmode: 0x0
XTagATM1, VCD: 1, VPI: 0, VCI: 32
OAM DISABLED, InARP DISABLED
InPkts: 38811, OutPkts: 38813, InBytes: 2911240, OutBytes: 2968834
InPRoc: 0, OutPRoc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
OAM F5 cells sent: 0, OAM cells received: 0
Status: ACTIVE
```

The table below describes the fields shown in the displays.

Table 2: show atm vc Field Descriptions

Field	Description
Interface	Interface slot and port.
VCD/Name	Virtual circuit descriptor (virtual circuit number). The connection name is displayed if the virtual circuit (VC) was configured using the **pvc** command and the name was specified.
VPI	Virtual path identifier.
VCI	Virtual channel identifier.

Field	Description
Type	Type of VC, either PVC, SVC, TVC, or multipoint SVC (MSVC). • MSVC (with no -*x*) indicates that VCD is a leaf of some other router's multipoint VC. • MSVC-*x* indicates there are *x* leaf routers for that multipoint VC opened by the root. Type of PVC detected from PVC discovery, either PVC-D, PVC-L, or PVC-M. • PVC-D indicates a PVC created due to PVC discovery. • PVC-L indicates that the corresponding peer of this PVC could not be found on the switch. • PVC-M indicates that some or all of the quality of service (QoS) parameters of this PVC do not match those of the corresponding peer on the switch. • TVC indicates a Tag VC.
Encaps	Type of ATM adaptation layer (AAL) and encapsulation.
PeakRate	Kilobits per second sent at the peak rate.
Average Rate	Kilobits per second sent at the average rate.
Burst Cells	Value that equals the maximum number of ATM cells the VC can send at peak rate.
Status	Status of the VC connection. • UP indicates that the connection is enabled for data traffic. • DN indicates that the connection is down (not ready for data traffic). When the Status field is DN (down), a State field is shown. • IN indicates that the interface is down (inactive). • ACTIVE indicates that the interface is in use and active.
etype	Encapsulation type.

Field	Description
Flags	Bit mask describing VC information. The flag values are summed to result in the displayed value. 0x10000 ABR VC 0x20000 CES VC 0x40000 TVC 0x100 TEMP (automatically created) 0x200 MULTIPOINT 0x400 DEFAULT_RATE 0x800 DEFAULT_BURST 0x10 ACTIVE 0x20 PVC 0x40 SVC 0x0 AAL5-SNAP 0x1 AAL5-NLPID 0x2 AAL5-FRNLPID 0x3 AAL5-MUX 0x4 AAL3/4-SMDS 0x5 QSAAL 0x6 AAL5-ILMI 0x7 AAL5-LANE 0x8 AAL5-XTAGATM 0x9 CES-AAL1 0xA F4-OAM
VCmode	AIP-specific or NPM-specific register describing the usage of the VC. This register contains values such as rate queue, peak rate, and AAL mode, which are also displayed in other fields.
OAM frequency	Seconds between OAM loopback messages, or DISABLED if OAM is not in use on this VC.
InARP frequency	Minutes between Inverse Address Resolution Protocol (InARP) messages, or DISABLED if InARP is not in use on this VC.
virtual-access	Virtual access interface identifier.
virtual-template	Virtual template identifier.
InPkts	Total number of packets received on this VC. This number includes all fast-switched and process-switched packets.
OutPkts	Total number of packets sent on this VC. This number includes all fast-switched and process-switched packets.
InBytes	Total number of bytes received on this VC. This number includes all fast-switched and process-switched packets.
OutBytes	Total number of bytes sent on this VC. This number includes all fast-switched and process-switched packets.
InPRoc	Number of process-switched input packets.
OutPRoc	Number of process-switched output packets.

Field	Description
Broadcasts	Number of process-switched broadcast packets.
InFast	Number of fast-switched input packets.
OutFast	Number of fast-switched output packets.
InAS	Number of autonomous-switched or silicon-switched input packets.
VC TxRingLimit	Transmit Ring Limit for this VC.
VC Rx Limit	Receive Ring Limit for this VC.
Transmit priority	ATM service class transmit priority for this VC.
InCells	Number of incoming cells on this VC.
OutCells	Number of outgoing cells on this VC.
InPktDrops	A non-zero value for the InPktDrops of a VC counter suggests that the ATM interface is running out of packet buffers for an individual VC, or is exceeding the total number of VC buffers that can be shared by the VCs.
OutPktDrops	The PA-A3 driver increments the OutPktDrops counter when a VC fills its individual transmit buffer quota. The purpose of the quota is to prevent a consistently oversubscribed VC from grabbing all of the packet buffer resources and hindering other VCs from transmitting normal traffic within their traffic contracts.
InCellDrops	Number of incoming cells dropped on this VC.
OutCellDrops	Number of outgoing cells dropped on this VC.
InByteDrops	Number of incoming bytes that are dropped on this VC.
OutByteDrops	Number of outgoing bytes that are dropped on this VC.
CrcErrors	Number of cyclic redundancy check (CRC) errors on this VC.
SarTimeOuts	Number of segmentation and reassembly sublayer time-outs on this VC.

Field	Description
OverSizedSDUs	Number of over-sized service data units on this VC
LengthViolation	Number of length violations on this VC. A length violation occurs when a reassembled packet is dropped without checking the CRC.
CPIErrors	The Common Part Indicator error field is a one octet field in the AAL5 encapsulation of an ATM cell and must be set to 0. If it is received with some other value, it is flagged as an error by the interface. For example, this error may indicate data corruption.
Out CLP	Number of packets or cells where the Output Cell Loss Priority bit is set.
OutAS	Number of autonomous-switched or silicon-switched output packets.
OAM cells received	Number of OAM cells received on this VC.
OAM cells sent	Number of OAM cells sent on this VC.
TTL	Time to live in ATM hops across the VC.
VC owner	IP Multicast address of the group.

Related Commands

Command	Description
atm nsap-address	Sets the NSAP address for an ATM interface using SVC mode.
show xtagatm vc	Displays information about the VCs on the extended MPLS ATM interfaces.

show bridge-domain

To display bridge-domain information, use the **show bridge-domain** command in privileged EXEC mode.

show bridge-domain [[*bridge-id*] **[c-mac]** [**mac** {**security** [**address**| **last violation**| **statistics**]| **static address**| **table** [*mac-address*| **aging-time**| **count**]}]| **split-horizon** [**group** {*group-number*| **all**| **none**}]| **stats**]

Syntax Description

bridge-id	(Optional) Identifier for the bridge-domain instance. Integer in the range 1 to Platform_Upper_Bound, where Platform_Upper_Bound is a platform-specific upper limit.
c-mac	(Optional) Displays a specified customer bridge domain.
mac	(Optional) Displays MAC address data. **Note** The **mac** keyword is not supported on the Cisco ASR 1000 Series Aggregation Services Router.
security	(Optional) Displays MAC security information.
address	(Optional) Displays addresses. • When used with the **security** keyword, displays secure addresses on a specified service instance. • When used with the **static** keyword, displays static addresses in a specified bridge domain. **Note** The **address** keyword is not supported on the Cisco ASR 1000 Series Aggregation Services Router.
last	(Optional) Displays the last violation recorded on the specified bridge domain.
violation	(Optional) Displays information about the last violation recorded on the specified bridge domain.
statistics	(Optional) Displays the number of secured MAC addresses and related statistics.
static	(Optional) Displays static MAC information.
table	(Optional) Displays commands related to the MAC address table.
mac-address	(Optional) Displays the MAC address.

aging-time	(Optional) Displays the time, in minutes, that an entry remains before aging out of the MAC address table.
count	(Optional) Displays the total number of addresses in a bridge-domain table.
split-horizon	(Optional) Displays bridge-domain information for a split-horizon.
group	(Optional) Displays bridge-domain information for a split-horizon group.
group-number	(Optional) Number of a specific split-horizon group for bridge-domain information display.
all	(Optional) Selects all ports in split-horizon groups for bridge-domain information display.
none	(Optional) Selects ports that do not belong to any split-horizon group for bridge-domain information display.
stats	(Optional) Displays bridge-domain statistics.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SRD	This command was introduced.
12.2(33)SRE	This command was modified. The **address**, **aging-time**, **count**, **static**, and **table** keywords and the *mac-address* argument were added.
Cisco IOS XE Release 3.5S	This command was integrated into Cisco IOS XE Release 3.5S to provide support for the Cisco ASR 903 Series Aggregation Services Router. This command was modified to provide support for Ethernet Flow Points (EFPs) on trunk ports (interfaces).
15.1(2)SNG	This command was implemented on the Cisco ASR 901 Series Aggregation Services Router.
15.3(1)S	This command was integrated into Cisco IOS Release 15.3(1)S. The command was modified to display the MAC address limit for the bridge domain.

Usage Guidelines

This command is useful for system monitoring and troubleshooting.

This command is available on both linecards and route processors. To invoke this command on a linecard, log in to the linecard. To invoke this command on a route processor, use the **remote command module** command; for example, **remote command module16 bridge-domain 25**.

Note The **remote command** command is not supported on the Cisco ASR 1000 Series Aggregation Services Router.

Examples

The following is sample output of the **show bridge-domain** command. The output varies slightly by platform. The fields are self-explanatory.

```
Device# show bridge-domain 10

Bridge-domain 10 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
    GigabitEthernet0/2/2 service instance 10
    GigabitEthernet0/2/3 service instance 10
   MAC address    Policy Tag           Age Pseudoport[VC-lbl,egr-intf]
   0000.5200.010E fwd    dynamic    300 GigabitEthernet0/2/3.EFP10
   0000.5200.010C fwd    dynamic    300 GigabitEthernet0/2/3.EFP10
   0000.5200.0107 fwd    dynamic    299 GigabitEthernet0/2/3.EFP10
   0000.5200.0104 fwd    dynamic    300 GigabitEthernet0/2/3.EFP10
```

The following is sample output where the MAC address limit is displayed:

```
Device# show bridge-domain 100 mac address

Bridge-domain 100 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 5 minute(s)
Maximum address limit: 10240      Current addresses: 300
    Ethernet0/0 service instance 100
    Maximum address limit: 200      Current addresses: 100
1 ports belonging to split-horizon group 1
    Ethernet0/0 service instance 101 (split-horizon group 1)
   Maximum address limit: 300      Current addresses: 150
  Software Bridging Info for Bridge Domain 100, contains 2 ports
   MAC address     Pseudoport
```

The table below describes the significant fields shown in the display.

Table 3: show bridge-domain Field Descriptions

Field	Description
Maximum address limit	The maximum MAC addresses configured for the bridge domain.
Current addresses	The current number of MAC addresses learned for the bridge domain. **Note** This information may not display for all platforms.

The following example shows the sample output where information of the Ethernet over Generic Routing Encapsulation (GRE) for a specific bridge domain are displayed:

```
Device# show bridge-domain 10

Bridge-domain 10 (2 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 180 second(s)
    GigabitEthernet2/0/0 service instance 1
    Virtual-Ethernet1 service instance 1
 MAC address    Policy    Tag     Age Pseudoport
   0000.0000.0002 forward dynamic   177 Virtual-Ethernet1.EFP1 sGRE src:11.1.1.1 dst:1.1.1.2

   0000.0000.0001 forward dynamic 180  GigabitEthernet2/0/0.EFP1
```

Related Commands

Command	Description
clear bridge-domain	Clears bridge-domain attributes that are not needed.
remote command	Executes a Cisco 7600 Series Router command directly on the console or a specified module without having to log into the Cisco 7600 Series Router first.
show ethernet service instance	Displays information about Ethernet service instances.
show ethernet service interface	Displays interface-only information about Ethernet customer service instances.

show connection

To display the status of interworking connections, use the **show connection** command in privileged EXEC mode.

show connection[**all**| *element*| **id** *startid*-[*endid*]| **name** *name*| **port** *port*]

Syntax Description

all	(Optional) Displays information about all interworking connections.
element	(Optional) Displays information about the specified connection element.
id	(Optional) Displays information about the specified connection identifier.
startid	Starting connection ID number.
endid	(Optional) Ending connection ID number.
name *name*	(Optional) Displays information about the specified connection name.
port *port*	(Optional) Displays information about all connections on an interface. (In Cisco IOS Release 12.0S, only ATM, serial, and Fast Ethernet are shown.)

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.1(2)T	This command was introduced as show connect (FR-ATM).
12.0(27)S	This command was integrated into Cisco IOS Release 12.0(27)S and updated to show all ATM, serial, and Fast Ethernet interworking connections.
12.4(2)T	The command output was modified to add Segment 1 and Segment 2 fields for Segment state and channel ID.
12.0(30)S	This command was integrated into Cisco IOS Release 12.0(30)S.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Release	Modification
12.4(8)	This command was integrated into Cisco IOS Release 12.4(8).
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.4(11)T	This command was integrated into Cisco IOS Release 12.4(11)T.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.2(33)SB	This command was updated to display High-Level Data Link Control (HDLC) local switching connections.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.
15.1(2)SNH	This command was implemented on the Cisco ASR 901 Series Aggregation Services Routers.

Examples

The following example shows the local interworking connections on a router:

```
Device# show connection

ID   Name               Segment 1            Segment 2           State
=============================================================================
1    conn1          ATM 1/0/0 AAL5 0/100   ATM 2/0/0 AAL5 0/100   UP
2    conn2          ATM 2/0/0 AAL5 0/300   Serial0/1 16           UP
3    conn3          ATM 2/0/0 AAL5 0/400   FA 0/0.1 10            UP
4    conn4          ATM 1/0/0 CELL 0/500   ATM 2/0/0 CELL 0/500   UP
5    conn5          ATM 1/0/0 CELL 100     ATM 2/0/0 CELL 100     UP
```

The table below describes the significant fields shown in the display.

Table 4: show connection Field Descriptions

Field	Description
ID	Arbitrary connection identifier assigned by the operating system.
Name	Name of the connection.

Field	Description
Segment 1 Segment 2	Information about the interworking segments: • Interface name and number. • Segment state, interface name and number, and channel ID. Segment state will displays nothing if the segment state is UP, "-" if the segment state is DOWN, and "***Card Removed***" if the segment state is DETACHED. • Type of encapsulation (if any) assigned to the interface. • Permanent virtual circuit (PVC) assigned to the ATM interface, data-link connection identifier (DLCI) assigned to the serial interface, or VLAN ID assigned to the Ethernet interface.
State	Status of the connection, which is one of the following: INVALID, UP, ADMIN UP, ADMIN DOWN, OPER DOWN, COMING UP, NOT VERIFIED, ERR.

Related Commands

Command	Description
connect (L2VPN local switching)	Connects two different or like interfaces on a router.
show atm pvc	Displays the status of ATM PVCs and SVCs.
show frame-relay pvc	Displays the status of Frame Relay interfaces.

show controllers vsi control-interface

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers vsi control-interface** is not available in Cisco IOS software.

To display information about an ATM interface configured with the **tag-control-protocol vsi** command to control an external switch (or if an interface is not specified, to display information about all Virtual Switch Interface [VSI] control interfaces), use the **show controllers vsi control-interface** command in user EXEC or privileged EXEC mode.

show controllers vsi control-interface [*interface*]

Syntax Description

interface	(Optional) Specifies the interface number.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.4(20)T	This command was removed.

Examples

The following is sample output from the **show controllers vsi control-interface** command:

```
Router# show controllers vsi control-interface
Interface:            ATM2/0        Connections:          14
```

The display shows the number of cross-connects currently on the switch that were established by the MPLS LSC through the VSI over the control interface.

The table below describes the significant fields shown in the display.

Table 5: show controllers vsi control-interface Field Descriptions

Field	Description
Interface	The (Cisco IOS) interface name.
Connections	The number of cross connections currently on the switch.

Related Commands

Command	Description
tag-control-protocol vsi	Configures the use of VSI on a control port.

show controllers vsi descriptor

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers vsi descriptor** command is not available in Cisco IOS software.

To display information about a switch interface discovered by the Multiprotocol Label Switching (MPLS) Label Switch Controller (LSC) through a Virtual Switch Interface (VSI), or if no descriptor is specified, about all such discovered interfaces, use the **show controllers vsi descriptor** command in user EXEC or privileged EXEC mode.

show controllers vsi descriptor [*descriptor*]

Syntax Description

descriptor	(Optional) Physical descriptor. For the Cisco BPX switch, the physical descriptor has the following form: *>slot.port .>0*

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.4(20)T	This command was removed.

Usage Guidelines

Specify an interface by its (switch-supplied) physical descriptor.

Per-interface information includes the following:

- Interface name
- Physical descriptor
- Interface status
- Physical interface state (supplied by the switch)
- Acceptable VPI and VCI ranges
- Maximum cell rate
- Available cell rate (forward/backward)
- Available channels

Similar information is displayed when you enter the **show controllers xtagatm** privileged EXEC command. However, you must specify a Cisco IOS interface name instead of a physical descriptor.

Examples

The following is sample output from the **show controllers vsi descriptor** command:

```
Router# show controllers vsi descriptor 12.2.0
Phys desc: 12.2.0
Log intf:  0x000C0200 (0.12.2.0)
Interface: XTagATM0
IF status: up                  IFC state: ACTIVE
Min VPI:   1                   Maximum cell rate:  10000
Max VPI:   259                 Available channels: 2000
Min VCI:   32                  Available cell rate (forward):  10000
Max VCI:   65535               Available cell rate (backward): 10000
```

The table below describes the significant fields shown in the display.

Table 6: show controllers vsi descriptor Field Descriptions

Field	Description
Phys desc	Physical descriptor. A string learned from the switch that identifies the interface.
Log intf	Logical interface ID. This 32-bit entity, learned from the switch, uniquely identifies the interface.
Interface	The (Cisco IOS) interface name.
IF status	Overall interface status. Can be “up,” “down,” or “administratively down.”
Min VPI	Minimum virtual path identifier. Indicates the low end of the VPI range configured on the switch.
Max VPI	Maximum virtual path identifier. Indicates the high end of the VPI range configured on the switch.
Min VCI	Minimum virtual path identifier. Indicates the high end of the VCI range configured on the switch.
Max VCI	Maximum virtual channel identifier. Indicates the high end of the VCI range configured on, or determined by, the switch.

Field	Description
IFC state	Operational state of the interface, according to the switch. Can be one of the following: • FAILED_EXT (that is, an external alarm) • FAILED_INT (indicates the inability of the MPLS LSC to communicate with the VSI slave controlling the interface, or another internal failure) • REMOVED (administratively removed from the switch)
Maximum cell rate	Maximum cell rate for the interface, which has been configured on the switch (in cells per second).
Available channels	Indicates the number of channels (endpoints) that are currently free to be used for cross-connects.
Available cell rate (forward)	Cell rate that is currently available in the forward (that is, ingress) direction for new cross-connects on the interface.
Available cell rate (backward)	Cell rate that is currently available in the backward (that is, egress) direction for new cross-connects on the interface.

Related Commands

Command	Description
show controllers xtagatm	Displays information about an extended MPLS ATM interface.

show controllers vsi session

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers vsi session** command is not available in Cisco IOS software.

To display information about all sessions with Virtual Switch Interface (VSI) slaves, use the **show controllers vsi session** command in user EXEC or privileged EXEC mode.

show controllers vsi session [*session-number* [**interface** *interface*]]

Syntax Description

session-number	(Optional) Specifies the session number.
interface *interface*	(Optional) Specifies the VSI control interface.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12,0(5)T	This command was introduced.
12.4(20)T	This command was removed.

Usage Guidelines

If a session number and an interface are specified, detailed information on the individual session is presented. If the session number is specified, but the interface is omitted, detailed information on all sessions with that number is presented. (Only one session can contain a given number, because multiple control interfaces are not supported.)

Note A session consists of an exchange of VSI messages between the VSI master (the LSC) and a VSI slave (an entity on the switch). There can be multiple VSI slaves for a switch. On the BPX, each port or trunk card assumes the role of a VSI slave.

Examples

The following is sample output from the **show controllers vsi session** command:

```
Router# show controllers vsi session
Interface    Session  VCD    VPI/VCI    Switch/Slave Ids   Session State

ATM0/0       0        1      0/40       0/1                ESTABLISHED
ATM0/0       1        2      0/41       0/2                ESTABLISHED
```

```
ATM0/0       2         3       0/42       0/3               DISCOVERY
ATM0/0       3         4       0/43       0/4               RESYNC-STARTING
ATM0/0       4         5       0/44       0/5               RESYNC-STOPPING
ATM0/0       5         6       0/45       0/6               RESYNC-UNDERWAY
ATM0/0       6         7       0/46       0/7               UNKNOWN
ATM0/0       7         8       0/47       0/8               UNKNOWN
ATM0/0       8         9       0/48       0/9               CLOSING
ATM0/0       9         10      0/49       0/10              ESTABLISHED
ATM0/0       10        11      0/50       0/11              ESTABLISHED
ATM0/0       11        12      0/51       0/12              ESTABLISHED
```

The table below describes the significant fields shown in the display.

Table 7: show controllers vsi session Field Descriptions

Field	Description
Interface	Control interface name.
Session	Session number (from 0 to <*n* -1>), where *n* is the number of sessions on the control interface.
VCD	Virtual circuit descriptor (virtual circuit number). Identifies the VC carrying the VSI protocol between the master and the slave for this session.
VPI/VCI	Virtual path identifier or virtual channel identifier (for the VC used for this session).
Switch/Slave Ids	Switch and slave identifiers supplied by the switch.
Session State	Indicates the status of the session between the master and the slave. • ESTABLISHED is the fully operational steady state. • UNKNOWN indicates that the slave is not responding. Other possible states include the following: • CONFIGURING • RESYNC-STARTING • RESYNC-UNDERWAY • RESYNC-ENDING • DISCOVERY • SHUTDOWN-STARTING • SHUTDOWN-ENDING • INACTIVE

In the following example, session number 9 is specified with the **show controllers vsi session** command:

```
Router# show controllers vsi session 9
Interface:            ATM1/0        Session number:       9
VCD:                  10            VPI/VCI:              0/49
Switch type:          BPX           Switch id:            0
Controller id:        1             Slave id:             10
Keepalive timer:      15            Powerup session id:   0x0000000A
Cfg/act retry timer:  8/8           Active session id:    0x0000000A
Max retries:          10            Ctrl port log intf:   0x000A0100
Trap window:          50            Max/actual cmd wndw:  21/21
Trap filter:          all           Max checksums:        19
Current VSI version:  1             Min/max VSI version:  1/1
Messages sent:        2502          Inter-slave timer:    4.000
Messages received:    2502          Messages outstanding: 0
```

The table below describes the significant fields shown in the display.

Table 8: show controllers vsi session Field Descriptions

Field	Description
Interface	Name of the control interface on which this session is configured.
Session number	A number from 0 to <*n* -1>, where *n* is the number of slaves. Configured on the MPLS LSC with the *slaves* option of the **tag-control-protocol vsi** command.
VCD	Virtual circuit descriptor (virtual circuit number). Identifies the VC that carries VSI protocol messages for this session.
VPI/VCI	Virtual path identifier or virtual channel identifier for the VC used for this session.
Switch type	Switch device (for example, the BPX).
Switch id	Switch identifier (supplied by the switch).
Controller id	Controller identifier. Configured on the LSC, and on the switch, with the **id** option of the **tag-control-protocol vsi** command.
Slave id	Slave identifier (supplied by the switch).
Keepalive timer	VSI master keepalive timeout period (in seconds). Configured on the MPLS LSC through the **keepalive** option of the **tag-control-protocol-vsi** command. If no valid message is received by the MPLS LSC within this time period, it sends a keepalive message to the slave.
Powerup session id	Session ID (supplied by the slave) used at powerup time.

Field	Description
Cfg/act retry timer	Configured and actual message retry timeout period (in seconds). If no response is received for a command sent by the master within the actual retry timeout period, the message is re-sent. This applies to most message transmissions. The configured retry timeout value is specified through the **retry** option of the **tag-control-protocol vsi** command. The actual retry timeout value is the larger of the configured value and the minimum retry timeout value permitted by the switch.
Active session id	Session ID (supplied by the slave) for the currently active session.
Max retries	Maximum number of times that a particular command transmission will be retried by the master. That is, a message may be sent up to <max_retries+1> times. Configured on the MPLS LSC through the **retry** option of the **tag-control-protocol vsi** command.
Ctrl port log intf	Logical interface identifier for the control port, as supplied by the switch.
Trap window	Maximum number of outstanding trap messages permitted by the master. This is advertised, but not enforced, by the LSC.
Max/actual cmd wndw	Maximum command window is the maximum number of outstanding (that is, unacknowledged) commands that may be sent by the master before waiting for acknowledgments. This number is communicated to the master by the slave. The command window is the maximum number of outstanding commands that are permitted by the master, before it waits for acknowledgments. This is always less than the maximum command window.
Trap filter	This is always “all” for the LSC, indicating that it wants to receive all traps from the slave. This is communicated to the slave by the master.
Max checksums	Maximum number of checksum blocks supported by the slave.
Current VSI version	VSI protocol version currently in use by the master for this session.

Field	Description
Min/max VSI version	Minimum and maximum VSI versions supported by the slave, as last reported by the slave. If both are zero, the slave has not yet responded to the master.
Messages sent	Number of commands sent to the slave.
Inter-slave timer	Timeout value associated by the slave for messages it sends to other slaves. On a VSI-controlled switch with a distributed slave implementation (such as the BPX), VSI messages may be sent between slaves to complete their processing. For the MPLS LSC VSI implementation to function properly, the value of its retry timer is forced to be at least two times the value of the interslave timer. (See "Cfg/act retry timer" in this table.)
Messages received	Number of responses and traps received by the master from the slave for this session.
Messages outstanding	Current number of outstanding messages (that is, commands sent by the master for which responses have not yet been received).

Related Commands

Command	Description
tag-control-protocol vsi	Configures the use of VSI on a control port.

show controllers vsi status

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers vsi status** command is not available in Cisco IOS software.

To display a one-line summary of each Virtual Switch Interface (VSI)-controlled interface, use the **show controllers vsi status** command in user EXEC or in privileged EXEC mode .

show controllers vsi status

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.4(20)T	This command was removed.

Usage Guidelines

If an interface is discovered by the LSC, but no extended Multiprotocol Label Switching (MPLS) ATM interface is associated with it through the **extended-port** command, then the interface name is marked <unknown>, and interface status is marked n/a.

Examples

The following is sample output from the **show controllers vsi status** command:

```
Router# show controllers vsi status
Interface Name                 IF Status   IFC State  Physical Descriptor
switch control port                 n/a      ACTIVE  12.1.0
XTagATM0                             up      ACTIVE  12.2.0
XTagATM1                             up      ACTIVE  12.3.0
<unknown>                           n/a  FAILED-EXT  12.4.0
```

The table below describes the significant fields shown in the display.

Table 9: show controllers vsi status Field Descriptions

Field	Description
Interface Name	The (Cisco IOS) interface name.
IF Status	Overall interface status. Can be “up,” “down,” or “administratively down.”

Field	Description
IFC State	The operational state of the interface, according to the switch. Can be one of the following: • FAILED-EXT (that is, an external alarm) • FAILED-INT (indicates the inability of the MPLS LSC to communicate with the VSI slave controlling the interface, or another internal failure) • REMOVED (administratively removed from the switch)
Physical Descriptor	A string learned from the switch that identifies the interface.

show controllers vsi traffic

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers vsi traffic** command is not available in Cisco IOS software.

To display traffic information about Virtual Switch Interface (VSI)-controlled interfaces, VSI sessions, or virtual circuits (VCs) on VSI-controlled interfaces, use the **show controllers vsi traffic** command in user EXEC or privileged EXEC mode.

show controllers vsi traffic {**descriptor** *descriptor*| **session** *session-number*| **vc** [**descriptor** *descriptor* [*vpi vci*]]}

Syntax Description

descriptor *descriptor*	Displays traffic statistics for the specified descriptor.
session *session-number*	Displays traffic statistics for the specified session.
vc	Displays traffic statistics for the specified VC.
descriptor descriptor *descriptor*	Specifies the name of the physical descriptor.
vpi	Virtual path identifier (0 to 4095).
vci	Virtual circuit identifier (0 to 65535).

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(4)T	The VPI range of values was extended to 4095.
12.4(20)T	This command was removed.

Usage Guidelines

If none of the keywords is specified, traffic for all interfaces is displayed. You can specify a single interface by its (switch-supplied) physical descriptor. For the BPX switch, the physical descriptor has the form

slot.port. 0

If a session number is specified, the output displays VSI protocol traffic by message type. The VC traffic display is also displayed by the **show xmplsatm vc cross-connect traffic descriptor** command.

Examples

The following is sample output from the **show controllers vsi traffic** command:

```
Router# show controllers vsi traffic
Phys desc: 10.1.0
Interface: switch control port
IF status: n/a
Rx cells: 304250             Rx cells discarded: 0
Tx cells: 361186             Tx cells discarded: 0
Rx header errors: 4294967254 Rx invalid addresses (per card): 80360
Last invalid address: 0/53

Phys desc: 10.2.0
Interface: XTagATM0
IF status: up
Rx cells: 202637             Rx cells discarded: 0
Tx cells: 194979             Tx cells discarded: 0
Rx header errors: 4294967258 Rx invalid addresses (per card): 80385
Last invalid address: 0/32

Phys desc: 10.3.0
Interface: XTagATM1
IF status: up
Rx cells: 182295             Rx cells discarded: 0
Tx cells: 136369             Tx cells discarded: 0
Rx header errors: 4294967262 Rx invalid addresses (per card): 80372
Last invalid address: 0/32
```

The table below describes the significant fields shown in the display.

Table 10: show controllers vsi traffic Field Descriptions

Field	Description
Phys desc	Physical descriptor of the interface.
Interface	The Cisco (IOS) interface name.
Rx cells	Number of cells received on the interface.
Tx cells	Number of cells transmitted on the interface.
Rx cells discarded	Number of cells received on the interface that were discarded due to traffic management.
Tx cells discarded	Number of cells that could not be transmitted on the interface due to traffic management and which were therefore discarded.
Rx header errors	Number of cells that were discarded due to ATM header errors.
Rx invalid addresses	Number of cells received with an invalid address (that is, an unexpected VPI/VCI combination). With the Cisco BPX switch, this count is of all such cells received on all interfaces in the port group of this interface.

Field	Description
Last invalid address	Number of cells received on this interface with ATM cell header errors.

The following sample output is displayed when you enter the **show controllers vsi traffic session 9** command:

```
Router# show controllers vsi traffic session 9
                         Sent                                Received
Sw Get Cnfg Cmd:          3656       Sw Get Cnfg Rsp:          3656
Sw Cnfg Trap Rsp:         0          Sw Cnfg Trap:             0
Sw Set Cnfg Cmd:          1          Sw Set Cnfg Rsp:          1
Sw Start Resync Cmd:      1          Sw Start Resync Rsp:      1
Sw End Resync Cmd:        1          Sw End Resync Rsp:        1
Ifc Getmore Cnfg Cmd:     1          Ifc Getmore Cnfg Rsp:     1
Ifc Cnfg Trap Rsp:        4          Ifc Cnfg Trap:            4
Ifc Get Stats Cmd:        8          Ifc Get Stats Rsp:        8
Conn Cmt Cmd:             73         Conn Cmt Rsp:             73
Conn Del Cmd:             50         Conn Del Rsp:             0
Conn Get Stats Cmd:       0          Conn Get Stats Rsp:       0
Conn Cnfg Trap Rsp:       0          Conn Cnfg Trap:           0
Conn Bulk Clr Stats Cmd: 0           Conn Bulk Clr Stats Rsp: 0
Gen Err Rsp:              0          Gen Err Rsp:              0
unused:                   0          unused:                   0
unknown:                  0          unknown:                  0
TOTAL:                    3795       TOTAL:                    3795
```

The table below describes the significant fields shown in the display.

Table 11: show controllers vsi traffic session Field Descriptions

Field	Description
Sw Get Cnfg Cmd	Number of VSI “get switch configuration command” messages sent.
Sw Cnfg Trap Rsp	Number of VSI “switch configuration asynchronous trap response” messages sent.
Sw Set Cnfg Cmd	Number of VSI “set switch configuration command” messages sent.
Sw Start Resync Cmd	Number of VSI “set resynchronization start command” messages sent.
Sw End Resync Cmd	Number of VSI “set resynchronization end command” messages sent.
Ifc Getmore Cnfg Cmd	Number of VSI “get more interfaces configuration command” messages sent.
Ifc Cnfg Trap Rsp	Number of VSI “interface configuration asynchronous trap response” messages sent.
Ifc Get Stats Cmd	Number of VSI “get interface statistics command” messages sent.

Field	Description
Conn Cmt Cmd	Number of VSI “set connection committed command” messages sent.
Conn Del Cmd	Number of VSI “delete connection command” messages sent.
Conn Get Stats Cmd	Number of VSI “get connection statistics command” messages sent.
Conn Cnfg Trap Rsp	Number of VSI “connection configuration asynchronous trap response” messages sent.
Conn Bulk Clr Stats Cmd	Number of VSI “bulk clear connection statistics command” messages sent.
Gen Err Rsp	Number of VSI “generic error response” messages sent or received.
Sw Get Cnfg Rsp	Number of VSI “get connection configuration command response” messages received.
Sw Cnfg Trap	Number of VSI “switch configuration asynchronous trap” messages received.
Sw Set Cnfg Rsp	Number of VSI “set switch configuration response” messages received.
Sw Start Resync Rsp	Number of VSI “set resynchronization start response” messages received.
Sw End Resync Rsp	Number of VSI “set resynchronization end response” messages received.
Ifc Getmore Cnfg Rsp	Number of VSI “get more interfaces configuration response” messages received.
Ifc Cnfg Trap	Number of VSI “interface configuration asynchronous trap” messages received.
Ifc Get Stats Rsp	Number of VSI “get interface statistics response” messages received.
Conn Cmt Rsp	Number of VSI “set connection committed response” messages received.
Conn Del Rsp	Number of VSI “delete connection response” messages received.

Field	Description
Conn Get Stats Rsp	Number of VSI “get connection statistics response” messages received.
Conn Cnfg Trap	Number of VSI “connection configuration asynchronous trap” messages received.
Conn Bulk Clr Stats Rsp	Number of VSI “bulk clear connection statistics response” messages received.
unused, unknown	“Unused” messages are those whose function codes are recognized as being part of the VSI protocol, but which are not used by the MPLS LSC and, consequently, are not expected to be received or sent. “Unknown” messages have function codes that the MPLS LSC does not recognize as part of the VSI protocol.
TOTAL	Total number of VSI messages sent or received.

show controllers xtagatm

Note Effective with Cisco IOS Release 12.4(20)T, the **show controllers xtagatm** command is not available in Cisco IOS software.

To display information about an extended Multiprotocol Label Switching (MPLS) ATM interface controlled through the Virtual Switch Interface (VSI) protocol (or, if an interface is not specified, to display information about all extended MPLS ATM interfaces controlled through the VSI protocol), use the **show controllers xtagatm** command in user EXEC or privileged EXEC mode.

show controllers xtagatm *if-number*

Syntax Description

if-number	Specifies the interface number.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.4(20)T	This command was removed.

Usage Guidelines

Per-interface information includes the following:

- Interface name
- Physical descriptor
- Interface status
- Physical interface state (supplied by the switch)
- Acceptable VPI and VCI ranges
- Maximum cell rate
- Available cell rate (forward/backward)
- Available channels

Similar information appears if you enter the **show controllers vsi descriptor** command. However, you must specify an interface by its (switch-supplied) physical descriptor, instead of its Cisco IOS interface name. For the Cisco BPX switch, the physical descriptor has the form *slot.port.0*.

Examples

In this example, the sample output is from the **show controllers xtagatm** command specifying interface 0:

```
Router# show controllers xtagatm 0
Interface XTagATM0 is up
Hardware is Tag-Controlled ATM Port (on BPX switch BPX-VSI1)
Control interface ATM1/0 is up
Physical descriptor is 10.2.0
Logical interface 0x000A0200 (0.10.2.0)
Oper state ACTIVE, admin state UP
VPI range 1-255, VCI range 32-65535
VPI is not translated at end of link
Tag control VC need not be strictly in VPI/VCI range
Available channels: ingress 30, egress 30
Maximum cell rate: ingress 300000, egress 300000
Available cell rate: ingress 300000, egress 300000
Endpoints in use: ingress 7, egress 8, ingress/egress 1
Rx cells 134747
rx cells discarded 0, rx header errors 0
rx invalid addresses (per card): 52994
last invalid address 0/32
Tx cells 132564
tx cells discarded: 0
```

The table below describes the significant fields shown in the display.

Table 12: show controllers xtagatm Field Descriptions

Field	Description
Interface XTagATM0 is up	Indicates the overall status of the interface. May be “up,” “down,” or “administratively down.”
Hardware is Tag-Controlled ATM Port	Indicates the hardware type. If the XTagATM was successfully associated with a switch port, a description of the form (on <switch_type> switch <name>) follows this field, where <switch_type> indicates the type of switch (for example, BPX), and the name is an identifying string learned from the switch. If the XTagATM interface was not bound to a switch interface (with the **extended-port** interface configuration command), then the label “Not bound to a control interface and switch port” appears. If the interface has been bound, but the target switch interface has not been discovered by the LSC, then the label “Bound to undiscovered switch port (id <number>)” appears, where <number> is the logical interface ID in hexadecimal notation.
Control interface ATM1/0 is up	Indicates that the XTagATM interface was bound (with the **extended-port** interface configuration command) to the VSI master whose control interface is ATM1/0 and that this control interface is up.

Field	Description
Physical descriptor is...	A string identifying the interface that was learned from the switch.
Logical interface	This 32-bit entity, learned from the switch, uniquely identifies the interface. It appears in both hexadecimal and dotted quad notation.
Oper state	Operational state of the interface, according to the switch. Can be one of the following: • ACTIVE • FAILED_EXT (that is, an external alarm) • FAILED_INT (indicates the inability of the MPLS LSC to communicate with the VSI slave controlling the interface, or another internal failure) • REMOVED (administratively removed from the switch)
admin state	Administrative state of the interface, according to the switch--either “Up” or “Down.”
VPI range 1 to 255	Indicates the allowable VPI range for the interface that was configured on the switch.
VCI range 32 to 65535	Indicates the allowable VCI range for the interface that was configured on, or determined by, the switch.
LSC control VC need not be strictly in VPI or VCI range	Indicates that the label control VC does not need to be within the range specified by VPI range, but may be on VPI 0 instead.
Available channels	Indicates the number of channels (endpoints) that are currently free to be used for cross-connects.
Maximum cell rate	Maximum cell rate for the interface, which was configured on the switch.
Available cell rate	Cell rate that is currently available for new cross-connects on the interface.

Field	Description
Endpoints in use	Number of endpoints (channels) in use on the interface, broken down by anticipated traffic flow, as follows: • Ingress--Endpoints carry traffic into the switch • Egress--Endpoints carry traffic away from the switch • Ingress/egress--Endpoints carry traffic in both directions
Rx cells	Number of cells received on the interface.
rx cells discarded	Number of cells received on the interface that were discarded due to traffic management actions (rx header errors).
rx header errors	Number of cells received on the interface with cell header errors.
rx invalid addresses (per card)	Number of cells received with invalid addresses (that is, unexpected VPI or VCI.). On the BPX, this counter is maintained per port group (not per interface).
last invalid address	Address of the last cell received on the interface with an invalid address (for example, 0/32).
Tx cells	Number of cells sent from the interface.
tx cells discarded	Number of cells intended for transmission from the interface that were discarded due to traffic management actions.

Related Commands

Command	Description
show controllers vsi descriptor	Displays information about a switch interface discovered by the MPLS LSC through the VSI.

show interface pseudowire

To display information about the pseudowire interface, use the **show interface pseudowire** command in privileged EXEC mode.

show interface pseudowire *number*

Syntax Description

number	Interface pseudowire number.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
Cisco IOS XE Release 3.7S	This command was introduced as part of the Multiprotocol Label Switching (MPLS)-based Layer 2 VPN (L2VPN) command modifications for cross-OS support.
15.3(1)S	This command was integrated as part of the Multiprotocol Label Switching (MPLS)-based Layer 2 VPN (L2VPN) command modifications for cross-OS support.

Examples

The following is sample output from the **show interface pseudowire** command. The output fields are self-explanatory.

```
Device# show interface pseudowire 100

pseudowire 100 is up
    Description: L2VPN Pseudowire
    MTU 1500 bytes, BW 1000000 Kbit
    Encapsulation: MPLS
                   Peer Address: 10.0.0.1, VC ID: 10
    RX
      0 unicast packets  0 multicast packets
      0 input packets  0 bit rate  0 packet rate
    TX
      0 unicast packets  0 multicast packets
      0 output packets  0 bits/sec  0 packets/sec
```

show interface tunnel configuration

To display the configuration of a mesh tunnel interface, use the **show interface tunnel configuration** command in privileged EXEC mode.

show interface tunnel *num* **configuration**

Syntax Description

num	Number of the mesh tunnel for which you want to display configuration information.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.0(27)S	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.

Usage Guidelines

The space before the *num* argument is optional.

Use this command to show the running configuration of the mesh tunnel interface.

Examples

The following command output shows the configuration of mesh tunnel interface 5:

```
Router# show interface tunnel 5 configuration
interface tunnel 5
 ip unnumbered Loopback0
 no ip directed-broadcast
 no keepalive
 tunnel destination access-list 1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng path-option 1 dynamic
```

The table below describes the significant fields shown in the display.

Table 13: show interface tunnel configuration Field Descriptions

Field	Description
ip unnumbered Loopback0	Indicates the type and number of another interface on which the router has an assigned IP address. It cannot be another unnumbered interface.
no ip directed-broadcast	Indicates that no IP broadcast addresses are used for the mesh tunnel interface.
no keepalive	Indicates that no keepalives are set for the mesh tunnel interface.
tunnel destination access-list 1	Indicates that access-list 1 is the access list that the template interface will use for obtaining the mesh tunnel interface destination address.
tunnel mode mpls traffic-eng	Indicates that the mode of the mesh tunnel is set to Multiprotocol Label Switching (MPLS) for traffic engineering.
tunnel mpls traffic-eng autoroute announce	Indicates that the Interior Gateway Protocol (IGP) should use the tunnel (if the tunnel is up) in its enhanced shortest path first (SPF) calculation.
tunnel mpls traffic-eng path-option 1 dynamic	Indicates that a path option (path-option1) for the label switch router (LSR) for the MPLS traffic engineering (TE) mesh tunnel is configured dynamically.

Related Commands

Command	Description
tunnel destination access-list	Specifies the access list that the template interface will use for obtaining the mesh tunnel interface destination address.

show interface virtual-ethernet

To display status and information about a virtual Ethernet interface, use the **show interface virtual-ethernet** command in user privileged EXEC mode.

show interface virtual-ethernet *num* [**switchport**| **transport**]

Syntax Description

num	The number of the virtual interface.
switchport	Show virtual Ethernet instance switchport information.
transport	Show virtual Ethernet instance transport information.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SXI4	This command was introduced.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Examples

The following example shows transport information for virtual Ethernet interface 1:

```
Router# show interface virtual-ethernet 1 transport

VLAN Transport type for the V-E instance: VPLS Mesh
  11 VPLS domains provisioned for this V-E instance
  VFI names : VFI[45-55]_
```

The following example shows switchport information for virtual Ethernet interface 1:

```
Router# show interface virtual-ethernet 1 switchport
Name: VE1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: up
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Trunking VLANs Enabled: 100,200
```

Related Commands

Command	Description
interface virtual-ethernet	Creates a virtual Ethernet interface.

show interface xtagatm

Note Effective with Cisco IOS Release 12.4(20)T, the **show interface xtagatm** command is not available in Cisco IOS software.

To display information about an extended Multiprotocol Label Switching (MPLS) ATM interface, use the **show interface xtagatm** command in user EXEC or privileged EXEC mode.

show interface xtagatm *if-number*

Syntax Description

if-number	Specifies the MPLS ATM interface number.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.3T	Sample command output was added for when an interface is down.
12.4(20)T	This command was removed.

Usage Guidelines

Extended MPLS ATM interfaces are virtual interfaces that are created on first reference like tunnel interfaces. Extended MPLS ATM interfaces are similar to ATM interfaces except that the former only supports LC-ATM encapsulation.

Examples

The following is sample command output when an interface is down:

```
Router# show interface xt92
XTagATM92 is down, line protocol is down
 Hardware is Tag-Controlled Switch Port
 Interface is unnumbered. Using address of Loopback1 (15.15.15.15)
 MTU 4470 bytes, BW 4240 Kbit, DLY 80 used,
 reliability 186/255, txload ½55, rxload ½55
 Encapsulation ATM, loopback not set
 Keepalive set (10 sec) [00:00:08/4]
 Encapsulation(s): AAL5
 Control interface: not configured
 0 terminating VCs
 Switch port traffic:
  ? cells input, ? cells output
 Last input 00:00:10, output never, output hang never
 Last clearing of "show interface" counters never
 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
```

```
 Queueing strategy: fifo
 Output queue: 0/0 (size/max)
 Terminating traffic:
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
  138 packets input, 9193 bytes, 0 no buffer
  Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 I
  00:05:46: %SYS-5-CONFIG_I: Configured from console by consolegnored, 0 abort
  142 packets output, 19686 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 output buffer failures, 0 output buffers swapped out
```

The following is sample command output when an interface is up:

```
Router# show interface xt92
XTagATM92 is up, line protocol is up
Hardware is Tag-Controlled Switch Port
Interface is unnumbered. Using address of Loopback1 (15.15.15.15)
MTU 4470 bytes, BW 4240 Kbit, DLY 80 used,
reliability 174/255, txload ½55, rxload ½55
Encapsulation ATM, loopback not set
Keepalive set (10 sec)
Encapsulation(s): AAL5
Control interface: ATM3/0, switch port: bpx 9.2
3 terminating VCs, 7 switch cross-connects
Switch port traffic:
275 cells input, 273 cells output
Last input 00:00:00, output never, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/0 (size/max)
Terminating traffic:
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
127 packets input, 8537 bytes, 0 no buffer
Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
131 packets output, 18350 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 output buffer failures, 0 output buffers swapped out
```

The table below describes the significant fields shown in the displays.

Table 14: show interface xtagatm Field Descriptions

Field	Description
XTagATM0 is up XTagATM0 is down	Interface is currently active (up) or inactive (down).
line protocol is up line protocol is down	Displays the line protocol as up or down.
Hardware is Tag-Controlled Switch Port	Specifies the hardware type.
Interface is unnumbered	Specifies that this is an unnumbered interface.
MTU	Maximum transmission unit of the extended MPLS ATM interface.
BW	Bandwidth of the interface (in kBps).
DLY	Delay of the interface in microseconds.

Field	Description
reliability	Reliability of the interface as a fraction of 255 (255/255 is 100% reliability), calculated as an exponential average over 5 minutes.
Encapsulation ATM	Encapsulation method.
loopback not set	Indicates that loopback is not set.
Keepalive set (10 sec) [00:00:08/4]	Indicates why the Xtag line is down. Valid values are: 1--Internal usage. 2--Administratively down. 3--Internal usage. 4--No extended port is configured. 5--Some cross-connects from an old session have been left operational. 6--No extended port or a wrong extended port was configured. 7--No control port was configured. 8--Internal usage. 9--Internal usage. 10--Internal usage. 11--Internal usage. 12--External port. The XTag is mapped to an invalid port on the switch. 13--External port. The XTag is mapped to a port that is down. 14--External port is mapped to the control panel on the switch. 15--OAM is being used to track the link state. The neighbor may be down or it is not responding to the OAM calls.
Encapsulation(s)	Identifies the ATM adaptation layer.
Control interface	Identifies the control port switch port with which the extended MPLS ATM interface has been associated through the **extended-port** interface configuration command.

Field	Description
n terminating VCs	Number of terminating VCs with an endpoint on this extended MPLS ATM interface. Packets are sent or received by the MPLS LSC on a terminating VC, or are forwarded between an LSC-controlled switch port and a router interface.
7 switch cross-connects	Number of switch cross-connects on the external switch with an endpoint on the switch port that corresponds to this interface. This includes cross-connects to terminating VCs that carry data to and from the LSC, and cross-connects that bypass the MPLS LSC and switch cells directly to other ports.
Switch port traffic	Number of cells received and sent on all cross-connects associated with this interface.
Terminating traffic	Indicates that counters below this line apply only to packets sent or received on terminating VCs.
5-minute input rate, 5-minute output rate	Average number of bits and packets sent per second in the last 5 minutes.
packets input	Total number of error-free packets received by the system.
bytes	Total number of bytes, including data and MAC encapsulation, in the error-free packets received by the system.
no buffer	Number of received packets discarded because there was no buffer space in the main system. Compare with ignored count. Broadcast storms on Ethernet systems and bursts of noise on serial lines are often responsible for no input buffer events.
broadcasts	Total number of broadcast or multicast packets received by the interface.
runts	Number of packets that are discarded because they are smaller than the medium's minimum packet size.
giants	Number of packets that are discarded because they exceed the medium's maximum packet size.
input errors	Total number of no buffer, runts, giants, CRCs, frame, overrun, ignored and abort counts. Other input-related errors can also increment the count, so that this sum may not balance with other counts.

Field	Description
CRC	Cyclic redundancy checksum generated by the originating LAN station or far-end device does not match the checksum calculated from the data received. On a LAN, this usually indicates noise or transmission problems on the LAN interface or the LAN bus. A high number of CRCs is usually the result of traffic collisions or a station sending bad data. On a serial link, CRCs usually indicate noise, gain hits, or other transmission problems on the data link.
frame	Number of packets received incorrectly having a CRC error and a noninteger number of octets.
overrun	Number of times the serial receiver hardware was unable to hand received data to a hardware buffer because the input rate exceeded the receiver's ability to handle the data.
ignored	Number of received packets ignored by the interface because the interface hardware ran low on internal buffers. These buffers are different from the system buffers mentioned previously in the buffer description. Broadcast storms and bursts of noise can cause the ignored count to be incremented.
abort	Illegal sequence of one bits on the interface. This usually indicates a clocking problem between the interface and the data-link equipment.
packets output	Total number of messages sent by the system.
bytes	Total number of bytes, including data and MAC encapsulation, sent by the system.
underruns	Number of times that the sender has been running faster than the router can handle data. This condition may never be reported on some interfaces.
output errors	Sum of all errors that prevented the final transmission of datagrams out of the interface being examined. Note that this may not balance with the sum of the enumerated output errors, because some datagrams may have more than one error, and others may have errors that do not fall into any of the specifically tabulated categories.

Field	Description
collisions	Number of messages re-sent due to an Ethernet collision. This is usually the result of an overextended LAN (Ethernet or transceiver cable too long, more than two repeaters between stations, or too many cascaded multiport transceivers). A packet that collides is counted only one time in output packets.
interface resets	Number of times an interface has been completely reset. Resets occur if packets queued for transmission were not sent within several seconds. On a serial line, this can be caused by a malfunctioning modem that is not supplying the transmit clock signal, or by a cable problem. If the system notices that the carrier detect line of a serial interface is up, but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.

Related Commands

Command	Description
interface xtagatm	Enters configuration mode for an extended MPLS ATM (XTagATM) interface.

show ip bgp l2vpn

To display Layer 2 Virtual Private Network (L2VPN) address family information from the Border Gateway Protocol (BGP) table, use the **show ip bgp l2vpn** command in user EXEC or privileged EXEC mode.

With BGP show Command Argument

show ip bgp l2vpn vpls {**all** | [**summary** | [**slow**]| **ve-id** *id-value*]| {**block-offset** | [*value*]}| **rd** {*route-distinguisher* | [**ve-id** | {**block-offset** | [*value*]}]}} [*bgp-keyword*]

With IP Prefix and Mask Length Syntax

show ip bgp l2vpn vpls {**all**| **rd** *route-distinguisher*} [*ip-prefix*/*length* [**[bestpath]**] [**longer-prefixes** [**[injected]**]] [**[multipaths]**] [**shorter-prefixes** [[*mask-length*]]] [**[subnets]**]]

With Network Address Syntax

show ip bgp l2vpn vpls {**all**| **rd** *route-distinguisher*} [*network-address* [*mask*| **bestpath**| **multipaths**] **[bestpath]** [**longer-prefixes [injected]**] **[multipaths]** [**shorter-prefixes** [*mask-length*]] **[subnets]**]

Syntax Description

vpls	Displays L2VPN address family database information for the Virtual Private LAN Service (VPLS) subsequent address family identifier (SAFI).
all	Displays the complete L2VPN database.
rd *route-distinguisher*	Displays prefixes that match the specified route distinguisher.
ve-id *id-value*	(Optional) Displays the target VPLS Endpoint (VE) ID and ID value.
summary	(Optional) Displays a summary of BGP neighbor status.
slow	(Optional) Displays a summary of slow-peer status.
block-offset *value*	Displays the target block-offset value.
bgp-keyword	(Optional) Argument representing a **show ip bgp** command keyword that can be added to this command. See the table below.
ip-prefix/*length*	(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
bestpath	(Optional) Displays the best path for the specified prefix.

longer-prefixes	(Optional) Displays the route and more specific routes.
injected	(Optional) Displays more specific routes that were injected because of the specified prefix.
multipaths	(Optional) Displays the multipaths for the specified prefix.
shorter-prefixes	(Optional) Displays the less specific routes.
mask-length	(Optional) The length of the mask as a number in the range from 0 to 32. Prefixes longer than the specified mask length are displayed.
subnets	(Optional) Displays the subnet routes for the specified prefix.
network-address	(Optional) The IP address of a network in the BGP routing table.
mask	(Optional) The mask of the network address, in dotted decimal format.

Command Default

If no arguments or keywords are specified, this command displays the complete L2VPN database.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SRB	This command was introduced.
Cisco IOS XE2.6	This command was integrated into Cisco IOS XE Release 2.6.
Cisco IOS XE3.8S	This command was modified. RFC4761 is fully supported in Cisco IOS XE Release 3.8S.

Usage Guidelines

The table below displays optional **show ip bgp** command keywords that can be configured with the **show ip bgp l2vpn** command. Replace the *bgp-keyword* argument with the appropriate keyword from the table. For more details about each command in its **show ip bgp** *bgp-keyword* form, see the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols,* Release 12.2.

Table 15: Optional show ip bgp Command Keywords and Descriptions

Keyword	Description
community	Displays routes that match a specified community.
community-list	Displays routes that match a specified community list.
dampening	Displays paths suppressed because of dampening (BGP route from peer is up and down).
extcommunity-list	Displays routes that match a specified extcommunity list.
filter-list	Displays routes that conform to the filter list.
inconsistent-as	Displays only routes that have inconsistent autonomous systems of origin.
neighbors	Displays details about TCP and BGP neighbor connections.
oer-paths	Displays all OER-managed path information.
paths [*regexp*]	Displays autonomous system path information. If the optional *regexp* argument is entered, the autonomous system paths that are displayed match the autonomous system path regular expression.
peer-group	Displays information about peer groups.
pending-prefixes	Displays prefixes that are pending deletion.
prefix-list	Displays routes that match a specified prefix list.
quote-regexp	Displays routes that match the quoted autonomous system path regular expression.
regexp	Displays routes that match the autonomous system path regular expression.
replication	Displays the replication status update groups.
route-map	Displays routes that match the specified route map.
rt-filter-list	Displays the specified inbound route target filter list.
summary	Displays a summary of BGP neighbor status.
update-group	Displays information on update groups.

Examples

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **all** keywords are used to display the complete L2VPN database:

```
Device# show ip bgp l2vpn vpls all

BGP table version is 5, local router ID is 192.168.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 45000:100
*> 45000:100:172.17.1.1/96
                    0.0.0.0                              32768 ?
*>i45000:100:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
Route Distinguisher: 45000:200
*> 45000:200:172.17.1.1/96
                    0.0.0.0                              32768 ?
*>i45000:200:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
```

The table below describes the significant fields shown in the display.

Table 16: show ip bgp l2vpn vpls all Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: • s—The table entry is suppressed. • d—The table entry is dampened. • h—The table entry is a historical entry. • *—The table entry is valid. • >—The table entry is the best entry to use for that network. • i—The table entry was learned via an internal BGP (iBGP) session. • r—The table entry failed to install in the routing information base (RIB) table. • S—The table entry is Stale (old). This entry is useful in BGP graceful restart situations.

Field	Description
Origin codes	Origin of the entry. The origin code is displayed at the end of each line in the table. It can be one of the following values: • i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. • e—Entry originated from an Exterior Gateway Protocol (EGP). • ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the interautonomous system metric.
LocPrf	Local preference value as set with the **set local-preference** command in route-map configuration mode. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.
Route Distinguisher	Route distinguisher that identifies a set of routing and forwarding tables used in virtual private networks.

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **all** keywords are used to display information about all VPLS BGP signaling prefixes (including local generated and received from remote):

```
Device#show ip bgp l2vpn vpls all

BGP table version is 14743, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

```
Network               Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:1
*>i 65000:1:VEID-3:Blk-1/136
                      3.3.3.3                  0    100       0 ?
*> 65000:1:VEID-4:Blk-1/136
                      0.0.0.0                             32768 ?
*>i 65000:1:VEID-5:Blk-1/136
                      2.2.2.2                  0    100       0 ?
*>i 65000:1:VEID-6:Blk-1/136
                      4.4.4.4                  0    100       0 ?
Route Distinguisher: 65000:2
*> 65000:2:VEID-20:Blk-20/136
                      0.0.0.0                             32768 ?
*>i 65000:2:VEID-21:Blk-20/136
                      2.2.2.2                  0    100       0 ?
*>i 65000:2:VEID-22:Blk-20/136
                      3.3.3.3                  0    100       0 ?
*>i 65000:2:VEID-23:Blk-20/136
                      4.4.4.4                  0    100       0 ?
```

The following example shows output for the **show ip bgp l2vpn** command when the **vpls**, **all** and **summary** keywords are used to display information about the L2VPN VPLS address family:

```
Device# show ip bgp l2vpn vpls all summary

BGP router identifier 10.1.1.1, local AS number 65000
BGP table version is 14743, main routing table version
14743
6552 network entries using 1677312 bytes of memory
6552 path entries using 838656 bytes of memory
3276/3276 BGP path/bestpath attribute entries using
760032 bytes of memory
1638 BGP extended community entries using 65520 bytes of
memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3341520 total bytes of memory
BGP activity 9828/3276 prefixes, 9828/3276 paths, scan
interval 60 secs
Neighbor       V          AS  MsgRcvd  MsgSent  TblVer  InQ OutQ  Up/Down
State/PfxRcd
10.2.2.2       4        65000   90518   90507    14743   0     0  8w0d     1638
10.3.3.3       4        65000    4901    4895    14743   0     0  2d01h    1638
10.4.4.4       4        65000    4903    4895    14743   0     0  2d01h    1638
```

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **rd** *rd* keywords are used to display information about all VPLS BGP signaling prefixes with the specified rd, i.e. the same VPLS instance:

```
Device# show ip bgp l2vpn vpls rd 65000:3

BGP table version is 14743, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:3
*> 65000:3:VEID-30:Blk-30/136
                      0.0.0.0                             32768 ?
*>i 65000:3:VEID-31:Blk-30/136
                      2.2.2.2                  0  100       0 ?
*>i 65000:3:VEID-32:Blk-30/136
                      3.3.3.3                  0  100       0 ?
*>i 65000:3:VEID-33:Blk-30/136
                      4.4.4.4                  0  100       0 ?
```

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **rd** keywords are used to display the L2VPN information that matches the route distinguisher 45000:100. Note that the information displayed is a subset of the information displayed using the **all** keyword.

```
Device# show ip bgp l2vpn vpls rd 45000:100

BGP table version is 5, local router ID is 192.168.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 45000:100
*> 45000:100:172.17.1.1/96
                    0.0.0.0                              32768 ?
*>i45000:100:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
```

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **all** keywords are used to display information about an individual prefix:

```
Device# show ip bgp l2vpn vpls all ve-id 31 block 30

BGP routing table entry for 65000:3:VEID-31:Blk-30/136, version 11
Paths: (1 available, best #1, table L2VPN-VPLS-BGP-Table)
   Not advertised to any peer
   Refresh Epoch 2
   Local
     2.2.2.2 (metric 2) from 2.2.2.2 (2.2.2.2)
       Origin incomplete, metric 0, localpref 100, valid, internal, best
       AGI version(0), VE Block Size(10) Label Base(16596)
       Extended Community: RT:65000:3 L2VPN L2:0x0:MTU-1500
       rx pathid: 0, tx pathid: 0x0
                0    100      0 ?
```

Related Commands

Command	Description
address-family l2vpn	Enters address family configuration mode to configure a routing session using L2VPN endpoint provisioning information.
show bgp l2vpn vpls	Displays L2VPN VPLS address family information from the BGP table.

show ip bgp labels

To display information about Multiprotocol Label Switching (MPLS) labels from the external Border Gateway Protocol (eBGP) route table, use the **show ip bgp labels** command in privileged EXEC mode.

show ip bgp labels

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(21)ST	This command was introduced.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series router.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
15.2(2)SNG	This command was integrated into Cisco ASR 901 Series Aggregation Services Routers.

Usage Guidelines

Use this command to display eBGP labels associated with an Autonomous System Boundary Router (ASBR).

This command displays labels for BGP routes in the default table only. To display labels in the Virtual Private Network (VPN) routing and forwarding (VRF) tables, use the **show ip bgp vpnv4** {**all** | **vrf** *vrf-name*} command with the optional **labels** keyword.

Examples

The following example shows output for an ASBR using BGP as a label distribution protocol:

```
Router# show ip bgp labels
Network          Next Hop         In Label/Out Label
10.3.0.0/16       0.0.0.0          imp-null/exp-null
10.15.15.15/32   10.15.15.15      18/exp-null
10.16.16.16/32   0.0.0.0          imp-null/exp-null
10.17.17.17/32   10.0.0.1         20/exp-null
```

```
10.18.18.18/32   10.0.0.1         24/31
10.18.18.18/32   10.0.0.1         24/33
```

The table below describes the significant fields shown in the display.

Table 17: show ip bgp labels Field Descriptions

Field	Description
Network	Displays the network address from the eBGP table.
Next Hop	Specifies the eBGP next hop address.
In Label	Displays the label (if any) assigned by this router.
Out Label	Displays the label assigned by the BGP next hop router.

Related Commands

Command	Description
show ip bgp vpnv4	Displays VPN address information from the BGP table.

show ip bgp neighbors

To display information about Border Gateway Protocol (BGP) and TCP connections to neighbors, use the **show ip bgp neighbors** command in user or privileged EXEC mode.

show ip bgp [**ipv4** {**multicast**| **unicast**}| **vpnv4 all**| **vpnv6 unicast all**] **neighbors** [**slow**| *ip-address*| *ipv6-address* [**advertised-routes**| **dampened-routes**| **flap-statistics**| **paths** [*reg-exp*]| **policy [detail]**| **received prefix-filter**| **received-routes**| **routes**]]

Syntax Description

ipv4	(Optional) Displays peers in the IPv4 address family.
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vpnv4 all	(Optional) Displays peers in the VPNv4 address family.
vpnv6 unicast all	(Optional) Displays peers in the VPNv6 address family.
slow	(Optional) Displays information about dynamically configured slow peers.
ip-address	(Optional) IP address of the IPv4 neighbor. If this argument is omitted, information about all neighbors is displayed.
ipv6-address	(Optional) IP address of the IPv6 neighbor.
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.
dampened-routes	(Optional) Displays the dampened routes received from the specified neighbor.
flap-statistics	(Optional) Displays the flap statistics of the routes learned from the specified neighbor (for external BGP peers only).
paths *reg-exp*	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.
policy	(Optional) Displays the policies applied to this neighbor per address family.

detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, access control lists (ACLs), and autonomous system path filter lists.
received prefix-filter	(Optional) Displays the prefix list (outbound route filter [ORF]) sent from the specified neighbor.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The output displayed when this keyword is entered is a subset of the output displayed by the **received-routes** keyword.

Command Default

The output of this command displays information for all neighbors.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Mainline and T Release	Modification
10.0	This command was introduced.
11.2	This command was modified. The **received-routes** keyword was added.
12.2(4)T	This command was modified. The **received** and **prefix-filter** keywords were added.
12.2(15)T	This command was modified. Support for the display of BGP graceful restart capability information was added.
12.3(7)T	This command was modified. The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.4(4)T	This command was modified. Support for the display of Bidirectional Forwarding Detection (BFD) information was added.
12.4(11)T	This command was modified. Support for the **policy** and **detail** keywords was added.
12.4(20)T	This command was modified. The output was modified to support BGP TCP path MTU discovery.

Mainline and T Release	Modification
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation was added.

Command History

S Release	Modification
12.0(18)S	This command was modifed. The output was modified to display the no-prepend configuration option.
12.0(21)ST	This command was modifed. The output was modified to display Multiprotocol Label Switching (MPLS) label information.
12.0(22)S	This command was modified. Support for the display of BGP graceful restart capability information was added. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was also added.
12.0(25)S	This command was modified. The **policy** and **detail** keywords were added.
12.0(27)S	This command was modified. The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.0(31)S	This command was modified. Support for the display of BFD information was added.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.2(18)SXE	This command was modified. Support for the display of BFD information was added.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was modified. The output was modified to support BGP TCP path Maximum Transmission Unit (MTU) discovery.
12.2(33)SRB	This command was modified. Support for the **policy** and **detail** keywords was added.

S Release	Modification
12.2(33)SXH	This command was modified. Support for displaying BGP dynamic neighbor information was added.
12.2(33)SRC	This command was modified. Support for displaying BGP graceful restart information was added.
12.2(33)SB	This command was modified. Support for displaying BFD and the BGP graceful restart per peer information was added, and support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.2(33)SB.
12.2(33)SXI1	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)SRE	This command was modified. Support for displaying BGP best external and BGP additional path features information was added. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
15.0(1)S	This command was modified. The **slow** keyword was added.
15.0(1)SY	This command was integrated into Cisco IOS Release 15.0(1)SY.
15.1(1)S	This command was modified. The Layer 2 VPN address family is displayed if graceful restart or nonstop forwarding (NSF) is enabled.
15.1(1)SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
15.2(4)S	This command was modified and implemented on the Cisco 7200 series router. The configured discard and treat-as-withdraw attributes are displayed, along with counts of incoming Updates with a matching discard attribute or treat-as-withdraw attribute, and number of times a malformed Update is treat-as-withdraw. The capabilities of the neighbor to send and receive additional paths that are advertised or received are added.
15.1(2)SNG	This command was implemented on the Cisco ASR 901 Series Aggregation Services Routers.
15.3(2)T	This command was integrated into Cisco IOS Release 15.3(2)T.
15.2(1)E	This command was integrated into Cisco IOS Release 15.2(1)E.

Command History

Cisco IOS XE	Modification
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.

Cisco IOS XE	Modification
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
Cisco IOS XE Release 3.1S	This command was modified. The **slow** keyword was added.
Cisco IOS XE Release 3.6S	This command was modified. Support for displaying BGP BFD multihop and C-bit information was added.
Cisco IOS XE Release 3.3SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
Cisco IOS XE Release 3.7S	This command was implemented on the Cisco ASR 903 router and the output modified. The configured discard and treat-as-withdraw attributes are displayed, along with counts of incoming Updates with a matching discard attribute or treat-as-withdraw attribute, and number of times a malformed Update is treat-as-withdraw. The capabilities of the neighbor to send and receive additional paths that are advertised or received are added.
Cisco IOS XE Release 3.8S	This command was modified. In support of the BGP Multi-Cluster ID feature, the cluster ID of a neighbor is displayed if the neighbor is assigned a cluster.

Usage Guidelines

Use the **show ip bgp neighbors** command to display BGP and TCP connection information for neighbor sessions. For BGP, this includes detailed neighbor attribute, capability, path, and prefix information. For TCP, this includes statistics related to BGP neighbor session establishment and maintenance.

Prefix activity is displayed based on the number of prefixes that are advertised and withdrawn. Policy denials display the number of routes that were advertised but then ignored based on the function or attribute that is displayed in the output.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538, for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Cisco IOS Releases 12.0(25)S, 12.4(11)T, 12.2(33)SRB, 12.2(33)SB, and Later Releases

When BGP neighbors use multiple levels of peer templates, determining which policies are applied to the neighbor can be difficult.

In Cisco IOS Release 12.0(25)S, 12.4(11)T, 12.2(33)SRB, 12.2(33)SB, and later releases, the **policy** and **detail** keywords were added to display the inherited policies and the policies configured directly on the specified neighbor. Inherited policies are policies that the neighbor inherits from a peer group or a peer policy template.

Examples

Example output is different for the various keywords available for the **show ip bgp neighbors** command. Examples using the various keywords appear in the following sections.

Examples

The following example shows output for the BGP neighbor at 10.108.50.2. This neighbor is an internal BGP (iBGP) peer. This neighbor supports the route refresh and graceful restart capabilities.

```
Device# show ip bgp neighbors 10.108.50.2

BGP neighbor is 10.108.50.2,  remote AS 1, internal link
  BGP version 4, remote router ID 192.168.252.252
  BGP state = Established, up for 00:24:25
  Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is
   60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    MPLS Label capability: advertised and received
    Graceful Restart Capability: advertised
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           113        112
    Route Refresh:          0          0
    Total:                116        115
  Default minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP additional-paths computation is enabled
  BGP advertise-best-external is enabled
  BGP table version 1, neighbor version 1/0
 Output queue size : 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               0          0
    Prefixes Total:                 0          0
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          0
    Used as multipath:            n/a          0
                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Total:                                0          0
  Number of NLRIs in the update sent: max 0, min 0
  Connections established 3; dropped 2
  Last reset 00:24:26, due to Peer closed the session
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 10.108.50.1, Local port: 179
Foreign host: 10.108.50.2, Foreign port: 42698
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x68B944):
Timer          Starts    Wakeups            Next
Retrans            27          0             0x0
TimeWait            0          0             0x0
```

```
AckHold            27         18            0x0
SendWnd             0          0            0x0
KeepAlive           0          0            0x0
GiveUp              0          0            0x0
PmtuAger            0          0            0x0
DeadWait            0          0            0x0
iss: 3915509457  snduna: 3915510016  sndnxt: 3915510016     sndwnd:  15826
irs:  233567076  rcvnxt:  233567616  rcvwnd:      15845  delrcvwnd:    539
SRTT: 292 ms, RTTO: 359 ms, RTV: 67 ms, KRTT: 0 ms
minRTT: 12 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
IP Precedence value : 6
Datagrams (max data segment is 1460 bytes):
Rcvd: 38 (out of order: 0), with data: 27, total data bytes: 539
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 08
```

The table below describes the significant fields shown in the display. Fields that are preceded by the asterisk character (*) are displayed only when the counter has a nonzero value.

Table 18: show ip bgp neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.
local AS 300 no-prepend (not shown in display)	Verifies that the local autonomous system number is not prepended to received external routes. This output supports the hiding of the local autonomous systems when a network administrator is migrating autonomous systems.
internal link	“internal link” is displayed for iBGP neighbors; “external link” is displayed for external BGP (eBGP) neighbors.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	Finite state machine (FSM) stage of session negotiation.
up for	Time, in hh:mm:ss, that the underlying TCP connection has been in existence.
Last read	Time, in hh:mm:ss, since BGP last received a message from this neighbor.
last write	Time, in hh:mm:ss, since BGP last sent a message to this neighbor.

Field	Description
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. “advertised and received” is displayed when a capability is successfully exchanged between two routers.
Route refresh	Status of the route refresh capability.
MPLS Label capability	Indicates that MPLS labels are both sent and received by the eBGP peer.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Revd	Total number of received messages.
Opens	Number of open messages sent and received.
Notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
For address family:	Address family to which the following fields refer.

Field	Description
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
1 update-group member	Number of the update-group member for this address family.
Prefix activity	Prefix statistics for this address family.
Prefixes Current	Number of prefixes accepted for this address family.
Prefixes Total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that a prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as best paths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.

Field	Description
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* Ext Community	Displays only outbound extended community policy denials.
* AS_PATH too long	Displays outbound AS_PATH length policy denials.
* AS_PATH loop	Displays outbound AS_PATH loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of autonomous system 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound nonlocal next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.
* ORIGINATOR loop	Displays outbound denials of local originated routes.
* unsuppress-map	Displays inbound denials due to an unsuppress map.
* advertise-map	Displays inbound denials due to an advertise map.
* VPN Imported prefix	Displays inbound denials of VPN prefixes.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the best path came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Deploys inbound denials because the best path came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a customer edge (CE) router.
* BGP distribute-list	Displays inbound denials due to a distribute list.

Field	Description
Number of NLRIs...	Number of network layer reachability attributes in updates.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time, in hh:mm:ss, since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be...	Indicates that the BGP time to live (TTL) security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
unread input bytes	Number of bytes of packets still to be processed.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number of times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.

Field	Description
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote neighbor.
irs:	Initial packet receive sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is higher than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT:	Shortest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Longest recorded round-trip timeout.
ACK hold:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.

Field	Description
out of order:	Number of packets received out of sequence.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
Second Congestion	Number of update packets with data sent.
Datagrams: Rcvd	Number of update packets received from a neighbor.
retransmit	Number of packets retransmitted.
fastretransmit	Number of duplicate acknowledgments retransmitted for an out of order segment before the retransmission timer expires.
partialack	Number of retransmissions for partial acknowledgments (transmissions before or without subsequent acknowledgments).
Second Congestion	Number of second retransmissions sent due to congestion.

Examples

The following partial example shows output for several external BGP neighbors in autonomous systems with 4-byte autonomous system numbers, 65536 and 65550. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router# show ip bgp neighbors

BGP neighbor is 192.168.1.2,  remote AS 65536, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 02:03:38, last write 02:03:38, hold time is 120, keepalive interval is 70
seconds
  Configured hold time is 120, keepalive interval is 70 seconds
  Minimum holdtime from neighbor is 0 seconds
.
.
.
BGP neighbor is 192.168.3.2,  remote AS 65550, external link
 Description: finance
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 02:03:48, last write 02:03:48, hold time is 120, keepalive interval is 70
seconds
  Configured hold time is 120, keepalive interval is 70 seconds
  Minimum holdtime from neighbor is 0 seconds
```

Examples

The following example displays routes advertised for only the 172.16.232.178 neighbor:

```
Device# show ip bgp neighbors 172.16.232.178 advertised-routes

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*>i10.0.0.0      172.16.232.179         0    100      0 ?
*> 10.20.2.0     10.0.0.0               0         32768 i
```

The table below describes the significant fields shown in the display.

Table 19: show ip bgp neighbors advertised-routes Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
local router ID	IP address of the local BGP speaker.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: • s—The table entry is suppressed. • d—The table entry is dampened and will not be advertised to BGP neighbors. • h—The table entry does not contain the best path based on historical information. • *—The table entry is valid. • >—The table entry is the best entry to use for that network. • i—The table entry was learned via an internal BGP (iBGP) session.

Field	Description
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: • i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • e—Entry originated from Exterior Gateway Protocol (EGP). • ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system used to forward a packet to the destination network. An entry of 0.0.0.0 indicates that there are non-BGP routes in the path to the destination network.
Metric	If shown, this is the value of the interautonomous system metric. This field is not used frequently.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

Examples

The following is sample output from the **show ip bgp neighbors** command entered with the **check-control-plane-failure** option configured:

```
Device# show ip bgp neighbors 10.10.10.1

BGP neighbor is 10.10.10.1,  remote AS 10, internal link
 Fall over configured for session
 BFD is configured. BFD peer is Up. Using BFD to detect fast fallover (single-hop) with
c-bit check-control-plane-failure.
 Inherits from template cbit-tps for session parameters
  BGP version 4, remote router ID 10.7.7.7
  BGP state = Established, up for 00:03:55
  Last read 00:00:02, last write 00:00:21, hold time is 180, keepalive interval is 60 seconds

  Neighbor sessions:
    1 active, is not multisession capable (disabled)
```

```
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
```

Examples

The following is sample output from the **show ip bgp neighbors** command entered with the **paths** keyword:

```
Device# show ip bgp neighbors 172.29.232.178 paths 10

Address    Refcount Metric Path
0x60E577B0        2     40 10 ?
```

The table below describes the significant fields shown in the display.

Table 20: show ip bgp neighbors paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	Autonomous system path for that route, followed by the origin code for that route.

Examples

The following example shows that a prefix list that filters all routes in the 10.0.0.0 network has been received from the 192.168.20.72 neighbor:

```
Device# show ip bgp neighbors 192.168.20.72 received prefix-filter

Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
   seq 5 deny 10.0.0.0/8 le 32
```

The table below describes the significant fields shown in the display.

Table 21: show ip bgp neighbors received prefix-filter Field Descriptions

Field	Description
Address family	Address family mode in which the prefix filter is received.
ip prefix-list	Prefix list sent from the specified neighbor.

Examples

The following sample output shows the policies applied to the neighbor at 192.168.1.2. The output displays both inherited policies and policies configured on the neighbor device. Inherited polices are policies that the neighbor inherits from a peer group or a peer-policy template.

```
Device# show ip bgp neighbors 192.168.1.2 policy

Neighbor: 192.168.1.2, Address-Family: IPv4 Unicast
Locally configured policies:
 route-map ROUTE in
Inherited polices:
 prefix-list NO-MARKETING in
 route-map ROUTE in
 weight 300
 maximum-prefix 10000
```

Examples

The following is sample output from the **show ip bgp neighbors** command that verifies that Bidirectional Forwarding Detection (BFD) is being used to detect fast fallover for the BGP neighbor that is a BFD peer:

```
Device# show ip bgp neighbors

BGP neighbor is 172.16.10.2,  remote AS 45000, external link
.
.
.
 Using BFD to detect fast fallover
```

Examples

The following is sample output from the **show ip bgp neighbors** command that verifies that BGP TCP path maximum transmission unit (MTU) discovery is enabled for the BGP neighbor at 172.16.1.2:

```
Device# show ip bgp neighbors 172.16.1.2

BGP neighbor is 172.16.1.2,  remote AS 45000, internal link
  BGP version 4, remote router ID 172.16.1.99
.
.
.
 For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5/0
.
.
.
  Address tracking is enabled, the RIB does have a route to 172.16.1.2
  Address tracking requires at least a /24 route to the peer
  Connections established 3; dropped 2
  Last reset 00:00:35, due to Router ID changed
  Transport(tcp) path-mtu-discovery is enabled
.
.
.
SRTT: 146 ms, RTTO: 1283 ms, RTV: 1137 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

Examples

The following is sample output from the **show ip bgp neighbors** command that verifies that the neighbor 192.168.3.2 is a member of the peer group group192 and belongs to the subnet range group 192.168.0.0/16, which shows that this BGP neighbor was dynamically created:

```
Device# show ip bgp neighbors 192.168.3.2
```

```
BGP neighbor is *192.168.3.2,  remote AS 50000, external link
 Member of peer-group group192 for session parameters
 Belongs to the subnet range group: 192.168.0.0/16
  BGP version 4, remote router ID 192.168.3.2
  BGP state = Established, up for 00:06:35
  Last read 00:00:33, last write 00:00:25, hold time is 180, keepalive intervals
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                0          0
    Keepalives:             7          7
    Route Refresh:          0          0
    Total:                  8          8
  Default minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1/0
  Output queue size : 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
  group192 peer-group member
.
.
.
```

Examples

The following is partial output from the **show ip bgp neighbors** command that verifies the status of the BGP graceful restart capability for the external BGP peer at 192.168.3.2. Graceful restart is shown as disabled for this BGP peer.

```
Device# show ip bgp neighbors 192.168.3.2

BGP neighbor is 192.168.3.2,  remote AS 50000, external link
 Inherits from template S2 for session parameters
  BGP version 4, remote router ID 192.168.3.2
  BGP state = Established, up for 00:01:41
  Last read 00:00:45, last write 00:00:45, hold time is 180, keepalive intervals
  Neighbor sessions:
    1 active, is multisession capable
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
.
.
.
Address tracking is enabled, the RIB does have a route to 192.168.3.2
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

Examples

The following is partial output from the **show ip bgp neighbors** command. For this release, the display includes the Layer 2 VFN address family information if graceful restart or NSF is enabled.

```
Device# show ip bgp neighbors

Load for five secs: 2%/0%; one minute: 0%; five minutes: 0%
Time source is hardware calendar, *21:49:17.034 GMT Wed Sep 22 2010
BGP neighbor is 10.1.1.3,  remote AS 2, internal link
```

```
 BGP version 4, remote router ID 10.1.1.3
 BGP state = Established, up for 00:14:32
Last read 00:00:30, last write 00:00:43, hold time is 180, keepalive interval is 60 seconds

 Neighbor sessions:
   1 active, is not multisession capable (disabled)
 Neighbor capabilities:
   Route refresh: advertised and received(new)
   Four-octets ASN Capability: advertised and received
   Address family IPv4 Unicast: advertised and received
   Address family L2VPN Vpls: advertised and received
   Graceful Restart Capability: advertised and received
     Remote Restart timer is 120 seconds
     Address families advertised by peer:
       IPv4 Unicast (was not preserved), L2VPN Vpls (was not preserved)
   Multisession Capability:
 Message statistics:
   InQ depth is 0
   OutQ depth is 0

                        Sent       Rcvd
   Opens:                  1          1
   Notifications:          0          0
   Updates:                4         16
   Keepalives:            16         16
   Route Refresh:          0          0
   Total:                 21         33
 Default minimum time between advertisement runs is 0 seconds
For address family: IPv4 Unicast
 Session: 10.1.1.3
 BGP table version 34, neighbor version 34/0
 Output queue size : 0
 Index 1, Advertise bit 0
 1 update-group member
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled
                              Sent       Rcvd
 Prefix activity:             ----       ----
   Prefixes Current:             2         11 (Consumes 572 bytes)
   Prefixes Total:               4         19
   Implicit Withdraw:            2          6
   Explicit Withdraw:            0          2
   Used as bestpath:           n/a          7
   Used as multipath:          n/a          0
                                Outbound    Inbound
 Local Policy Denied Prefixes:  --------    -------
   NEXT_HOP is us:                   n/a          1
   Bestpath from this peer:           20        n/a
   Bestpath from iBGP peer:            8        n/a
   Invalid Path:                      10        n/a
   Total:                             38          1
 Number of NLRIs in the update sent: max 2, min 0
 Last detected as dynamic slow peer: never
 Dynamic slow peer recovered: never
For address family: L2VPN Vpls
 Session: 10.1.1.3
 BGP table version 8, neighbor version 8/0
 Output queue size : 0
 Index 1, Advertise bit 0
 1 update-group member
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled
                              Sent       Rcvd
 Prefix activity:             ----       ----
   Prefixes Current:             1          1 (Consumes 68 bytes)
   Prefixes Total:               2          1
   Implicit Withdraw:            1          0
   Explicit Withdraw:            0          0
   Used as bestpath:           n/a          1
   Used as multipath:          n/a          0
                                Outbound    Inbound
 Local Policy Denied Prefixes:  --------    -------
   Bestpath from this peer:             4        n/a
```

```
    Bestpath from iBGP peer:              1        n/a
    Invalid Path:                         2        n/a
    Total:                                7          0
  Number of NLRIs in the update sent: max 1, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Address tracking is enabled, the RIB does have a route to 10.1.1.3
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is enabled, restart-time 120 seconds, stalepath-time 360 seconds
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Mininum incoming TTL 0, Outgoing TTL 255
Local host: 10.1.1.1, Local port: 179
Foreign host: 10.1.1.3, Foreign port: 48485
Connection tableid (VRF): 0
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0xE750C):
Timer          Starts    Wakeups            Next
Retrans            18          0             0x0
TimeWait            0          0             0x0
AckHold            22         20             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
iss: 3196633674  snduna: 3196634254  sndnxt: 3196634254     sndwnd:  15805
irs: 1633793063  rcvnxt: 1633794411  rcvwnd:      15037  delrcvwnd:   1347
SRTT: 273 ms, RTTO: 490 ms, RTV: 217 ms, KRTT: 0 ms
minRTT: 2 ms, maxRTT: 300 ms, ACK hold: 200 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable
Datagrams (max data segment is 1436 bytes):
Rcvd: 42 (out of order: 0), with data: 24, total data bytes: 1347
Sent: 40 (retransmit: 0 fastretransmit: 0),with data: 19, total data bytes: 579
```

Examples

The following is sample output from the **show ip bgp neighbors** command that indicates the discard attribute values and treat-as-withdraw attribute values configured. It also provides a count of received Updates matching a treat-as-withdraw attribute, a count of received Updates matching a discard attribute, and a count of received malformed Updates that are treat-as-withdraw.

```
Device# show ip bgp vpnv4 all neighbors 10.0.103.1

BGP neighbor is 10.0.103.1,  remote AS 100, internal link
 Path-attribute treat-as-withdraw inbound
 Path-attribute treat-as-withdraw value 128
 Path-attribute treat-as-withdraw 128 in: count 2
 Path-attribute discard 128 inbound
 Path-attribute discard 128 in: count 2

      Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    MALFORM treat as withdraw:            0          1
    Total:                                0          1
```

Examples

The following output indicates that the neighbor is capable of advertising additional paths and sending additional paths it receives. It is also capable of receiving additional paths and advertised paths.

```
Device# show ip bgp neighbors 10.108.50.2

BGP neighbor is 10.108.50.2,  remote AS 1, internal link
  BGP version 4, remote router ID 192.168.252.252
  BGP state = Established, up for 00:24:25
```

```
Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is 60 seconds

Neighbor capabilities:
  Additional paths Send: advertised and received
  Additional paths Receive: advertised and received
  Route refresh: advertised and received(old & new)
  Graceful Restart Capabilty: advertised and received
  Address family IPv4 Unicast: advertised and received
```

Examples

In the following output, the cluster ID of the neighbor is displayed. (The vertical bar and letter “i” for “include” cause the device to display only lines that include the user's input after the “i”, in this case, “cluster-id.”) The cluster ID displayed is the one directly configured through a neighbor or a template.

```
Device# show ip bgp neighbors 192.168.2.2 | i cluster-id

 Configured with the cluster-id 192.168.15.6
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
bgp enhanced-error	Restores the default behavior of treating Update messages that have a malformed attribute as withdrawn, or includes iBGP peers in the Enhanced Attribute Error Handling feature.
neighbor path-attribute discard	Configures the device to discard unwanted Update messages from the specified neighbor that contain a specified path attribute.
neighbor path-attribute treat-as-withdraw	Configures the device to withdraw from the specified neighbor unwanted Update messages that contain a specified attribute.
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
neighbor send-label explicit-null	Enables a BGP router to send MPLS labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router.
router bgp	Configures the BGP routing process.

show ip bgp vpnv4

To display VPN Version 4 (VPNv4) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in user EXEC or privileged EXEC mode.

show ip bgp vpnv4 {**all**| **rd** *route-distinguisher*| **vrf** *vrf-name*} [[*ip-prefix*/*length* [*mask*| **bestpath**| **multipaths**]| *network-address* [*mask*| **bestpath**| **longer-prefixes**| **multipaths**| **shorter-prefixes**| **subnets**]]| **cidr-only**| **cluster-ids**| **community**| **community-list**| **dampening**| **extcommunity-list** *extcommunity-list-name*| **filter-list**| **inconsistency nexthop-label**| **inconsistent-as**| **labels**| **neighbors** [{*ip-address*| *ipv6-address*} [**advertised-routes**| **dampened-routes**| **flap-statistics**| **paths**| **policy [detail]**| **received**| **received-routes**| **routes**]| **slow**]| **nexthops**| **oer-paths**| **path-attribute** {**discard**| **unknown**}| **paths** [*line*]| **peer-group**| **pending-prefixes**| **prefix-list** *prefix-list-name*| **quote-regexp**| **regexp**| **replication** [*update-group-index*] [*update-group-member-address*]| **rib-failure**| **route-map** *route-map-name*| **summary**| **update-group**| **update-source**| **version** {*version-number*| **recent** *offset-value*}]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays Network Layer Reachability Information (NLRI) prefixes that match the named route distinguisher.
vrf *vrf-name*	Displays NLRI prefixes associated with the named VPN routing and forwarding (VRF) instance.
ip-prefix/*length*	(Optional) IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter and all entries that match the prefix in a “longest-match” sense. That is, prefixes for which the specified prefix is an initial substring.
network-address	(Optional) IP address of a network in the BGP routing table.
mask	(Optional) Mask of the network address, in dotted decimal format.
cidr-only	(Optional) Displays only routes that have nonclassful netmasks.
cluster-ids	(Optional) Displays configured cluster IDs.
community	(Optional) Displays routes that match this community.

community-list	(Optional) Displays routes that match this community list.
dampening	(Optional) Displays paths suppressed because of dampening (BGP route from peer is up and down).
extcommunity-list *extended-community-list-name*	(Optional) Displays routes that match the extended community list.
filter-list	(Optional) Displays routes that conform to the filter list.
inconsistency nexthop-label	(Optional) Displays all inconsistent paths.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
labels	(Optional) Displays incoming and outgoing BGP labels for each NLRI prefix.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
ip-address	(Optional) Displays information about the neighbor at this IPv4 address.
ipv6-address	(Optional) Displays information about the neighbor at this IPv6 address.
advertised-routes	(Optional) Displays advertised routes from the specified neighbor.
dampened-routes	(Optional) Displays dampened routes from the specified neighbor.
flap-statistics	(Optional) Displays flap statistics about the specified neighbor.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
policy [detail]	(Optional) Displays configured policies for the specified neighbor.
slow	(Optional) Displays BGP slow peer information.
nexthops	(Optional) Displays nexthop address table.

oer-paths	(Optional) Displays all OER-controlled paths.
path-attribute	(Optional) Displays path-attribute-specific information.
discard	(Optional) Displays prefixes with discarded path attribute.
unknown	(Optional) Displays prefixes with unknown path attribute.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
peer-group	(Optional) Displays information about peer groups.
pending-prefixes	(Optional) Displays prefixes that are pending deletion.
prefix-list *prefix-list*	(Optional) Displays routes that match the prefix list.
quote-regexp	(Optional) Displays routes that match the autonomous system path regular expression.
regexp	(Optional) Displays routes that match the autonomous system path regular expression.
replication	(Optional) Displays replication status of update group(s).
rib-failure	(Optional) Displays BGP routes that failed to install in the VRF table.
route-map	(Optional) Displays routes that match the route map.
summary	(Optional) Displays BGP neighbor status.
update-group	(Optional) Displays information on update groups.
update-source	(Optional) Displays update source interface table.
version	(Optional) Displays prefixes with matching version numbers.
version-number	(Optional) If the **version** keyword is specified, either a *version-number* or the **recent** keyword and an *offset-value* are required.

recent *offset-value*	(Optional) Displays prefixes with matching version numbers.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	This command was modified. The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.
12.0(21)ST	This command was modified. The **tags** keyword was replaced by the **labels** keyword to conform to the MPLS guidelines.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.0(27)S	This command was modified. The output of the **show ip bgp vpnv4 all labels** command was enhanced to display explicit-null label information.
12.3	This command was modified. The **rib-failure** keyword was added for VRFs.
12.2(22)S	This command was modified. The output of the **show ip bgp vpnv4 vrf** *vrf-name* **labels** command was modified so that directly connected VRF networks no longer display as aggregate; no label appears instead.
12.2(25)S	This command was updated to display MPLS VPN nonstop forwarding information.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series router. The display output was modified to indicate whether BGP nonstop routing (NSR) with stateful switchover (SSO) is enabled and the reason the last BGP lost SSO capability.
12.2(33)SRA	This command was modified. The output was modified to support per-VRF assignment of the BGP router ID.
12.2(31)SB2	This command was modified. The output was modified to support per-VRF assignment of the BGP router ID.

Release	Modification
12.2(33)SXH	This command was modified. The output was modified to support per-VRF assignment of the BGP router ID. **Note** In Cisco IOS Release 12.2(33)SXH, the command output does not display on the standby Route Processor in NSF/SSO mode.
12.4(20)T	This command was modified. The output was modified to support per-VRF assignment of the BGP router ID.
15.0(1)M	This command was modified. The output was modified to support the BGP Event-Based VPN Import feature.
12.2(33)SRE	This command was modified. The command output was modified to support the BGP Event-Based VPN Import, BGP best external, and BGP additional path features.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.
15.0(1)SY	This command was integrated into Cisco IOS Release 15.0(1)SY.
15.2(3)T	This command was integrated into Cisco IOS Release 15.2(3)T.
15.2(4)S	This command was implemented on the Cisco 7200 series router and the output was modified to display unknown attributes and discarded attributes associated with a prefix.
Cisco IOS XE Release 3.7S	This command was implemented on the Cisco ASR 903 router and the output modified to display unknown attributes and discarded attributes associated with a prefix.
15.2(2)SNG	This command was implemented on the Cisco ASR 901 Series Aggregation Services Routers.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The **show ip bgp vpnv4 all** command displays all available VPNv4 information. The **show ip bgp vpnv4 all summary** command displays BGP neighbor status. The **show ip bgp vpnv4 all labels** command displays explicit-null label information.

Examples

The following example shows all available VPNv4 information in a BGP routing table:

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 10.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
```

```
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i10.6.6.6/32       10.0.0.21              11    100      0 ?
*> 10.7.7.7/32       10.150.0.2             11         32768 ?
*>i10.69.0.0/30      10.0.0.21               0    100      0 ?
*> 10.150.0.0/24     0.0.0.0                 0         32768 ?
```

The table below describes the significant fields shown in the display.

Table 22: show ip bgp vpnv4 all Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRI prefixes that have a route distinguisher value of 100:1.

```
Router# show ip bgp vpnv4 rd 100:1 labels

Network            Next Hop       In label/Out label
Route Distinguisher: 100:1 (vrf1)
   10.0.0.0         10.20.0.60      34/nolabel
   10.0.0.0         10.20.0.60      35/nolabel
   10.0.0.0         10.20.0.60      26/nolabel
                    10.20.0.60      26/nolabel
   10.0.0.0         10.15.0.15      nolabel/26
```

The table below describes the significant fields shown in the display.

Table 23: show ip bgp vpnv4 rd labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next-hop router.

The following example shows VPNv4 routing entries for the VRF named vpn1:

```
Router# show ip bgp vpnv4 vrf vpn1

BGP table version is 18, local router ID is 10.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf test1)
*> 10.1.1.1/32      192.168.1.1              0             0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
*> 10.2.2.2/32      192.168.1.1                           0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
*> 172.16.1.0/24    192.168.1.1              0             0 100 i
* i                 10.4.4.4                 0    100      0 100 i
r> 192.168.1.0      192.168.1.1              0             0 100 i
rbi                 10.4.4.4                 0    100      0 100 i
*> 192.168.3.0      192.168.1.1                           0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
```

The table below describes the significant fields shown in the display.

Table 24: show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows attributes for network 192.168.9.0 that include multipaths, best path, and a recursive-via-host flag:

```
Router# show ip bgp vpnv4 vrf vpn1 192.168.9.0 255.255.255.0

BGP routing table entry for 100:1:192.168.9.0/24, version 44
Paths: (2 available, best #2, table test1)
  Additional-path
  Advertised to update-groups:
     2
  100, imported path from 400:1:192.168.9.0/24
    10.8.8.8 (metric 20) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Extended Community: RT:100:1 RT:200:1 RT:300:1 RT:400:1
      Originator: 10.8.8.8, Cluster list: 10.5.5.5 , recursive-via-host
      mpls labels in/out nolabel/17
  100, imported path from 300:1:192.168.9.0/24
    10.7.7.7 (metric 20) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:1 RT:200:1 RT:300:1 RT:400:1
      Originator: 10.7.7.7, Cluster list: 10.5.5.5 , recursive-via-host
```

```
mpls labels in/out nolabel/17
```

The table below describes the significant fields shown in the display.

Table 25: show ip bgp vpnv4 all network-address Field Descriptions

Field	Description
BGP routing table entry for ... version	Internal version number of the table. This number is incremented whenever the table changes.
Paths	Number of autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath	Indicates the maximum paths configured (iBGP or eBGP).
Advertised to non peer-group peers	IP address of the BGP peers to which the specified route is advertised.
10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)	Indicates the next hop address and the address of the gateway that sent the update.
Origin	Indicates the origin of the entry. It can be one of the following values: • IGP—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • incomplete—Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the **redistribute** router configuration command. • EGP—Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is internal if the path is learned via iBGP. The field is external if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.

Field	Description
best	If multiple paths exist, one of the multipaths is designated the best path and this path is advertised to neighbors.
Extended Community	Route Target value associated with the specified route.
Originator	The router ID of the router from which the route originated when route reflector is used.
Cluster list	The router ID of all the route reflectors that the specified route has passed through.

The following example shows routes that BGP could not install in the VRF table:

```
Router# show ip bgp vpnv4 vrf xyz rib-failure

Network            Next Hop                      RIB-failure   RIB-NH Matches
Route Distinguisher: 2:2 (default for vrf bar)
10.1.1.2/32        10.100.100.100      Higher admin distance              No
10.111.111.112/32  10.9.9.9            Higher admin distance             Yes
```

The table below describes the significant fields shown in the display.

Table 26: show ip bgp vpnv4 vrf rib-failure Field Descriptions

Field	Description
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
RIB-failure	Cause of the Routing Information Base (RIB) failure. Higher admin distance means that a route with a better (lower) administrative distance, such as a static route, already exists in the IP routing table.

Field	Description
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and the **bgp suppress-inactive** command is configured for the address family being used. There are three choices: • Yes—Means that the route in the RIB has the same next hop as the BGP route or that the next hop recurses down to the same adjacency as the BGP next hop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that the **bgp suppress-inactive** command is not configured for the address family being used.

The following example shows the information displayed on the active and standby Route Processors when they are configured for NSF/SSO: MPLS VPN.

Note In Cisco IOS Release 12.2(33)SXH, the Cisco IOS Software Modularity: MPLS Layer 3 VPNs feature incurred various infrastructure changes. The result of those changes affects the output of this command on the standby Route Processor (RP). In Cisco IOS Release 12.2(33)SXH, the standby RP does not display any output from the **show ip bgp vpnv4** command.

```
Router# show ip bgp vpnv4 all labels

Network          Next Hop   In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0    16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0    17/aggregate(vpn1)
Route Distinguisher: 609:1 (vpn0)
10.13.13.13/32   0.0.0.0    18/aggregate(vpn0)

Router# show ip bgp vpnv4 vrf vpn1 labels

Network          Next Hop   In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0    16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0    17/aggregate(vpn1)


Router# show ip bgp vpnv4 all labels

Network       Masklen   In label
Route Distinguisher: 100:1
10.12.12.12   /32       16
10.0.0.0      /8        17
Route Distinguisher: 609:1
10.13.13.13   /32       18

Router# show ip bgp vpnv4 vrf vpn1 labels
```

```
Network        Masklen   In label
Route Distinguisher: 100:1
10.12.12.12    /32       16
10.0.0.0       /8        17
```

The table below describes the significant fields shown in the display.

Table 27: show ip bgp vpnv4 labels Field Descriptions

Field	Description
Network	The network address from the BGP table.
Next Hop	The BGP next-hop address.
In label	The label (if any) assigned by this router.
Out label	The label assigned by the BGP next-hop router.
Masklen	The mask length of the network address.

The following example displays output, including the explicit-null label, from the **show ip bgp vpnv4 all labels** command on a CSC-PE router:

```
Router# show ip bgp vpnv4 all labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
   10.0.0.0/24       10.0.0.0        19/aggregate(v1)
   10.0.0.1/32       10.0.0.0        20/nolabel
   10.1.1.1/32       10.0.0.0        21/aggregate(v1)
   10.10.10.10/32    10.0.0.1        25/exp-null

   10.168.100.100/32
                     10.0.0.1        23/exp-null
   10.168.101.101/32
                     10.0.0.1        22/exp-null
```

The table below describes the significant fields shown in the display.

Table 28: show ip bgp vpnv4 all labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next-hop router.
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.

The following example displays separate router IDs for each VRF in the output from an image in Cisco IOS Release 12.2(31)SB2, 12.2(33)SRA, 12.2(33)SXH, 12.4(20)T, Cisco IOS XE Release 2.1, and later releases with the Per-VRF Assignment of BGP Router ID feature configured. The router ID is shown next to the VRF name.

```
Router# show ip bgp vpnv4 all

BGP table version is 5, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf vrf_trans) VRF Router ID 10.99.1.2
*> 192.168.4.0      0.0.0.0                  0         32768 ?
Route Distinguisher: 42:1 (default for vrf vrf_user) VRF Router ID 10.99.1.1
*> 192.168.5.0      0.0.0.0                  0         32768 ?
```

The table below describes the significant fields shown in the display.

Table 29: show ip bgp vpnv4 all (VRF Router ID) Field Descriptions

Field	Description
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.
vrf	Name of the VRF.
VRF Router ID	Router ID for the VRF.

In the following example, the BGP Event-Based VPN Import feature is configured in Cisco IOS Release 15.0(1)M, 12.2(33)SRE, and later releases. When the **import path selection** command is configured, but the **strict** keyword is not included, then a safe import path selection policy is in effect. When a path is imported as the best available path (when the best path or multipaths are not eligible for import), the imported path includes the wording “imported safety path,” as shown in the output.

```
Router# show ip bgp vpnv4 all 172.17.0.0

BGP routing table entry for 45000:1:172.17.0.0/16, version 10
Paths: (1 available, best #1, table vrf-A)
Flag: 0x820
   Not advertised to any peer
   2, imported safety path from 50000:2:172.17.0.0/16
     10.0.101.1 from 10.0.101.1 (10.0.101.1)
       Origin IGP, metric 200, localpref 100, valid, internal, best
       Extended Community: RT:45000:100
```

In the following example, BGP Event-Based VPN Import feature configuration information is shown for Cisco IOS Release 15.0(1)M, 12.2(33)SRE, and later releases. When the **import path selection** command is configured with the **all** keyword, any path that matches an RD of the specified VRF will be imported, even though the path does not match the Route Targets (RT) imported by the specified VRF. In this situation, the imported path is marked as “not-in-vrf” as shown in the output. Note that on the net for vrf-A, this path is not the best path because any paths that are not in the VRFs appear less attractive than paths in the VRF.

```
Router# show ip bgp vpnv4 all 172.17.0.0
```

```
BBGP routing table entry for 45000:1:172.17.0.0/16, version 11
Paths: (2 available, best #2, table vrf-A)
Flag: 0x820
   Not advertised to any peer
   2
     10.0.101.2 from 10.0.101.2 (10.0.101.2)
       Origin IGP, metric 100, localpref 100, valid, internal, not-in-vrf
       Extended Community: RT:45000:200
       mpls labels in/out nolabel/16
   2
     10.0.101.1 from 10.0.101.1 (10.0.101.1)
       Origin IGP, metric 50, localpref 100, valid, internal, best
       Extended Community: RT:45000:100
mpls labels in/out nolabel/16
```

In the following example, the unknown attributes and discarded attributes associated with the prefix are displayed.

```
Device# show ip bgp vpnv4 all 10.0.0.0/8

BGP routing table entry for 100:200:10.0.0.0/8, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.0.103.1 from 10.0.103.1 (10.0.103.1)
      Origin IGP, localpref 100, valid, internal
      Extended Community: RT:1:100
      Connector Attribute: count=1
       type 1 len 12 value 22:22:10.0.101.22
      mpls labels in/out nolabel/16
      unknown transitive attribute: flag E0 type 129 length 32
        value 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
0000
      unknown transitive attribute: flag E0 type 140 length 32
        value 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
0000
      unknown transitive attribute: flag E0 type 120 length 32
        value 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
0000
      discarded unknown attribute: flag C0 type 128 length 32
        value 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
0000
```

The following example is based on the BGP—VPN Distinguisher Attribute feature. The output displays an Extended Community attribute, which is the VPN distinguisher (VD) of 104:1.

```
Device# show ip bgp vpnv4 unicast all 1.4.1.0/24

BGP routing table entry for 104:1:1.4.1.0/24, version 28
Paths: (1 available, best #1, no table)
   Advertised to update-groups:
      1
   Refresh Epoch 1
   1001
     19.0.101.1 from 19.0.101.1 (19.0.101.1)
       Origin IGP, localpref 100, valid, external, best
       Extended Community: VD:104:1
       mpls labels in/out nolabel/16
       rx pathid: 0, tx pathid: 0x0
```

The following example includes “allow-policy” in the output, indicating that the BGP—Support for iBGP Local-AS feature was configured for the specified neighbor by configuring the **neighbor allow-policy** command.

```
Device# show ip bgp vpnv4 all neighbors 192.168.3.3 policy
```

```
Neighbor: 192.168.3.3, Address-Family: VPNv4 Unicast
Locally configured policies:
 route-map pe33 out
 route-reflector-client
 allow-policy
 send-community both
```

Related Commands

Command	Description
import path limit	Specifies the maximum number of BGP paths, per VRF importing net, that can be imported from an exporting net.
import path selection	Specifies the BGP import path selection policy for a specific VRF instance.
neighbor allow-policy	Allows iBGP policies to be configured for the specified neighbor.
set extcommunity vpn-distinguisher	Sets a VPN distinguisher attribute to routes that pass a route map.
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip explicit-paths

To display the configured IP explicit paths, use the **show ip explicit-paths** command in user EXEC or privileged EXEC mode.

show ip explicit-paths [**name** *pathname*| **identifier** *number*] **[detail]**

Syntax Description

name *pathname*	(Optional) Displays the pathname of the explicit path.
identifier *number*	(Optional) Displays the number of the explicit path. The range is 1 to 65535.
detail	(Optional) Displays, in the long form, information about the configured IP explicit paths.

Command Default

If you enter the command without entering an optional keyword, all configured IP explicit paths are displayed.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(5)S	This command was introduced.
12.1(3)T	This command was integrated into Cisco IOS Release 12.1(3)T.
12.0(10)ST	This command was integrated into Cisco IOS Release 12.0(10)ST.
12.2(28)SB	The command output was enhanced to display SLRG-related information.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.
Cisco IOS XE Release 2.3	This command was integrated into Cisco IOS XE Release 2.3.

Usage Guidelines

An IP explicit path is a list of IP addresses, each representing a node or link in the explicit path.

Examples

The following is sample output from the **show ip explicit-paths** command:

```
Router# show ip explicit-paths
PATH 200 (strict source route, path complete, generation 6)
    1: next-address 10.3.28.3
    2: next-address 10.3.27.3
```

The table below describes the significant fields shown in the display.

Table 30: show ip explicit-paths Field Descriptions

Field	Description
PATH	Pathname or number, followed by the path status.
1: next-address	First IP address in the path.
2: next-address	Second IP address in the path.

Related Commands

Command	Description
append-after	Inserts a path entry after a specific index number.
index	Inserts or modifies a path entry at a specific index.
ip explicit-path	Enters the subcommand mode for IP explicit paths so that you can create or modify the named path.
list	Displays all or part of the explicit paths.
next-address	Specifies the next IP address in the explicit path.

show ip multicast mpls vif

To display the virtual interfaces (VIFs) that are created on the Multiprotocol Label Switching (MPLS) traffic engineering (TE) point-to-multipoint (P2MP) tailend router, use the **show ip multicast mpls vif** command in privileged EXEC mode.

show ip multicast mpls vif

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SRE	This command was introduced.

Examples

The following example shows information about the virtual interfaces:

```
Router# show ip multicast mpls vif
Interface   Next-hop             Application     Ref-Count   Table / VRF name
 Lspvif0     10.1.0.1             Traffic-eng      1           default
 Lspvif4     10.2.0.1             Traffic-eng      1           default
```

The table below describes the significant fields shown in the display.

Table 31: show ip multicast mpls vif Field Descriptions

Field	Description
Interface	The name of the virtual interface
Next-hop	For P2MP TE, the source address of the TE P2MP tunnel. Only one label switched path (LSP) VIF is created for all TE P2MP tunnels that have the same source address.
Application	The name of the multicast application that creates the VIF.
Table/VRF name	The multicast virtual routing and forwarding (VRF) table used.

Related Commands

Command	Description
show ip mroute	Displays IP multicast traffic.

show ip ospf database opaque-area

To display lists of information related to traffic engineering opaque link-state advertisements (LSAs), also known as Type-10 opaque link area link states, use the **show ip ospf database opaque-area** command in user EXEC or privileged EXEC mode.

show ip ospf database opaque-area

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.0(8)S	This command was introduced.
12.1(3)T	This command was integrated into Cisco IOS Release 12.1(3)T.
12.0(10)ST	This command was integrated into Cisco IOS Release 12.0(10)ST.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

The following is sample output from the **show ip ospf database opaque-area** command:

```
Router# show ip ospf database opaque-area
OSPF Router with ID (10.3.3.3) (Process ID 1)

                Type-10 Opaque Link Area Link States (Area 0)

  LS age: 12
  Options: (No TOS-capability, DC)
  LS Type: Opaque Area Link
  Link State ID: 10.0.0.0
  Opaque Type: 1
  Opaque ID: 0
  Advertising Router: 172.16.8.8
  LS Seq Number: 80000004
  Checksum: 0xD423
  Length: 132
  Fragment number : 0

    MPLS TE router ID: 172.16.8.8

    Link connected to Point-to-Point network
```

```
    Link ID : 10.2.2.2
    Interface Address : 192.168.1.1
```

The table below describes the significant fields shown in the display.

Table 32: show ip ospf database opaque-area Field Descriptions

Field	Description
LS age	Link-state age.
Options	Type of service options.
LS Type	Type of the link state.
Link State ID	Router ID number.
Opaque Type	Opaque link-state type.
Opaque ID	Opaque LSA ID number.
Advertising Router	Advertising router ID.
LS Seq Number	Link-state sequence number that detects old or duplicate link state advertisements (LSAs).
Checksum	Fletcher checksum of the complete contents of the LSA.
Length	Length (in bytes) of the LSA.
Fragment number	Arbitrary value used to maintain multiple traffic engineering LSAs.
MPLS TE router ID	Unique MPLS traffic engineering ID.
Link ID	Index of the link being described.
Interface Address	Address of the interface.

Related Commands

Command	Description
mpls traffic-eng area	Configures a router running OSPF MPLS to flood traffic engineering for an indicated OSPF area.
mpls traffic-eng router-id	Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.

Command	Description
show ip ospf mpls traffic-eng	Provides information about the links available on the local router for traffic engineering.

show ip ospf mpls ldp interface

To display information about interfaces belonging to an Open Shortest Path First (OSPF) process that is configured for Multiprotocol Label Switching (MPLS) Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP), use the **show ip ospf mpls ldp interface** command in privileged EXEC mode.

show ip ospf [*process-id*] **mpls ldp interface** [*interface*]

Syntax Description

process-id	(Optional) Process ID. Includes information only for the specified routing process.
interface	(Optional) Defines the interface for which MPLS LDP-IGP synchronization information is displayed.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(30)S	This command was introduced.
12.3(14)T	This command was integrated into Cisco IOS Release 12.3(14)T.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SB	This command was integrated into Cisco IOS Release 12.2(33)SB.
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.
Cisco IOS XE Release 3.6S	This command was implemented on the Cisco ASR 903 series routers.

Usage Guidelines

This command shows MPLS LDP-IGP synchronization information for specified interfaces or OSPF processes. If you do not specify an argument, information is displayed for each interface that was configured for MPLS LDP-IGP synchronization.

Examples

The following is sample output from the **show ip ospf mpls ldp interface** command:

```
Router# show ip ospf mpls ldp interface
Serial1/2.4
  Process ID 2, Area 0
  LDP is configured through LDP autoconfig
  LDP-IGP Synchronization : Not required
  Holddown timer is disabled
  Interface is up
Serial1/2.11
```

```
  Process ID 6, VRF VFR1, Area 2
  LDP is configured through LDP autoconfig
  LDP-IGP Synchronization : Not required
  Holddown timer is disabled
  Interface is up
Ethernet2/0
  Process ID 1, Area 0
  LDP is configured through LDP autoconfig
  LDP-IGP Synchronization : Required
  Holddown timer is configured : 1 msecs
  Holddown timer is not running
  Interface is up
Loopback1
  Process ID 1, Area 0
  LDP is not configured through LDP autoconfig
  LDP-IGP Synchronization : Not required
  Holddown timer is disabled
  Interface is up
Serial1/2.1
  Process ID 1, Area 10.0.1.44
  LDP is configured through LDP autoconfig
  LDP-IGP Synchronization : Required
  Holddown timer is configured : 1 msecs
  Holddown timer is not running
  Interface is up
```

The table below describes the significant fields shown in the display.

Table 33: show ip ospf mpls ldp interface Field Descriptions

Field	Description
Process ID	The number of the OSPF process to which the interface belongs.
Area	The OSPF area to which the interface belongs.
LDP is configured through	The means by which LDP was configured on the interface. LDP can be configured on the interface by the**mpls ip** or **mpls ldp** command.
LDP-IGP Synchronization	Indicates whether MPLS LDP-IGP synchronization was enabled on this interface.
Holddown timer	Indicates whether the hold-down timer was specified for this interface.

Related Commands

Command	Description
debug mpls ldp igp sync	Displays events related to MPLS LDP-IGP synchronization.
show mpls ldp igp sync	Displays the status of the MPLS LDP-IGP synchronization process.

show ip ospf mpls traffic-eng

To display information about the links available on the local router for traffic engineering, use the **show ip ospf mpls traffic-eng** command in user EXEC or privileged EXEC mode.

show ip ospf [**process-id [area-id] mpls traffic-eng [link]| fragment**]

Syntax Description

process-id	(Optional) Internal identification number that is assigned locally when the OSPF routing process is enabled. The value can be any positive integer.
area-id	(Optional) Area number associated with OSPF.
link	(Optional) Provides detailed information about the links over which traffic engineering is supported on the local router.
fragment	(Optional) Provides detailed information about the traffic engineering fragments on the local router.

Command Default

No default behavior or values.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.0S	This command was introduced.
12.1(3)T	This command was integrated into Cisco IOS Release 12.1(3)T.
12.0(10)ST	This command was integrated into Cisco IOS Release 12.0(10)ST.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

The following is sample output from the **show ip ospf mpls traffic-eng** command:

```
Router# show ip ospf mpls traffic-eng link
OSPF Router with ID (10.0.0.1) (Process ID 1)

  Area 0 has 2 MPLS TE links. Area instance is 14.

  Links in hash bucket 8.
    Link is associated with fragment 1. Link instance is 14
      Link connected to Point-to-Point network
      Link ID :197.0.0.1
      Interface Address :172.16.0.1
      Neighbor Address :172.16.0.2
      Admin Metric :97
      Maximum bandwidth :128000
      Maximum reservable bandwidth :250000
      Number of Priority :8
      Priority 0 :250000      Priority 1 :250000
      Priority 2 :250000      Priority 3 :250000
      Priority 4 :250000      Priority 5 :250000
      Priority 6 :250000      Priority 7 :212500
      Affinity Bit :0x0
    Link is associated with fragment 0. Link instance is 14
      Link connected to Broadcast network
      Link ID :192.168.1.2
      Interface Address :192.168.1.1
      Neighbor Address :192.168.1.2
      Admin Metric :10
      Maximum bandwidth :1250000
      Maximum reservable bandwidth :2500000
      Number of Priority :8
      Priority 0 :2500000     Priority 1 :2500000
      Priority 2 :2500000     Priority 3 :2500000
      Priority 4 :2500000     Priority 5 :2500000
      Priority 6 :2500000     Priority 7 :2500000
      Affinity Bit :0x0
```

The table below describes the significant fields shown in the display.

Table 34: show ip ospf mpls traffic-eng Field Descriptions

Field	Description
OSPF Router with ID	Router identification number.
Process ID	OSPF process identification.
Area instance	Number of times traffic engineering information or any link changed.
Link instance	Number of times any link changed.
Link ID	Link-state ID.
Interface Address	Local IP address on the link.
Neighbor Address	IP address that is on the remote end of the link.
Admin Metric	Traffic engineering link metric.

Field	Description
Maximum bandwidth	Bandwidth set by the **bandwidth interface** command in the interface configuration mode.
Maximum reservable bandwidth	Bandwidth available for traffic engineering on this link. This value is set in the **ip rsvp** command in the interface configuration mode.
Number of priority	Number of priorities that are supported.
Priority	Bandwidth (in bytes per second) that is available for traffic engineering at certain priorities.
Affinity Bit	Affinity bits (color) assigned to the link.

show ip protocols vrf

To display the routing protocol information associated with a Virtual Private Network (VPN) routing and forwarding (VRF) instance, use the **show ip protocols vrf** command in user EXEC or privileged EXEC mode.

show ip protocols vrf vrf-name [summary]

Syntax Description

vrf-name	Name assigned to a VRF.
summary	Optional. Displays the routing protocol information in summary format.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.0(21)ST	This command was integrated into Cisco IOS Release 12.0(21)ST.
12.0(22)S	The **summary** keyword was added. EIGRP VRF support was added.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use this command to display routing information associated with a VRF.

Examples

The following example shows information about a VRF named vpn1:

```
Router# show ip protocols vrf vpn1
Routing Protocol is "bgp 100"
  Sending updates every 60 seconds, next due in 0 sec
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
```

```
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Redistributing:connected, static
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.13.13.13          200      02:20:54
    10.18.18.18          200      03:26:15
  Distance:external 20 internal 200 local 200
```

The table below describes the significant fields shown in the display.

Table 35: show ip protocols vrf Field Descriptions

Field	Description
Gateway	Displays the IP address of the router identifier for all routers in the network.
Distance	Displays the metric used to access the destination route.
Last Update	Displays the last time the routing table was updated from the source.

Related Commands

Command	Description
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip route

To display contents of the routing table, use the **show ip route** command in user EXEC or privileged EXEC mode.

show ip route [*ip-address* [**repair-paths**| **next-hop-override [dhcp]**| *mask* **[longer-prefixes]**]| *protocol* [*process-id*]| **list** [*access-list-number* | *access-list-name*]| **static download**| **update-queue**]

Syntax Description

ip-address	(Optional) IP address for which routing information should be displayed.
repair-paths	(Optional) Displays the repair paths.
next-hop-override	(Optional) Displays the Next Hop Resolution Protocol (NHRP) next-hop overrides that are associated with a particular route and the corresponding default next hops.
dhcp	(Optional) Displays routes added by the Dynamic Host Configuration Protocol (DHCP) server.
mask	(Optional) Subnet mask.
longer-prefixes	(Optional) Displays output for longer prefix entries.
protocol	(Optional) The name of a routing protocol or the keyword **connected**, **mobile**, **static**, or **summary**. If you specify a routing protocol, use one of the following keywords: **bgp**, **eigrp**, **hello**, **isis**, **odr**, **ospf**, **nhrp**, or **rip**.
process-id	(Optional) Number used to identify a process of the specified protocol.
list	(Optional) Filters output by an access list name or number.
access-list-number	(Optional) Access list number.
access-list-name	(Optional) Access list name.
static	(Optional) Displays static routes.
download	(Optional) Displays routes installed using the authentication, authorization, and accounting (AAA) route download function. This keyword is used only when AAA is configured.

update-queue	(Optional) Displays Routing Information Base (RIB) queue updates.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
9.2	This command was introduced.
10.0	This command was modified. The "D—EIGRP, EX—EIGRP, N1—SPF NSSA external type 1 route" and "N2—OSPF NSSA external type 2 route" codes were included in the command output.
10.3	This command was modified. The *process-id* argument was added.
11.0	This command was modified. The **longer-prefixes** keyword was added.
11.1	This command was modified. The "U—per-user static route" code was included in the command output.
11.2	This command was modified. The "o—on-demand routing" code was included in the command output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the **update-queue** keyword was added.
11.3	This command was modified. The command output was enhanced to display the origin of an IP route in Intermediate System-to-Intermediate System (IS-IS) networks.
12.0(1)T	This command was modified. The "M—mobile" code was included in the command output.
12.0(3)T	This command was modified. The "P—periodic downloaded static route" code was included in the command output.
12.0(4)T	This command was modified. The "ia—IS-IS" code was included in the command output.
12.2(2)T	This command was modified. The command output was enhanced to display information on multipaths to the specified network.

Release	Modification
12.2(13)T	This command was modified. The *egp* and *igrp* arguments were removed because the Exterior Gateway Protocol (EGP) and the Interior Gateway Routing Protocol (IGRP) were no longer available in Cisco software.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.3(2)T	This command was modified. The command output was enhanced to display route tag information.
12.3(8)T	This command was modified. The command output was enhanced to display static routes using DHCP.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRE	This command was modified. The **dhcp** and **repair-paths** keywords were added.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5. The **next-hop-override** and **nhrp** keywords were added.
15.2(2)S	This command was modified. The command output was enhanced to display route tag values in dotted decimal format.
Cisco IOS XE Release 3.6S	This command was modified. The command output was enhanced to display route tag values in dotted decimal format.
15.2(4)S	This command was implemented on the Cisco 7200 series router.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Examples

Examples

The following is sample output from the **show ip route** command when an IP address is not specified:

```
Device# show ip route

Codes: R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
```

```
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route
Gateway of last resort is 10.119.254.240 to network 10.140.0.0
O E2 10.110.0.0 [160/5] via 10.119.254.6, 0:01:00, Ethernet2
E    10.67.10.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
O E2 10.68.132.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
O E2 10.130.0.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
E    10.128.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.129.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.65.129.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.10.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.75.139.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.16.208.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.84.148.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.31.223.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.44.236.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.141.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.140.0.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
```

The following sample output from the **show ip route** command includes routes learned from IS-IS Level 2:

```
Device# show ip route

Codes: R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route
Gateway of last resort is not set
     10.89.0.0 is subnetted (mask is 255.255.255.0), 3 subnets
C       10.89.64.0 255.255.255.0 is possibly down,
          routing via 10.0.0.0, Ethernet0
i L2    10.89.67.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0
i L2    10.89.66.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0
```

The following is sample output from the **show ip route** *ip-address mask* **longer-prefixes** command. When this keyword is included, the address-mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed. The logical AND operation is performed on the source address 10.0.0.0 and the mask 10.0.0.0, resulting in 10.0.0.0. Each destination in the routing table is also logically ANDed with the mask and compared with 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Device# show ip route 10.0.0.0 10.0.0.0 longer-prefixes

Codes: R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is not set

S    10.134.0.0 is directly connected, Ethernet0
S    10.10.0.0 is directly connected, Ethernet0
S    10.129.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
S    10.49.246.0 is directly connected, Ethernet0
S    10.160.97.0 is directly connected, Ethernet0
S    10.153.88.0 is directly connected, Ethernet0
S    10.76.141.0 is directly connected, Ethernet0
S    10.75.138.0 is directly connected, Ethernet0
S    10.44.237.0 is directly connected, Ethernet0
S    10.31.222.0 is directly connected, Ethernet0
S    10.16.209.0 is directly connected, Ethernet0
S    10.145.0.0 is directly connected, Ethernet0
```

```
S    10.141.0.0 is directly connected, Ethernet0
S    10.138.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
     10.19.0.0 255.255.255.0 is subnetted, 1 subnets
C       10.19.64.0 is directly connected, Ethernet0
     10.69.0.0 is variably subnetted, 2 subnets, 2 masks
C       10.69.232.32 255.255.255.240 is directly connected, Ethernet0
S       10.69.0.0 255.255.0.0 is directly connected, Ethernet0
```

The following sample outputs from the **show ip route** command display all downloaded static routes. A “p” indicates that these routes were installed using the AAA route download function.

```
Device# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR, P - periodic downloaded static route
       T - traffic engineered route

Gateway of last resort is 172.16.17.1 to network 10.0.0.0

        172.31.0.0/32 is subnetted, 1 subnets
P       172.31.229.41 is directly connected, Dialer1 10.0.0.0/8 is subnetted, 3 subnets
P       10.1.1.0 [200/0] via 172.31.229.41, Dialer1
P       10.1.3.0 [200/0] via 172.31.229.41, Dialer1
P       10.1.2.0 [200/0] via 172.31.229.41, Dialer1

Device# show ip route static

     172.16.4.0/8 is variably subnetted, 2 subnets, 2 masks
P       172.16.1.1/32 is directly connected, BRI0
P       172.16.4.0/8 [1/0] via 10.1.1.1, BRI0
S    172.31.0.0/16 [1/0] via 172.16.114.65, Ethernet0
S    10.0.0.0/8 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
     172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S       172.16.114.201/32 is directly connected, BRI0
S       172.16.114.205/32 is directly connected, BRI0
S       172.16.114.174/32 is directly connected, BRI0
S       172.16.114.12/32 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
P    10.1.0.0/16 is directly connected, BRI0
P    10.2.2.0/24 is directly connected, BRI0
S*   0.0.0.0/0 [1/0] via 172.16.114.65, Ethernet0
S    172.16.0.0/16 [1/0] via 172.16.114.65, Ethernet0
```

The following sample output from the **show ip route static download** command displays all active and inactive routes installed using the AAA route download function:

```
Device# show ip route static download

Connectivity: A - Active, I - Inactive

A     10.10.0.0 255.0.0.0 BRI0
A     10.11.0.0 255.0.0.0 BRI0
A     10.12.0.0 255.0.0.0 BRI0
A     10.13.0.0 255.0.0.0 BRI0
I     10.20.0.0 255.0.0.0 172.21.1.1
I     10.22.0.0 255.0.0.0 Serial0
I     10.30.0.0 255.0.0.0 Serial0
I     10.31.0.0 255.0.0.0 Serial1
I     10.32.0.0 255.0.0.0 Serial1
A     10.34.0.0 255.0.0.0 192.168.1.1
A     10.36.1.1 255.255.255.255 BRI0 200 name remote1
I     10.38.1.9 255.255.255.0 192.168.69.1
```

The following sample outputs from the **show ip route nhrp** command display shortcut switching on the tunnel interface:

```
Device# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
Gateway of last resort is not set
10.0.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       10.1.1.0/24 is directly connected, Tunnel0
C       172.16.22.0 is directly connected, Ethernet1/0
H       172.16.99.0 [250/1] via 10.1.1.99, 00:11:43, Tunnel0
     10.11.0.0/24 is subnetted, 1 subnets
C       10.11.11.0 is directly connected, Ethernet0/0

Device# show ip route nhrp

H       172.16.99.0 [250/1] via 10.1.1.99, 00:11:43, Tunnel0
```

The following are sample outputs from the **show ip route** command when the **next-hop-override** keyword is used. When this keyword is included, the NHRP next-hop overrides that are associated with a particular route and the corresponding default next hops are displayed.

```
================================================================
1) Initial configuration
================================================================

Device# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
```

```
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip cef

Prefix               Next Hop             Interface
.
.
.
10.2.1.255/32         receive              Loopback1
10.10.10.0/24        attached             Tunnel0  <<<<<<<<
10.11.11.0/24        attached             Ethernet0/0
172.16.0.0/12          drop
.
.
.
============================================================
2) Add a next-hop override
   address = 10.10.10.0
   mask = 255.255.255.0
   gateway = 10.1.1.1
   interface = Tunnel0
============================================================

Device# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets

S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets

S        10.10.10.0 is directly connected, Tunnel0
                  [NHO][1/0] via 10.1.1.1, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip cef

Prefix               Next Hop             Interface
.
.
.
10.2.1.255/32         receive              Loopback110.10.10.0/24
```

```
10.10.10.0/24       10.1.1.1              Tunnel0

10.11.11.0/24       attached            Ethernet0/0
10.12.0.0/16 drop
.
.
.
==============================================================
3) Delete a next-hop override
   address = 10.10.10.0
   mask = 255.255.255.0
   gateway = 10.11.1.1
   interface = Tunnel0
==============================================================

Device# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set
      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Device# show ip cef

Prefix              Next Hop            Interface
.
.
.
10.2.1.255/32         receive              Loopback110.10.10.0/24

10.10.10.0/24       attached            Tunnel0
10.11.11.0/24       attached            Ethernet0/0
10.120.0.0/16 drop
.
.
.
```

The table below describes the significant fields shown in the displays:

Table 36: show ip route Field Descriptions

Field	Description
Codes (Protocol)	Indicates the protocol that derived the route. It can be one of the following values: • B—BGP derived • C—Connected • D—Enhanced Interior Gateway Routing Protocol (EIGRP) • EX—EIGRP external • H—NHRP • i—IS-IS derived • ia—IS-IS • L—Local • M—Mobile • o—On-demand routing • O—Open Shortest Path First (OSPF) derived • P—Periodic downloaded static route • R—Routing Information Protocol (RIP) derived • S—Static • U—Per-user static route • +—Replicated route
Codes (Type)	Type of route. It can be one of the following values: • *—Indicates the last path used when a packet was forwarded. This information is specific to nonfast-switched packets. • E1—OSPF external type 1 route • E2—OSPF external type 2 route • IA—OSPF interarea route • L1—IS-IS Level 1 route • L2—IS-IS Level 2 route • N1—OSPF not-so-stubby area (NSSA) external type 1 route • N2—OSPF NSSA external type 2 route

Field	Description
10.110.0.0	Indicates the address of the remote network.
[160/5]	The first number in brackets is the administrative distance of the information source; the second number is the metric for the route.
via 10.119.254.6	Specifies the address of the next device to the remote network.
0:01:00	Specifies the last time the route was updated (in hours:minutes:seconds).
Ethernet2	Specifies the interface through which the specified network can be reached.

Examples

The following is sample output from the **show ip route** command when an IP address is specified:

```
Device# show ip route 10.0.0.1

Routing entry for 10.0.0.1/32
    Known via "isis", distance 115, metric 20, type level-1
    Redistributing via isis
    Last update from 10.191.255.251 on Fddi1/0, 00:00:13 ago
    Routing Descriptor Blocks:
    * 10.22.22.2, from 10.191.255.247, via Serial2/3
       Route metric is 20, traffic share count is 1
       10.191.255.251, from 10.191.255.247, via Fddi1/0
       Route metric is 20, traffic share count is 1
```

When an IS-IS router advertises its link-state information, the router includes one of its IP addresses to be used as the originator IP address. When other routers calculate IP routes, they store the originator IP address with each route in the routing table.

The preceding example shows the output from the **show ip route** command for an IP route generated by IS-IS. Each path that is shown under the Routing Descriptor Blocks report displays two IP addresses. The first address (10.22.22.2) is the next-hop address. The second is the originator IP address from the advertising IS-IS router. This address helps you determine the origin of a particular IP route in your network. In the preceding example, the route to 10.0.0.1/32 was originated by a device with IP address 10.191.255.247.

The table below describes the significant fields shown in the display.

Table 37: show ip route with IP Address Field Descriptions

Field	Description
Routing entry for 10.0.0.1/32	Network number and mask.
Known via...	Indicates how the route was derived.
Redistributing via...	Indicates the redistribution protocol.

Field	Description
Last update from 10.191.255.251	Indicates the IP address of the router that is the next hop to the remote network and the interface on which the last update arrived.
Routing Descriptor Blocks	Displays the next-hop IP address followed by the information source.
Route metric	This value is the best metric for this Routing Descriptor Block.
traffic share count	Indicates the number of packets transmitted over various routes.

The following sample output from the **show ip route** command displays the tag applied to the route 10.22.0.0/16. You must specify an IP prefix to see the tag value. The fields in the display are self-explanatory.

```
Device# show ip route 10.22.0.0

Routing entry for 10.22.0.0/16
  Known via "isis", distance 115, metric 12
  Tag 120, type level-1
  Redistributing via isis
  Last update from 172.19.170.12 on Ethernet2, 01:29:13 ago
  Routing Descriptor Blocks:
    * 172.19.170.12, from 10.3.3.3, via Ethernet2
        Route metric is 12, traffic share count is 1
        Route tag 120
```

Examples

The following example shows that IP route 10.8.8.0 is directly connected to the Internet and is the next-hop (option 3) default gateway. Routes 10.1.1.1 [1/0], 10.3.2.1 [24/0], and 172.16.2.2 [1/0] are static, and route 10.0.0.0/0 is a default route candidate. The fields in the display are self-explanatory.

```
Device# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
Gateway of last resort is 10.0.19.14 to network 0.0.0.0
10.0.0.0/24 is subnetted, 1 subnets
C 10.8.8.0 is directly connected, Ethernet1
  10.0.0.0/32 is subnetted, 1 subnets
S 10.1.1.1 [1/0] via 10.8.8.1
  10.0.0.0/32 is subnetted, 1 subnets
S 10.3.2.1 [24/0] via 10.8.8.1
  172.16.0.0/32 is subnetted, 1 subnets
S 172.16.2.2 [1/0] via 10.8.8.1
  10.0.0.0/28 is subnetted, 1 subnets
C 10.0.19.0 is directly connected, Ethernet0
  10.0.0.0/24 is subnetted, 1 subnets
C 10.15.15.0 is directly connected, Loopback0
S* 10.0.0.0/0 [1/0] via 10.0.19.14
```

The following sample output from the **show ip route repair-paths** command shows repair paths marked with the tag [RPR]. The fields in the display are self-explanatory:

```
Device# show ip route repair-paths

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
C        10.1.1.1 is directly connected, Loopback0
B        10.2.2.2 [200/0] via 172.16.1.2, 00:31:07
                  [RPR][200/0] via 192.168.1.2, 00:31:07
B        10.9.9.9 [20/0] via 192.168.1.2, 00:29:45
                  [RPR][20/0] via 192.168.3.2, 00:29:45
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.16.1.0/24 is directly connected, Ethernet0/0
L        172.16.1.1/32 is directly connected, Ethernet0/0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial2/0
L        192.168.1.1/32 is directly connected, Serial2/0
B     192.168.3.0/24 [200/0] via 172.16.1.2, 00:31:07
                     [RPR][200/0] via 192.168.1.2, 00:31:07
B     192.168.9.0/24 [20/0] via 192.168.1.2, 00:29:45
                     [RPR][20/0] via 192.168.3.2, 00:29:45
B     192.168.13.0/24 [20/0] via 192.168.1.2, 00:29:45
                      [RPR][20/0] via 192.168.3.2, 00:29:45

Device# show ip route repair-paths 10.9.9.9

>Routing entry for 10.9.9.9/32
>  Known via "bgp 100", distance 20, metric 0
>  Tag 10, type external
>  Last update from 192.168.1.2 00:44:52 ago
>  Routing Descriptor Blocks:
>  * 192.168.1.2, from 192.168.1.2, 00:44:52 ago, recursive-via-conn
>      Route metric is 0, traffic share count is 1
>      AS Hops 2
>      Route tag 10
>      MPLS label: none
>    [RPR]192.168.3.2, from 172.16.1.2, 00:44:52 ago
>      Route metric is 0, traffic share count is 1
>      AS Hops 2
>      Route tag 10
>      MPLS label: none
```

Related Commands

Command	Description
show interfaces tunnel	Displays tunnel interface information.
show ip route summary	Displays the current state of the routing table in summary format.

show ip route vrf

To display the IP routing table associated with a specific VPN routing and forwarding (VRF) instance, use the **show ip route vrf** command in user EXEC or privileged EXEC mode.

show ip route vrf *vrf-name* [**connected**| *protocol* [*as-number*]| **list** [*list-number*]| **profile**| **static**| **summary**| [*ip-prefix*| *ip-address*] [*mask*| **longer-prefixes**]| **repair-paths**| **dhcp**| **supernets-only**| **tag** {*tag-value*| *tag-value-dotted-decimal* [*mask*]}]

Syntax Description

vrf-name	Name of the VRF.
connected	(Optional) Displays all connected routes in a VRF.
protocol	(Optional) Routing protocol. To specify a routing protocol, use one of the following keywords: **bgp**, **egp**, **eigrp**, **hello**, **igrp**, **isis**, **ospf**, or **rip**.
as-number	(Optional) Autonomous system number.
list *number*	(Optional) Specifies the IP access list to be displayed.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
ip-prefix	(Optional) Network for which routing information is displayed.
ip-address	(Optional) Address for which routing information is displayed.
mask	(Optional) Network mask.
longer-prefixes	(Optional) Displays longer prefix entries.
repair-paths	(Optional) Displays repair paths.
dhcp	(Optional) Displays routes added by the DHCP server.
supernets-only	(Optional) Displays only supernet entries.
tag	(Optional) Displays information about route tags in the VRF table.
tag-value	(Optional) Route tag values as a plain decimals.

tag-value-dotted-decimal	(Optional) Route tag values as a dotted decimals.
mask	(Optional) Route tag wildcard mask.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	This command was modified. The *ip-prefix* argument was added. The command output was enhanced to display information on multipaths to the specified network.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was modified. Support for Enhanced Interior Gateway Routing Protocol (EIGRP) VRFs was added.
12.2(15)T	This command was modified. Support for EIGRP VRFs was added.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH. The output was enhanced to display remote label information and corresponding Multiprotocol Label Switching (MPLS) flags for prefixes that have remote labels stored in the Routing Information Base (RIB).
12.2(33)SRE	This command was modified. The **repair-paths**, **dhcp**, and **supernets-only** keywords were added. Support for the Border Gateway Protocol (BGP) Best External and BGP Additional Path features was added.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.
15.2(2)S	This command was modified. The **tag** keyword and *tag-value*, *tag-value-dotted-decimal*, and *mask* arguments were added to enable the display of route tags as plain or dotted decimals in the command output.
Cisco IOS XE Release 3.6S	This command was modified. The **tag** keyword and *tag-value*, *tag-value-dotted-decimal*, and *mask* arguments were added to enable the display of route tags as plain or dotted decimals in the command output.
15.2(4)S	This command was implemented on the Cisco 7200 series router.

Release	Modification
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Examples

The following sample output displays the IP routing table associated with the VRF named vrf1:

```
Device# show ip route vrf vrf1

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       I - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR
       T - traffic engineered route

Gateway of last resort is not set

B    10.0.0.0/8 [200/0] via 10.13.13.13, 00:24:19
C    10.0.0.0/8 is directly connected, Ethernet1/3
B    10.0.0.0/8 [20/0] via 10.0.0.1, 02:10:22
B    10.0.0.0/8 [200/0] via 10.13.13.13, 00:24:20
```

This following sample output shows BGP entries in the IP routing table associated with the VRF named vrf1:

```
Device# show ip route vrf vrf1 bgp

B  10.0.0.0/8 [200/0] via 10.13.13.13, 03:44:14
B  10.0.0.0/8 [20/0] via 10.0.0.1, 03:44:12
B  10.0.0.0/8 [200/0] via 10.13.13.13, 03:43:14
```

The following sample output displays the IP routing table associated with a VRF named PATH:

```
Device# show ip route vrf PATH 10.22.22.0

Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 200, metric 0
  Tag 22, type internal
  Last update from 10.22.5.10 00:01:07 ago
  Routing Descriptor Blocks:
  * 10.22.7.8 (Default-IP-Routing-Table), from 10.11.3.4, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.1.9 (Default-IP-Routing-Table), from 10.11.1.2, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.6.10 (Default-IP-Routing-Table), from 10.11.6.7, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.4.10 (Default-IP-Routing-Table), from 10.11.4.5, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.5.10 (Default-IP-Routing-Table), from 10.11.5.6, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
```

The following sample output from the **show ip route vrf** *vrf-name* **tag** command displays route tag information for routes associated with vrf1. The route tags in the sample output are displayed in dotted decimal format.

```
Device# show ip route vrf vrf1 tag 5

Routing Table: vrf1
Routing entry for 10.0.0.1/24
  Known via "static", distance 1, metric 0 (connected)
  Tag 0.0.0.5
  Routing Descriptor Blocks:
```

```
  * directly connected, via Null0
      Route metric is 0, traffic share count is 1
      Route tag 0.0.0.5
```

The following sample outputs from the **show ip route vrf** command include recursive-via-host and recursive-via-connected flags:

```
Device# show ip route vrf v2 10.2.2.2

Routing Table: v2
Routing entry for 10.2.2.2/32
  Known via "bgp 10", distance 20, metric 0
  Tag 100, type external
  Last update from 192.168.1.1 00:15:54 ago
  Routing Descriptor Blocks:
  * 192.168.1.1, from 192.168.1.1, 00:15:54 ago, recursive-via-conn
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none

Device# show ip route vrf v2 10.2.2.2

Routing Table: v2
Routing entry for 10.2.2.2/32
  Known via "bgp 10", distance 200, metric 0
  Tag 100, type internal
  Last update from 10.3.3.3 00:18:11 ago
  Routing Descriptor Blocks:
  * 10.3.3.3 (default), from 10.5.5.5, 00:18:11 ago, recursive-via-host
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: 16
      MPLS Flags: MPLS Required
```

The table below describes the significant fields shown in the displays.

Table 38: show ip route vrf Field Descriptions

Field	Description
Routing entry for 10.22.22.0/24	Network number.
Known via ...	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	Metric used to reach the destination network.
Tag	Integer used to tag the route.
type	Indicates whether the route is an L1 type or L2 type of route.
Last update from 10.22.5.10	Indicates the IP address of the device that is the next hop to the remote network and identifies the interface on which the last update arrived.
00:01:07 ago	Specifies the last time the route was updated (in hours:minutes:seconds).

Field	Description
Routing Descriptor Blocks	Displays the next-hop IP address followed by the information source.
10.22.6.10, from 10.11.6.7, 00:01:07 ago	Indicates the next-hop address, the address of the gateway that sent the update, and the time that has elapsed since this update was received (in hours:minutes:seconds).
Route metric	This value is the best metric for this routing descriptor block.
Traffic share count	Indicates the number of packets transmitted over various routes.
AS Hops	Number of hops to the destination or to the device where the route first enters internal BGP (iBGP).

The following is sample output from the **show ip route vrf** command on devices using the Cisco IOS Software Modularity for Layer 3 VPNs feature. The output includes remote label information and corresponding MPLS flags for prefixes that have remote labels stored in the RIB if BGP is the label distribution protocol.

```
Device# show ip route vrf v2 10.2.2.2

Routing entry for 10.2.2.2/32
  Known via "bgp 1", distance 200, metric 0, type internal
  Redistributing via ospf 2
  Advertised by ospf 2 subnets
  Last update from 10.0.0.4 00:22:59 ago
  Routing Descriptor Blocks:
  * 10.0.0.4 (Default-IP-Routing-Table), from 10.0.0.31, 00:22:59 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: 1300
      MPLS Flags: MPLS Required
```

The table below describes the significant fields shown in the display.

Table 39: show ip route vrf Field Descriptions

Field	Description
MPLS label	Displays the BGP prefix from the BGP peer. The output shows one of the following values: • A label value (16–1048575). • A reserved label value, such as explicit-null or implicit-null. • The word "none" if no label is received from the peer. The MPLS label field is not displayed if any of the following conditions is true: • BGP is not the Label Distribution Protocol (LDP). However, Open Shortest Path First (OSPF) prefixes learned via sham links display an MPLS label. • MPLS is not supported. • The prefix is imported from another VRF, where the prefix was an Interior Gateway Protocol (IGP) prefix and LDP provided the remote label for it.
MPLS Flags	Name of the MPLS flag. One of the following MPLS flags is displayed: • MPLS Required—Indicates that packets are forwarded to this prefix because of the presence of the MPLS label stack. If MPLS is disabled on the outgoing interface, the packets are dropped. • No Global—Indicates that MPLS packets for this prefix are forwarded from the VRF interface and not from the interface in the global table. VRF interfaces prevent loops in scenarios that use iBGP multipaths. • NSF—Indicates that the prefix is from a nonstop forwarding (NSF)-aware neighbor. If the routing information temporarily disappears due to a disruption in the control plane, packets for this prefix are preserved.

The following sample output from the **show ip route vrf** command shows repair paths in the routing table. The fields in the display are self-explanatory.

```
Device> show ip route vrf test1 repair-paths 192.168.3.0

Routing Table: test1
Routing entry for 192.168.3.0/24
  Known via "bgp 10", distance 20, metric 0
  Tag 100, type external
  Last update from 192.168.1.1 00:49:39 ago
  Routing Descriptor Blocks:
  * 192.168.1.1, from 192.168.1.1, 00:49:39 ago, recursive-via-conn
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none
    [RPR]10.4.4.4 (default), from 10.5.5.5, 00:49:39 ago, recursive-via-host
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: 29
MPLS Flags: MPLS Required, No Global
```

Related Commands

Command	Description
show ip cache	Displays the Cisco Express Forwarding table associated with a VRF.
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip rsvp fast bw-protect

To display information about whether backup bandwidth protection is enabled and the status of backup tunnels that may be used to provide that protection, use the **showiprsvpfastbw-protect** command in user EXEC or privileged EXEC mode.

show ip rsvp fast bw-protect [detail] [**filter** [**destination** *ip-address*| *hostname*] [**dst-port** *port-number*] [**source** *ip-address*| *hostname*] [**src-port** *port-number*]]

Syntax Description

detail	(Optional) Specifies additional receiver information.
filter	(Optional) Specifies a subset of the receivers to display .
destination *ip-address*	(Optional) Specifies the destination IP address of the receiver.
hostname	(Optional) Specifies the hostname of the receiver.
dst-port *port-number*	(Optional) Specifies the destination port number. Valid destination port numbers must be in the range from 0 to 65535.
source *ip-address*	(Optional) Specifies the source IP address of the receiver.
src-port *port-number*	(Optional) Specifies the source port number. Valid source port numbers must be in the range from 0 to 65535.

Command Default

The backup bandwidth protection and backup tunnel status information is not displayed.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(29)S	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Release	Modification
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T

Examples

The following is sample output from the **showiprsvpfastbw-protect** command:

```
Router# show ip rsvp fast bw-protect

Primary          Protect  BW         Backup
Tunnel           I/F      BPS:Type   Tunnel:Label  State   BW-P    Type
--------------   -------  --------   ----------    -----   ----    ----
PRAB-72-5_t500   PO2/0    500K:S     Tu501:19      Ready   ON      Nhop
PRAB-72-5_t601   PO2/0    103K:S     Tu501:20      Ready   OFF     Nhop
PRAB-72-5_t602   PO2/0    70K:S      Tu501:21      Ready   ON      Nhop
PRAB-72-5_t603   PO2/0    99K:S      Tu501:22      Ready   ON      Nhop
PRAB-72-5_t604   PO2/0    100K:S     Tu501:23      Ready   OFF     Nhop
PRAB-72-5_t605   PO2/0    101K:S     Tu501:24      Ready   OFF     Nhop
```

The table below describes the significant fields shown in the display.

Table 40: show ip rsvp fast bw-protect Field Descriptions

Field	Description
Primary Tunnel	Identification of the tunnel being protected.
Protect I/F	Interface name.
BW BPS:Type	Bandwidth, in bits per second, and type of bandwidth. Possible values are the following: • S--Subpool • G--Global pool
Backup Tunnel:Label	Identification of the backup tunnel.
State	Status of backup tunnel. Valid values are the following: • Ready--Data is passing through the primary tunnel, but the backup tunnel is ready to take over if the primary tunnel goes down. • Active--The primary tunnel is down, so the backup tunnel is used for traffic. • None--There is no backup tunnel.
BW-P	Status of backup bandwidth protection. Possible values are ON and OFF.

Field	Description
Type	Type of backup tunnel. Possible values are the following: • Nhop--Next hop • NNHOP--Next-next hop

Related Commands

Command	Description
tunnel mpls traffic-eng fast-reroute bw-protect	Enables an MPLS TE tunnel to use an established backup tunnel in the event of a link or node failure.

show ip rsvp fast detail

To display specific information for Resource Reservation Protocol (RSVP) categories, use the **showiprsvpfastdetail**command in user EXEC or privileged EXEC mode.

show ip rsvp fast detail [**filter** [**destination** *ip-address*| *hostname*] [**dst-port** *port-number*] [**source** *ip-address*| *hostname*] [**src-port** *port-number*]]

Syntax Description

filter	(Optional) Specifies a subset of the receivers to display .
destination *ip-address*	(Optional) Specifies the destination IP address of the receiver.
hostname	(Optional) Specifies the hostname of the receiver.
dst-port *port-number*	(Optional) Specifies the destination port number. Valid destination port numbers must be in the range from 0 to 65535.
source *ip-address*	(Optional) Specifies the source IP address of the receiver.
src-port *port-number*	(Optional) Specifies the source port number. Valid source port numbers must be in the range from 0 to 65535.

Command Default

Specific information for RSVP categories is not displayed.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.0(29)S	Bandwidth Prot desired was added in the Flag field of the command output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.

Examples

The following is sample output from the **showiprsvpfastdetail** command:

```
Router# show ip rsvp fast detail

PATH:
  Tun Dest:   10.0.0.7  Tun ID: 500  Ext Tun ID: 10.0.0.5
  Tun Sender: 10.0.0.5  LSP ID: 8
  Path refreshes:
    sent:     to   NHOP 10.5.6.6 on POS2/0
  Session Attr:
    Setup Prio: 7, Holding Prio: 7
    Flags: Local Prot desired, Label Recording, SE Style, Bandwidth Prot desired
    Session Name: PRAB-72-5_t500
  ERO: (incoming)
    10.0.0.5 (Strict IPv4 Prefix, 8 bytes, /32)
    10.0.5.6 (Strict IPv4 Prefix, 8 bytes, /32)
    10.6.7.7 (Strict IPv4 Prefix, 8 bytes, /32)
    10.0.0.7 (Strict IPv4 Prefix, 8 bytes, /32)
  ERO: (outgoing)
    10.5.6.6 (Strict IPv4 Prefix, 8 bytes, /32)
    10.6.7.7 (Strict IPv4 Prefix, 8 bytes, /32)
    10.0.0.7 (Strict IPv4 Prefix, 8 bytes, /32)
  Traffic params - Rate: 500K bits/sec, Max. burst: 1K bytes
    Min Policed Unit: 0 bytes, Max Pkt Size 4294967295 bytes
  Fast-Reroute Backup info:
    Inbound  FRR: Not active
    Outbound FRR: Ready -- backup tunnel selected
      Backup Tunnel: Tu501      (label 19)
      Bkup Sender Template:
        Tun Sender: 10.5.6.5  LSP ID: 8
      Bkup FilerSpec:
        Tun Sender: 10.5.6.5, LSP ID: 8
  Path ID handle: 04000405.
  Incoming policy: Accepted. Policy source(s): MPLS/TE
  Status: Proxied
  Output on POS2/0. Policy status: Forwarding. Handle: 02000406
```

The table below describes the significant fields shown in the display.

Table 41: show ip rsvp fast detail Field Descriptions

Field	Description
Tun Dest	IP address of the receiver.
Tun ID	Tunnel identification number.
Ext Tun ID	Extended tunnel identification number.
Tun Sender	IP address of the sender.
LSP ID	Label-switched path identification number.
Setup Prio	Setup priority.
Holding Prio	Holding priority.
Flags	Backup bandwidth protection has been configured for the label-switched path (LSP).

Field	Description
Session Name	Name of the session.
ERO (incoming)	EXPLICIT_ROUTE object of incoming path messages.
ERO (outgoing)	EXPLICIT_ROUTE object of outgoing path messages.
Traffic params Rate	Average rate, in bits per second.
Max. burst	Maximum burst size, in bytes.
Min Policed Unit	Minimum policed units, in bytes.
Max Pkt Size	Maximum packet size, in bytes.
Inbound FRR	Status of inbound Fast Reroute (FRR) backup tunnel. If this node is downstream from a rerouted LSP (for example, at a merge point for this LSP), the state is Active.
Outbound FRR	Status of outbound FRR backup tunnel. If this node is a point of local repair (PLR) for an LSP, there are three possible states: • Active--This LSP is actively using its backup tunnel, presumably because there has been a downstream failure. • No Backup--This LSP does not have local (Fast Reroute) protection. No backup tunnel has been selected for it to use in case of a failure. • Ready--This LSP is ready to use a backup tunnel in case of a downstream link or node failure. A backup tunnel has been selected for it to use.
Backup Tunnel	If the Outbound FRR state is Ready or Active, this field indicates the following: • Which backup tunnel has been selected for this LSP to use in case of a failure. • The inbound label that will be prepended to the LSP's data packets for acceptance at the backup tunnel tail (the merge point).

Field	Description
Bkup Sender Template	If the Outbound FRR state is Ready or Active, SENDER_TEMPLATE and FILTERSPEC objects are shown. These objects will be used in RSVP messages sent by the backup tunnel if or when the LSP starts actively using the backup tunnel. They differ from the original (prefailure) objects only in that the node (the PLR) substitutes its own IP address for that of the original sender. For example, path and pathTear messages will contain the new SENDER_TEMPLATE. Resv and resvTear messages will contain the new FILTERSPEC object. If this LSP begins actively using the backup tunnel, the display changes.
Bkup FilerSpec	If the Outbound FRR state is Ready or Active, SENDER_TEMPLATE and FILTERSPEC objects are shown. These objects will be used in RSVP messages sent by the backup tunnel if or when the LSP starts actively using the backup tunnel. They differ from the original (prefailure) objects only in that the node (the PLR) substitutes its own IP address for that of the original sender. For example, path and pathTear messages will contain the new SENDER_TEMPLATE. Resv and resvTear messages will contain the new FILTERSPEC object. If this LSP begins actively using the backup tunnel, the display changes.
Path ID handle	Protection Switch Byte (PSB) identifier.
Incoming policy	Policy decision of the LSP. If RSVP policy was not granted for the incoming path message for the tunnel, the LSP does not come up. Accepted is displayed.
Policy source(s)	For FRR LSPs, this value always is MPLS/TE for the policy source.
Status	For FRR LSPs, valid values are as follows: • Proxied--Headend routers. • Proxied Terminated--Tailend routers. For midpoint routers, the field always is blank.

Related Commands

Command	Description
mpls traffic-eng fast-reroute backup-prot-preemption	Changes the backup protection preemption algorithm to minimize the amount of bandwidth that is wasted.

show ip rsvp hello

To display hello status and statistics for Fast Reroute, reroute (hello state timer), and graceful restart, use the **showiprsvphello** command in user EXEC or privileged EXEC mode.

show ip rsvp hello

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.0(29)S	The command output was modified to include graceful restart, reroute (hello state timer), and Fast Reroute information.
12.2(18)SXD1	This command was integrated into Cisco IOS Release 12.2(18)SXD1.
12.2(33)SRA	The command output was modified to show whether graceful restart is configured and full mode was added.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.2(33)SRC	The command output was modified to include Bidirectional Forwarding Detection (BFD) protocol information.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Examples

The following is sample output from the **showiprsvphello** command:

```
Router# show ip rsvp hello
Hello:
 RSVP Hello for Fast-Reroute/Reroute: Enabled
  Statistics: Disabled
 BFD for Fast-Reroute/Reroute: Enabled
 RSVP Hello for Graceful Restart: Disabled
```

The table below describes the significant fields shown in the display. The fields describe the processes for which hello is enabled or disabled.

Table 42: show ip rsvp hello Field Descriptions

Field	Description
RSVP Hello for Fast-Reroute/Reroute	Status of Fast-Reroute/Reroute: • Disabled--Fast reroute and reroute (hello for state timer) are not activated (disabled). • Enabled--Fast reroute and reroute (hello for state timer) are activated (enabled).
Statistics	Status of hello statistics: • Disabled--Hello statistics are not configured. • Enabled--Statistics are configured. Hello packets are time-stamped when they arrive in the hello input queue for the purpose of recording the time required until they are processed. • Shutdown--Hello statistics are configured but not operational. The input queue is too long (that is, more than 10,000 packets are queued).
BFD for Fast-Reroute/Reroute	Status of BFD for Fast-Reroute/Reroute: • Disabled--BFD is not configured. • Enabled--BFD is configured.
Graceful Restart	Restart capability: • Disabled--Restart capability is not activated. • Enabled--Restart capability is activated for a router (full mode) or its neighbor (help-neighbor).

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables hello globally on the router.
ip rsvp signalling hello statistics	Enables hello statistics on the router.
show ip rsvp hello statistics	Displays how long hello packets have been in the hello input queue.

show ip rsvp hello bfd nbr

To display information about all Multiprotocol Label Switching (MPLS) traffic engineering (TE) clients that use the Bidirectional Forwarding Detection (BFD) protocol, use the **show ip rsvp hello bfd nbr** command in user EXEC or privileged EXEC mode.

show ip rsvp hello bfd nbr

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.2(33)SRC	This command was introduced.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.
15.2(2)SNG	This command was integrated into Cisco ASR 901 Series Aggregation Services Routers.
15.3(1)S	This command was integrated into Cisco IOS Release 15.3(1)S.

Usage Guidelines

The command output is the same as the **show ip rsvp hello bfd nbr summary** command output.

Examples

The following is sample output from the **show ip rsvp hello bfd nbr** command.

```
Router# show ip rsvp hello bfd nbr
Client  Neighbor    I/F     State  LostCnt  LSPs
FRR     10.0.0.6    Gi9/47  Up     0        1
```

The table below describes the significant fields shown in the display.

Table 43: show ip rsvp hello bfd nbr Field Descriptions

Field	Description
Client	MPLS TE feature that is using the BFD protocol.
Neighbor	IP address of the next-hop (that is, the neighbor).
I/F	Outbound (egress) interface name.
State	Status of the BFD session (Up, Down, or Lost).

Field	Description
LostCnt	Number of times that the BFD session is lost (dropped) on this interface.
LSPs	Number of label-switched paths (LSPs) that BFD is protecting on this interface.

Related Commands

Command	Description
clear ip rsvp hello bfd	Globally resets to zero the number of times that the BFD protocol was dropped on an interface or the number of times that a link was down.
ip rsvp signalling hello bfd (configuration)	Enables the BFD protocol globally on the router for MPLS TE link and node protection.
ip rsvp signalling hello bfd (interface)	Enables the BFD protocol on an interface for MPLS TE link and node protection.
show ip rsvp hello bfd nbr detail	Displays detailed information about all MPLS TE clients that use the BFD protocol.
show ip rsvp hello bfd nbr summary	Displays summarized information about all MPLS TE clients that use the BFD protocol.

show ip rsvp hello bfd nbr detail

To display detailed information about all Multiprotocol Label Switching (MPLS) traffic engineering (TE) clients that use the Bidirectional Forwarding Detection (BFD) protocol, use the **show ip rsvp hello bfd nbr detail** command in user EXEC or privileged EXEC mode.

show ip rsvp hello bfd nbr detail

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.2(33)SRC	This command was introduced.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.
15.3(1)S	This command was integrated into Cisco IOS Release 15.3(1)S.

Examples

The following is sample output from the **show ip rsvp hello bfd nbr detail** command:

```
Router# show ip rsvp hello bfd nbr detail
 Hello Client Neighbors
 Remote addr 10.0.0.6, Local addr  10.0.0.7
  Type: Active
  I/F: Gi9/47
  State: Up (for 00:09:41)
  Clients: FRR
  LSPs protecting: 1 (frr: 1, hst upstream: 0 hst downstream: 0)
  Communication with neighbor lost: 0
```

The table below describes the significant fields shown in the display.

Table 44: show ip rsvp hello bfd nbr detail Field Descriptions

Field	Description
Remote addr	IP address of the next hop interface.
Local addr	IP address of the outbound interface.
Type	Type of signaling that is in effect (Active or Passive).
I/F	Interface name.
State	Status of the BFD session (Up, Down, or Lost).

Field	Description
Clients	Software that is using the BFD protocol.
LSPs protecting	Number of label switched paths (LSPs) that the BFD protocol is protecting.
Communication with neighbor lost	Number of times the BFD protocol detected that a link was down.

Related Commands

Command	Description
clear ip rsvp hello bfd	Globally resets to zero the number of times that the BFD protocol was dropped on an interface or the number of times that a link was down.
ip rsvp signalling hello bfd (configuration)	Enables the BFD protocol globally on the router for MPLS TE link and node protection.
ip rsvp signalling hello bfd (interface)	Enables the BFD protocol on an interface for MPLS TE link and node protection.
show ip rsvp hello bfd nbr	Displays information about all MPLS TE clients that use the BFD protocol.
show ip rsvp hello bfd nbr summary	Displays summarized information about all MPLS TE clients that use the BFD protocol.

show ip rsvp hello bfd nbr summary

To display summarized information about all Multiprotocol Label Switching (MPLS) traffic engineering (TE) clients that use the Bidirectional Forwarding Detection (BFD) protocol, use the **show ip rsvp hello bfd nbr summary** command in user EXEC or privileged EXEC mode.

show ip rsvp hello bfd nbr summary

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.2(33)SRC	This command was introduced.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.
15.3(1)S	This command was integrated into Cisco IOS Release 15.3(1)S.

Usage Guidelines

The command output is the same as the **show ip rsvp hello bfd nbr** command output.

Examples

The following is sample output from the **show ip rsvp hello bfd nbr summary** command.

```
Router# show ip rsvp hello bfd nbr summary

Client  Neighbor   I/F     State  LostCnt  LSPs
FRR     10.0.0.6   Gi9/47  Up     0        1
```

The table below describes the significant fields shown in the display.

Table 45: show ip rsvp hello bfd nbr summary Field Descriptions

Field	Description
Client	MPLS TE feature that uses the BFD protocol.
Neighbor	IP address of the next hop (that is, the neighbor).
I/F	Interface type and slot or port.
State	Status of the BFD session (Up, Down, or Lost).
LostCnt	Number of times that the BFD session is lost (dropped) on this interface.

Field	Description
LSPs	Number of label switched paths (LSPs) that BFD is protecting on this interface.

Related Commands

Command	Description
clear ip rsvp hello bfd	Globally resets to zero the number of times that the BFD protocol was dropped on an interface or the number of times that a link was down.
ip rsvp signalling hello bfd (configuration)	Enables the BFD protocol globally on the router for MPLS TE link and node protection.
ip rsvp signalling hello bfd (interface)	Enables the BFD protocol globally on an interface for MPLS TE link and node protection.
show ip rsvp hello bfd nbr	Displays information about all MPLS TE clients that use the BFD protocol.
show ip rsvp hello bfd nbr detail	Displays detailed information about all MPLS TE clients that use the BFD protocol.

show ip rsvp hello instance detail

To display detailed information about a hello instance, use the **showiprsvphelloinstancedetail**command in user EXEC or privileged EXEC mode.

show ip rsvp hello instance detail [**filter destination** *ip-address*]

Syntax Description

filter destination *ip-address*	(Optional) IP address of the neighbor node.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.0(29)S	The command output was modified to include graceful restart, hello state timer (reroute), and fast reroute information.
12.2(18)SXD1	This command was integrated into Cisco IOS Release 12.2(18)SXD1.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.

Usage Guidelines

Use the **showiprsvphelloinstancedetail** command to display information about the processes (clients) currently configured.

Examples

The following is sample output from the **showiprsvphelloinstancedetail** command:

```
Router# show ip rsvp hello instance detail
Neighbor 10.0.0.3  Source  10.0.0.2
    Type: Active    (sending requests)
    I/F:  Serial2/0
    State:   Up        (for 2d19h2d19h)
    Clients: ReRoute
    LSPs protecting: 1
    Missed acks: 4, IP DSCP: 0x30
    Refresh Interval (msec)
      Configured: 6000
      Statistics: (from 40722 samples)
        Min:      6000
        Max:      6064
        Average:  6000
        Waverage: 6000 (Weight = 0.8)
```

```
      Current:  6000
  Last sent Src_instance: 0xE617C847
  Last recv nbr's Src_instance: 0xFEC28E95
  Counters:
    Communication with neighbor lost:
      Num times:                    0
      Reasons:
        Missed acks:                0
        Bad Src_Inst received:      0
        Bad Dst_Inst received:      0
        I/F went down:              0
        Neighbor disabled Hello:    0
    Msgs Received:   55590
         Sent:       55854
         Suppressed: 521
Neighbor 10.0.0.8 Source  10.0.0.7
  Type: Passive   (responding to requests)
  I/F:  Serial2/1
  Last sent Src_instance: 0xF7A80A52
  Last recv nbr's Src_instance: 0xD2F1B7F7
  Counters:
    Msgs Received:   199442
         Sent:       199442
```

The table below describes the significant fields shown in the display.

Table 46: show ip rsvp hello instance detail Field Descriptions

Field	Description
Neighbor	IP address of the adjacent node.
Source	IP address of the node that is sending the hello message.
Type	Values are Active (node is sending a request) and Passive (node is responding to a request).
I/F	Interface from which hellos are sent for this instance. Any means that the hellos can be sent out any interface.
State	Status of communication. Values are as follows: • Up--Node is communicating with its neighbor. • Lost--Communication has been lost. • Init--Communication is being established.
Clients	Clients that created this hello instance; they include graceful restart, ReRoute (hello state timer), and Fast Reroute.
LSPs protecting	Number of LSPs that are being protected by this hello instance.
Missed acks	Number of times that communication was lost due to missed acknowledgments (ACKs).

Field	Description
IP DSCP	IP differentiated services code point (DSCP) value used in the hello IP header.
Refresh Interval (msec)	The frequency (in milliseconds) with which a node generates a hello message containing a Hello Request object for each neighbor whose status is being tracked.
Configured	Configured refresh interval.
Statistics	Refresh interval statistics from a specified number of samples (packets).
Min	Minimum refresh interval.
Max	Maximum refresh interval.
Average	Average refresh interval.
Waverage	Weighted average refresh interval.
Current	Current refresh interval.
Last sent Src_instance	The last source instance sent to a neighbor.
Last recv nbr's Src_instance	The last source instance field value received from a neighbor. (0 means none received.)
Counters	Incremental information relating to communication with a neighbor.
Num times	Total number of times that communication with a neighbor was lost.
Reasons	Subsequent fields designate why communication with a neighbor was lost.
Missed acks	Number of times that communication was lost due to missed ACKs.
Bad Src_Inst received	Number of times that communication was lost due to bad source instance fields.
Bad Dst_Inst received	Number of times that communication was lost due to bad destination instance fields.
I/F went down	Number of times that the interface became unoperational.

Field	Description
Neighbor disabled Hello	Number of times that a neighbor disabled hello messages.
Msgs Received	Number of messages that were received.
Sent	Number of messages that were sent.
Suppressed	Number of messages that were suppressed due to optimization.

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables hello globally on the router.
ip rsvp signalling hello statistics	Enables hello statistics on the router.
show ip rsvp hello	Displays hello status and statistics for Fast reroute, reroute (hello state timer), and graceful restart.
show ip rsvp hello instance summary	Displays summary information about a hello instance.

show ip rsvp hello instance summary

To display summary information about a hello instance, use the **showiprsvphelloinstancesummary** command in user EXEC or privileged EXEC mode.

show ip rsvp hello instance summary

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC Privileged EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.0(29)S	The command output was modified to include graceful restart, reroute (hello state timer), and fast reroute information.
12.2(18)SXD1	This command was integrated into Cisco IOS Release 12.2(18)SXD1.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.

Examples

The following is sample output from the **showiprsvphelloinstancesummary** command:

```
Router# show ip rsvp hello instance summary
Active Instances:
  Client  Neighbor        I/F        State      LostCnt  LSPs Interval
  RR      10.0.0.3        Se2/0      Up               0     1 6000
  GR      10.1.1.1        Any        Up              13     1 10000
  GR      10.1.1.5        Any        Lost             0     1 10000
  GR      10.2.2.1        Any        Init             1     0 5000
Passive Instances:
  Neighbor        I/F
  10.0.0.1        Se2/1
Active = Actively tracking neighbor state on behalf of clients:
         RR = ReRoute, FRR = Fast ReRoute, or GR = Graceful Restart
Passive = Responding to hello requests from neighbor
```

The table below describes the significant fields shown in the display.

Table 47: show ip rsvp hello instance summary Field Descriptions

Field	Description
Active Instances	Active nodes that are sending hello requests.

Field	Description
Client	Clients on behalf of which hellos are sent; they include GR (graceful restart), RR (reroute = hello state timer), and FRR (Fast Reroute).
Neighbor	IP address of the adjacent node. For graceful restart, this is the neighbor router's ID; for Fast Reroute and hello state timer (reroute), this is one of the neighbor's interface addresses.
I/F	Interface from which hellos are sent for this instance. Any means that the hellos can be sent out any interface.
State	Status of communication. Values are as follows: • Up--Node is communicating with its neighbor. • Lost--Communication has been lost. • Init--Communication is being established.
LostCnt	Number of times that communication was lost with the neighbor.
LSPs	Number of label-switched paths (LSPs) protected by this hello instance.
Interval	Hello refresh interval in milliseconds.
Passive Instances	Passive nodes that are responding to hello requests.
Neighbor	IP address of adjacent node. For graceful restart, this is the neighbor router's ID; for Fast Reroute and hello state timer (reroute), this is one of the neighbor's interface addresses.
I/F	Interface from which hellos are sent for this instance. Any means that the hellos can be sent out any interface.

Related Commands

Command	Description
ip rsvp signalling hello (configuration)	Enables hello globally on the router.
ip rsvp signalling hello statistics	Enables hello statistics on the router.

Command	Description
show ip rsvp hello	Displays hello status and statistics for fast reroute, reroute (hello state timer), and graceful restart.
show ip rsvp hello instance detail	Displays detailed information about a hello instance.

show ip rsvp hello statistics

To display how long hello packets have been in the Hello input queue, use the **showiprsvphellostatistics**command in privileged EXEC mode.

show ip rsvp hello statistics

Syntax Description

This command has no arguments or keywords.

Command Default

Information about how long hello packets have been in the Hello input queue is not displayed.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.2(18)SXD1	This command was integrated into Cisco IOS Release 12.2(18)SXD1.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T

Usage Guidelines

You can use this command to determine if the Hello refresh interval is too small. If the interval is too small, communication may falsely be declared as lost.

Examples

The following is sample output from the **showiprsvphellostatistics** command:

```
Router# show ip rsvp hello statistics

Status: Enabled
  Packet arrival queue:
    Wait times (msec)
      Current:0
      Average:0
      Weighted Average:0 (weight = 0.8)
      Max:4
    Current length: 0 (max:500)
  Number of samples taken: 2398525
```

The table below describes the significant fields shown in the display.

Table 48: show ip rsvp hello statistics Field Descriptions

Field	Description
Status	Indicator of whether Hello has been enabled globally on the router.
Current	Amount of time, in milliseconds, that the current hello packet has been in the Hello input queue.
Average	Average amount of time, in milliseconds, that hello packets are in the Hello input queue.
Max	Maximum amount of time, in milliseconds, that hello packets have been in the Hello input queue.
Current length	Current amount of time, in milliseconds, that hello packets have been in the Hello input queue.
Number of samples taken	Number of packets for which these statistics were compiled.

Related Commands

Command	Description
clear ip rsvp hello instance statistics	Clears Hello statistics for an instance.
clear ip rsvp hello statistics	Globally clears Hello statistics.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.
ip rsvp signalling hello statistics	Enables Hello statistics on the router.

show ip rsvp high-availability database

To display contents of Resource Reservation Protocol (RSVP) high availability (HA) read and write databases used in traffic engineering (TE), use the **show ip rsvp high-availability database** command in user EXEC or privileged EXEC mode.

show ip rsvp high-availability database {**hello**| **if-autotun**| **link-management** {**interfaces** [**fixed**| **variable**]| **system**}| **lsp** [**filter** [**destination** *ip-address*]| [**lsp-id** *lsp-id*]| [**source** *ip-address*]| [**tunnel-id** *tunnel-id*]]| **lsp-head** [**filter** *number*]| **summary**}

Syntax Description

hello	Displays information about hello entries in read and write databases.
if-autotun	Displays information about TE HA autotunnel interface entries in read and write databases.
link-management	Displays information about link-management entries in the read and write databases.
interfaces	Displays information about link-management interfaces in the read and write databases.
fixed	(Optional) Displays information about link-management fixed interfaces in the read and write databases.
variable	(Optional) Displays information about link-management variable interfaces in the read and write databases.
system	Displays information about the link-management system in the read and write databases.
lsp	Displays information about label switched path (LSP) entries in the read and write databases.
filter destination *ip-address*	(Optional) Displays filtered information on the IP address of the destination (tunnel tail).
filter lsp-id *lsp-id*	(Optional) Displays filtered information on a specific LSP ID designated by a number from 0 to 65535.
filter source *ip-address*	(Optional) Displays filtered information on the IP address of the source (tunnel head).
filter tunnel-id *tunnel-id*	(Optional) Displays filtered information on a specific tunnel ID designated by a number from 0 to 65535.

lsp-head	Displays information about LSP-head entries in the read and write databases.
filter *number*	(Optional) Displays filtered information on a specific LSP-head router designated by a number from 0 to 65535.
summary	Displays cumulative information about entries in read and write databases.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SRA	This command was introduced.
12.2(33)SRB	The command output was modified to display the result of a loose hop expansion performed on the router.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.2(33)SRC	This command was integrated into Cisco IOS Release 12.2(33)SRC. The command output was modified to include path protection information specified by the **lsp-head** keyword.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S. The command output was modified to distinguish database-entry information for point-to-point (P2P) tunnels from that for point-to-multipoint (P2MP) tunnels and to display error database information.
12.2(50)SY	This command was integrated into Cisco IOS Release 12.2(50)SY.
Cisco IOS XE Release 3.5S	This command was integrated into Cisco IOS XE Release 3.5S.
15.2(2)S	This command was modified. The **if-autotun** keyword was added. The output for the **show ip rsvp high-availability database lsp**, the **show ip rsvp high-availability database lsp-head**, and the **show ip rsvp high-availability database summary** commands was enhanced to display checkpoint information for MPLS TE autotunnel and automesh stateful switchover (SSO) tunnels.

Release	Modification
Cisco IOS XE Release 3.6S	This command was modified. The **if-autotun** keyword was added. The output for the **show ip rsvp high-availability database lsp**, the **show ip rsvp high-availability database lsp-head**, and the **show ip rsvp high-availability database summary** commands was enhanced to display checkpoint information for MPLS TE autotunnel and automesh stateful switchover (SSO) tunnels.

Usage Guidelines

Use the **show ip rsvp high-availability database** command to display information about entries in the read and write databases.

Use the **show ip rsvp high-availability database lsp** command to display loose hop information. A loose hop expansion can be performed on a router when the router processes the explicit router object (ERO) for an incoming path message. After the router removes all local IP addresses from the incoming ERO, it finds the next hop. If the ERO specifies that the next hop is loose instead of strict, the router consults the TE topology database and routing to determine the next hop and output interface to forward the path message. The result of the calculation is a list of hops; the list is placed in the outgoing ERO and checkpointed with the LSP data as the loose hop information.

In Cisco IOS Release 15.0(1)S and later releases, the **show ip rsvp high-availability database lsp** command displays sub-LSP information. If any sub-LSP, whether P2MP or P2P, fails to recover after a stateful switchover (SSO), the failure is noted in an error database for troubleshooting. You can use the **show ip rsvp high-availability database lsp** command to display error database entries.

You can use the **show ip rsvp high-availability database lsp-head** command only on a headend router; this command gives no information on other routers

Examples

Examples

The following is sample output from the **show ip rsvp high-availability database hello** command on an active Route Processor (RP):

```
Router# show ip rsvp high-availability database hello

HELLO WRITE DB
   Header:
     State: Checkpointed     Action: Add
     Seq #: 1                Flags: 0x0
   Data:
     Last sent Src_instance: 0xDE435865
HELLO READ DB
```

The table below describes the significant fields shown in the display.

Table 49: show ip rsvp high-availability database hello—Active RP Field Descriptions

Field	Description
HELLO WRITE DB	Storage area for active RP hello data consisting of checkpointed RSVP-TE information that is sent to the standby RP when it becomes the active RP and needs to recover LSPs. This field is blank on a standby RP.

Field	Description
Header	Header information.
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are waiting to be sent.
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This is a temporary action that takes place while the active RP awaits an acknowledgment (ack) of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acknowledgments (acks) and negative acknowledgments (nacks) to sent messages.
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
Last sent Src_instance	Last sent source instance identifier.
HELLO READ DB	Storage area for standby RP hello data. This field is blank on an active RP, except when it is in recovery mode.

Examples

The following is sample output from the **show ip rsvp high-availability database hello** on a standby RP:

```
Router# show ip rsvp high-availability database hello

HELLO WRITE DB
HELLO READ DB
   Header:
     State: Checkpointed     Action: Add
     Seq #: 1                Flags: 0x0
   Data:
     Last sent Src_instance: 0xDE435865
```

These fields are the same as those for the active RP described in the table except they are now in the read database for the standby RP.

Examples

The following is sample output from the **show ip rsvp high-availability database if-autotun** command on an active RP.

```
Router# show ip rsvp high-availability database if-autotun
IF_AUTOTUN WRITE DB

  Header:
    State: Checkpointed     Action: Add
    Seq #: 1                Flags: 0x0
  Data:
    Tunnel ID: 1000 (if_handle: 85), prot_if_handle: 14
    template_unit: n/a, dest: 22.22.22.22, flags=0x0

  Header:
    State: Checkpointed     Action: Add
    Seq #: 61               Flags: 0x0
  Data:
    Tunnel ID: 2000 (if_handle: 86), prot_if_handle: 14
    template_unit: n/a, dest: 22.22.22.22, flags=0x1

  Header:
    State: Checkpointed     Action: Add
    Seq #: 1                Flags: 0x0
  Data:
    Tunnel ID: 3000 (if_handle: 87), prot_if_handle: 0
    template_unit: 1, dest: 22.22.22.22, flags=0x2

  Header:
    State: Checkpointed     Action: Add
    Seq #: 1                Flags: 0x0
  Data:
    Tunnel ID: 3001 (if_handle: 88), prot_if_handle: 0
    template_unit: 1, dest: 172.16.255.128, flags=0x2

  Header:
    State: Checkpointed     Action: Add
    Seq #: 1                Flags: 0x0
  Data:
    Tunnel ID: 3002 (if_handle: 89), prot_if_handle: 0
    template_unit: 1, dest: 200.0.0.0, flags=0x2


IF_AUTOTUN READ DB
```

The table below describes the significant fields shown in the display.

Table 50: show ip rsvp high-availability database if-autotun—Active RP Field Descriptions

Field	Description
IF_AUTOTUN WRITE DB	Storage area for active RP autotunnel interface information. This field is blank on a standby RP.
Header	Header information.

Field	Description
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are still waiting to be sent.
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This action appears temporarily while the active RP awaits an ack of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acks and nacks to sent messages.
Flags	Attributes used to identify or track data.
Data	Information about the last transmission.
Tunnel ID	Tunnel identifier.
if_handle	Internal number representing the autotunnel interface. For the same tunnel ID, this if_handle value should always be the same for the record in the Standby READ DB as in the Active WRITE DB.
prot_if_handle	For autotunnel mesh tunnels, this value should always be zero. For autotunnel primary tunnels, this is an internal number representing the egress interface of the autotunnel primary. For autotunnel backup tunnels, this is an internal number representing the interface that the backup is protecting. In all three cases, for the same tunnel ID, this value should always be the same for the record in the Standby READ DB as in the Active WRITE DB.

Field	Description
template_unit	For autotunnel mesh, this represents the auto-template interface number that the mesh tunnel was created from. For autotunnel primary and backup, this should be "n/a."
dest	Destination IP address of the autotunnel.
flags	Encodings have these values: • 0 = autotunnel primary • 1 = autotunnel backup • 2 = autotunnel mesh
IF_AUTOTUN READ DB	Storage area for standby RP autotunnel interface information. This field is blank on an active RP.

The fields for a standby RP are the same as those described in the table except that they are now in the interface autotunnel read database instead of the interface autotunnel write database that is used by an active RP.

Examples

The following is sample output from the **show ip rsvp high-availability database link-management interfaces** command on an active RP:

```
Router# show ip rsvp high-availability database link-management interfaces

TE LINK WRITE DB
Flooding Protocol: ospf  IGP Area ID: 0  Link ID: 0 (GigabitEthernet3/2)
  Header:
    State: Checkpointed     Action: Add
    Seq #: 4                Flags: 0x0
  Data:
        Ifnumber: 5  Link Valid Flags: 0x193B
        Link Subnet Type: Broadcast
        Local Intfc ID: 0  Neighbor Intf ID: 0
        Link IP Address: 172.16.3.1
        Neighbor IGP System ID: 172.16.3.2  Neighbor IP Address: 10.0.0.0
        IGP Metric: 1  TE Metric: 1
        Physical Bandwidth: 1000000 kbits/sec
        Res. Global BW: 3000 kbits/sec
        Res. Sub BW: 0 kbits/sec
        Upstream::
                                 Global Pool   Sub Pool
                                 -----------   ----------
        Reservable Bandwidth[0]:           0            0 kbits/sec
        Reservable Bandwidth[1]:           0            0 kbits/sec
        Reservable Bandwidth[2]:           0            0 kbits/sec
        Reservable Bandwidth[3]:           0            0 kbits/sec
        Reservable Bandwidth[4]:           0            0 kbits/sec
        Reservable Bandwidth[5]:           0            0 kbits/sec
        Reservable Bandwidth[6]:           0            0 kbits/sec
        Reservable Bandwidth[7]:           0            0 kbits/sec
        Downstream::
                                 Global Pool   Sub Pool
                                 -----------   ----------
        Reservable Bandwidth[0]:        3000            0 kbits/sec
        Reservable Bandwidth[1]:        3000            0 kbits/sec
```

```
Reservable Bandwidth[2]:        3000            0 kbits/sec
Reservable Bandwidth[3]:        3000            0 kbits/sec
Reservable Bandwidth[4]:        3000            0 kbits/sec
Reservable Bandwidth[5]:        3000            0 kbits/sec
Reservable Bandwidth[6]:        3000            0 kbits/sec
Reservable Bandwidth[7]:        2900            0 kbits/sec
Affinity Bits: 0x0
Protection Type: Capability 0,  Working Priority 0
Number of TLVs: 0
```

The table below describes the significant fields shown in the display.

Table 51: show ip rsvp high-availability database link-management interfaces—Active RP Field Descriptions

Field	Description
TE LINK WRITE DB	Storage area for active TE RP link data. This field is blank on a standby RP.
Flooding Protocol	Protocol that is flooding information for this area. OSPF = Open Shortest Path First.
IGP Area ID	Interior Gateway Protocol (IGP) identifier for the area being flooded.
Link ID	Link identifier and interface for the area being flooded.
Header	Header information.
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are waiting to be sent.
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This action appears temporarily while the active RP awaits an ack of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acks and nacks to sent messages.

Field	Description
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
Ifnumber	Interface number.
Link Valid Flags	Attributes used to identify or track links.
Link Subnet Type	Subnet type of the link. Values are as follows: • Broadcast—Data for multiple recipients. • Nonbroadcast Multiaccess--A network in which data is transmitted directly from one computer to another over a virtual circuit or across a switching fabric. • Point-to-Multipoint—Unidirectional connection in which a single source end system (known as a root node) connects to multiple destination end systems (known as leaves). • Point-to-Point—Unidirectional or bidirectional connection between two end systems. • Unknown subnet type—Subnet type not identified.
Local Intfc ID	Local interface identifier.
Neighbor Intf ID	Neighbor's interface identifier.
Link IP Address	IP address of the link.
Neighbor IGP System ID	Neighbor system identifier configured using IGP.
Neighbor IP Address	Neighbor's IP address.
IGP Metric	Metric value for the TE link configured using IGP.
TE Metric	Metric value for the TE link configured using Multiprotocol Label Switching (MPLS) TE.
Physical Bandwidth	Link bandwidth capacity in kilobits per second (kb/s).
Res. Global BW	Amount of reservable global pool bandwidth (in kb/s) on this link.
Res. Sub BW	Amount of reservable subpool bandwidth (in kb/s) on this link.

Field	Description
Upstream	Header for the following section of bandwidth values.
Global Pool	Global pool bandwidth (in kb/s) on this link.
Sub Pool	Subpool bandwidth (in kb/s) on this link.
Reservable Bandwidth [1]	Amount of bandwidth (in kb/s) available for reservations in the global TE topology and subpools.
Downstream	Header for the following section of bandwidth values.
Affinity Bits	Link attributes required in tunnels.
Protection Type	LSPs protected by fast reroute (FRR). • Capability = LSPs capable of using FRR. • Working Priority = LSPs actually using FRR.
Number of TLVs	Number of type, length, values (TLVs).

The fields for a standby RP are the same as those described in the table except that they are now in the TE link read database instead of the TE link write database that is used by an active RP.

Examples

The following is sample output from the **show ip rsvp high-availability database link-management system** command on an active RP:

```
Router# show ip rsvp high-availability database link-management system

TE SYSTEM WRITE DB
Flooding Protocol: OSPF  IGP Area ID: 0
  Header:
    State: Checkpointed     Action: Modify
    Seq #: 4                Flags: 0x0
  Data:
    LM Flood Data::
      LSA Valid flags: 0x0  Node LSA flag: 0x0
      IGP System ID: 172.16.3.1  MPLS TE Router ID: 10.0.0.3
      Flooded links: 1  TLV length: 0 (bytes)
      Fragment id: 0
TE SYSTEM READ DB
```

The table below describes the significant fields shown in the display.

Table 52: show ip rsvp high-availability database link-management system—Active RP Field Descriptions

Field	Description
TE SYSTEM WRITE DB	Storage area for active TE RP system data. This field is blank on a standby RP.

Field	Description
Flooding Protocol	Protocol that is flooding information for this area. OSPF = Open Shortest Path First.
IGP Area ID	IGP identifier for the area being flooded.
Header	Header information.
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are waiting to be sent.
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This action appears temporarily while the active RP awaits an ack of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acks and nacks to messages sent.
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
LM Flood Data	Link management (LM) flood data.
LSA Valid flags	Link-state advertisement (LSA) attributes.
Node LSA flag	LSA attributes used by a router.
IGP System ID	Identification (IP address) that IGP flooding uses in this area to identify this node.
MPLS TE Router ID	MPLS TE router identifier (IP address).
Flooded links	Number of flooded links.

Field	Description
TLV length	TLV length in bytes.
Fragment id	Fragment identifier for this link.
TE SYSTEM READ DB	Storage area for standby TE RP system data. This field is blank on a standby RP.

The fields for a standby RP are the same as those described in the table except that they are now in the TE system read database instead of the TE system write database that is used by an active RP.

Examples

The following is sample output from the **show ip rsvp high-availability database lsp** command on an active RP for a P2P tunnel:

```
Router# show ip rsvp high-availability database lsp

Tun ID: 0   LSP ID: 10   (P2P)
  SubGrp ID:   -
  SubGrp Orig: -
  Dest:   10.3.0.1
  Sender: 10.1.0.1     Ext. Tun ID: 10.1.0.1
  Header:
    State: Checkpointed     Action: Add
    Seq #: 2                Flags: 0x0
  Data:
    PathSet ID: -
    Lspvif if_num: -
    InLabel: -
    Out I/F: Se2/0
    Next-Hop: 10.1.3.2
    OutLabel: 16
    Loose hop info: None (0)
```

Examples

The following is sample output from the **show ip rsvp high-availability database lsp** command on an active RP for a P2MP tunnel:

```
Router# show ip rsvp high-availability database lsp

Tun ID: 1   LSP ID: 127   (P2MP)
  SubGrp ID:   1
  SubGrp Orig: 10.1.0.1
  Dest:   10.2.0.1
  Sender: 10.1.0.1     Ext. Tun ID: 10.1.0.1
  Header:
    State: Checkpointed     Action: Add
    Seq #: 30               Flags: 0x0
  Data:
    PathSet ID: 0x1A000003
    Lspvif if_num: 35 (Lspvif0)
    InLabel: 19
    Out I/F: None
    Next-Hop: -
    OutLabel: -
    Loose hop info: None (0)
```

The table below describes the significant fields shown in the display.

Table 53: show ip rsvp high-availability database lsp—Active RP Field Descriptions

Field	Description
P2P/P2MP	Tunnel type.
Subgrp ID	Subgroup identifier (valid only for P2MP TE LSPs).
Subgrp Orig	Subgroup origin IP address (valid only for P2MP TE LSPs).
Lspvif if_num	Interface number of the LSPVIF (valid only for P2MP TE tailends).
PathSet ID	Path set identifier (valid only for P2MP TE LSPs)
LSP WRITE DB	Storage area for active RP LSP data. This field is blank on a standby RP.
Tun ID	Tunnel identifier.
LSP ID	LSP identifier.
Dest	Tunnel destination IP address.
Sender	Tunnel sender IP address.
Ext. Tun ID	Extended tunnel identifier; usually set to 0 or the sender's IP address.
Header	Header information.
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent, but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are waiting to be sent.

Field	Description
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This action appears temporarily while the active RP awaits an ack of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acks and nacks to messages sent.
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
InLabel	Incoming label identifier.
Out I/F	Outgoing interface.
Next-Hop	Next hop IP address.
OutLabel	Outgoing label identifier.
Loose hop info	Lists the loose hop expansions performed on the router, or specifies None.
LSP READ DB	Storage area for standby RP LSP data. This field is blank on an active RP.

The fields for a standby RP are the same as those described in the table except that they are now in the LSP read database instead of the LSP write database that is used by an active RP.

Examples

The following is sample output from the **show ip rsvp high-availability database lsp-head** command on an active RP for a P2P tunnel:

```
Router# show ip rsvp high-availability database lsp-head

LSP_HEAD WRITE DB
  Tun ID: 0  (P2P)
  Header:
    State: Checkpointed     Action: Add
    Seq #: 2                Flags: 0x0
  Data:
    lsp_id: 10, bandwidth: 5, thead_flags: 0x1, popt: 1
    feature flags: none
```

```
    output_if_num: 11, output_nhop: 10.1.3.2
    RRR path setup info
      Destination: 10.3.0.1, Id: 10.3.0.1 Router Node (ospf) flag:0x0
      IGP: ospf, IGP area: 0, Number of hops: 3, metric: 128
      Hop 0: 10.1.3.2, Id: 10.2.0.1 Router Node (ospf), flag:0x0
      Hop 1: 10.2.3.3, Id: 10.3.0.1 Router Node (ospf), flag:0x0
      Hop 2: 10.3.0.1, Id: 10.3.0.1 Router Node (ospf), flag:0x0
```

Examples

The following is sample output from the **show ip rsvp high-availability database lsp-head** command on an active RP for a P2MP tunnel:

```
Router# show ip rsvp high-availability database lsp-head

LSP_HEAD WRITE DB
Tun ID: 1  (P2MP)
  Destination: 10.2.0.1
  Header:
    State: Checkpointed     Action: Add
    Seq #: 3                Flags: 0x0
  Data:
    lsp_id: 11, bandwidth: 100, thead_flags: 0x1, popt: 1
    Subgrp_id: 1
    feature flags: none
    output_if_num: 3, output_nhop: 10.1.2.2
    RRR path setup info
      Destination: 10.2.0.1, Id: 10.2.0.1 Router Node (ospf) flag:0x0
      IGP: ospf, IGP area: 0, Number of hops: 3, metric: 10
      Hop 0: 10.1.2.1, Id: 10.1.0.1 Router Node (ospf), flag:0x0
      Hop 1: 10.1.2.2, Id: 10.2.0.1 Router Node (ospf), flag:0x0
      Hop 2: 10.2.0.1, Id: 10.2.0.1 Router Node (ospf), flag:0x0
```

The table below describes the significant fields shown in the display.

Table 54: show ip rsvp high-availability database lsp-head—Active RP Field Descriptions

Field	Description
LSP_HEAD WRITE DB	Storage area for active RP LSP-head data. This field is blank on a standby RP.
P2P/P2MP	Tunnel type.
Tun ID	Tunnel identifier.
Header	Header information.
State	Status of an entry. Values are as follows: • Ack-Pending—Entries have been sent, but not acknowledged. • Checkpointed—Entries have been sent and acknowledged by the standby RP. • Send-Pending—Entries are waiting to be sent.

Field	Description
Action	Action taken. Values are as follows: • Add—Adding an item to the standby RP. • Delete—Deleting an item from the standby RP. This is a temporary action that takes place while the active RP awaits an ack of the delete operation. • Modify—Modifying an item on the standby RP. • Remove—Removing an item from the standby RP.
Seq #	Number used by the active and standby RPs to synchronize message acks and nacks to messages sent.
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
lsp_id	LSP identifier.
bandwidth	Bandwidth on the LSP (in kb/s).
thead_flags	Tunnel head attribute used to identify or track data.
popt	Parsing option number.
feature_flags	Indicates whether the LSP being used to forward traffic is the secondary LSP using the path protection path option. Valid values are as follows: • none • path protection active
output_if_num	Output interface number.
output_nhop	Output next hop IP address.
RRR path setup info	Routing with Resource Reservation (RRR) path information.
Destination	Destination IP address.

Field	Description
Id	IP address and protocol of the routing node. Values are as follows: • ISIS = Intermediate System-to-Intermediate System • OSPF = Open Shortest Path First
flag	Attribute used to track data.
IGP	Interior Gateway Protocol. OSPF = Open Shortest Path First.
IGP area	IGP area identifier.
Number of hops	Number of connections or routers.
metric	Routing cost.
Hop	Hop's number and IP address.
LSP_HEAD READ DB	Storage area for standby RP LSP-head data. This field is blank on an active RP.

The fields for a standby RP are the same as those described in the table except that they are now in the LSP_head read database instead of the LSP_head write database that is used by an active RP.

Examples

The following is sample output from the **show ip rsvp high-availability database summary** command on an active RP:

```
Router# show ip rsvp high-availability database summary

Write DB:
  Send-Pending:     0
  Ack-Pending :     0
  Checkpointed:    10
  Total       :    10
Read DB:
  Total       :     0
```

The table below describes the significant fields shown in the display.

Table 55: show ip rsvp high-availability database summary—Active RP Field Descriptions

Field	Description
Write DB	Storage area for active RP summary data. This field is blank on a standby RP.
Send-Pending	Entries are waiting to be sent.

Field	Description
Ack-Pending	Entries have been sent, but are waiting to be acknowledged.
Checkpointed	Entries have been sent and acknowledged.
Total	Total number of entries in the write database.
Total	Total number of entries in the read database.

Examples

The following is sample output from the **show ip rsvp high-availability database summary** command on a standby RP:

```
Router# show ip rsvp high-availability database summary

Write DB:
  Send-Pending:     0
  Ack-Pending :     0
  Checkpointed:     0
  Total       :     0
Read DB:
  Total       :    10
```

The table below describes the significant fields shown in the display.

Table 56: show ip rsvp high-availability database summary—Standby RP Field Descriptions

Field	Description
Write DB	Storage area for active RP summary data.
Send-Pending	Entries are waiting to be sent.
Ack-Pending	Entries have been sent but are waiting to be acknowledged.
Checkpointed	Entries have been sent and acknowledged.
Total	Total number of entries in the write database.
Total	Total number of entries in the read database.

Related Commands

Command	Description
show ip rsvp high-availability counters	Displays all RSVP HA counters that are being maintained by an RP.
show ip rsvp high-availability summary	Displays summary information for an RSVP HA RP.

show ip rsvp host

To display specific information for a Resource Reservation Protocol (RSVP) host, use the **showiprsvphost** command in user EXEC or privileged EXEC mode.

show ip rsvp host {**receivers**| **senders**} [*hostname*| *group-address*]

Syntax Description

senders	RSVP-related sender information currently in the database.
receivers	RSVP-related receiver information currently in the database.
hostname	(Optional) Hostname of the source or destination.
group-address	(Optional) IP address of the source or destination.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(3)T	This command was introduced.
12.4(6)T	This command was modified. The command output was modified to display RSVP identity information when configured.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
Cisco IOS XE Release 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

Use the **showiprsvphost** command to display static RSVP senders and receivers. If a router has any local host receivers or senders that have RSVP identities configured, the application IDs that they use are also displayed.

Examples

In the following example from the**showiprsvphostsenders**command, no RSVP identities are configured for the local sender:

```
Router# show ip rsvp host senders
To            From          Pro DPort Sport Prev Hop      I/F      BPS
192.168.104.3 192.168.104.1 UDP 1     1                           10K
  Mode(s): Host CLI
```

The table below describes the significant fields shown in the display.

Table 57: show ip rsvp host senders (No RSVP Identities Configured) Field Descriptions

Field	Description
To	IP address of the receiver.
From	IP address of the sender.
Pro	Protocol code. IP protocol such as TCP or UDP.
DPort	Destination port number. Code 1 indicates an IP protocol such as TCP or UDP.
Sport	Source port number. Code 1 indicates an IP protocol such as TCP or UDP.
Prev Hop	IP address of the previous hop. Blank means no previous hop.
I/F	Interface of the previous hop.
BPS	Reservation rate, in bits per second (bps).
Mode(s)	Any of the following strings: • Host--The router is acting as the host system or RSVP endpoint for this reservation. • LSP-Tunnel--The reservation is for a traffic engineering (TE) tunnel. • MIB--The reservation was created via an Simple Network Management Protocol (SNMP) SET directive from a remote management station. • CLI--The reservation was created via a local RSVP command. • Host CLI--A combination of the host and command line interface (CLI) strings meaning that the static sender being displayed was created by the **iprsvpsender-host** command.

In the following example from the**showiprsvphostsenders**command, an RSVP identity is configured for the local sender:

```
Router# show ip rsvp host senders
To            From          Pro DPort Sport Prev Hop      I/F      BPS
192.168.104.3 192.168.104.1 UDP 1     1                             10K
  Mode(s): Host CLI
  Identity: voice100
    Locator: GUID=www.cisco.com,APP=voice,VER=100.0
    ID Type: Application
```

The table below describes the significant fields shown in the display.

Table 58: show ip rsvp host senders (RSVP Identity Configured) Field Descriptions

Field	Description
To	IP address of the receiver.
From	IP address of the sender.
Pro	Protocol code. IP protocol such as TCP or UDP.
DPort	Destination port number. Code 1 indicates IP protocol such as TCP or UDP.
Sport	Source port number. Code 1 indicates IP protocol such as TCP or UDP.
Prev Hop	IP address of the previous hop. Blank means no previous hop.
I/F	Interface of the previous hop.
BPS	Reservation rate in bits per second (bps).
Mode(s)	Any of the following strings: • CLI--The reservation was created via a local RSVP command. • Host--The router is acting as the host system or RSVP endpoint for this reservation. • Host CLI--A combination of the host and CLI strings meaning that the static sender being displayed was created by the **iprsvpsender-host** command. • LSP-Tunnel--The reservation is for a Traffic Engineering (TE) tunnel. • MIB--The reservation was created via an SNMP SET directive from a remote management station.
Identity	The alias string for the RSVP application ID.
Locator	The application ID that is being signaled in the RSVP PATH message for this statically-configured sender.

Field	Description
ID Type	Types of identities. RSVP defines two types: application IDs (Application) and user IDs (User). Cisco IOS software and Cisco IOS XE software support application IDs only.

Related Commands

Command	Description
ip rsvp sender-host	Enables a router to simulate a host generating an RSVP PATH message.

show ip rsvp interface detail

To display the hello configuration for all interface types, use the **show ip rsvp interface detail**command in user EXEC or privileged EXEC mode.

show ip rsvp interface detail [*type number*]

Syntax Description

type number	(Optional) The type and number of the interface for which you want to display the hello configuration.

Command Default

The hello configuration for all interfaces is displayed.

Command Modes

User EXEC (>) Privileged EXEC (#)

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.2(18)SXD1	This command was integrated into Cisco IOS Release 12.2(18)SXD1.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.2(33)SRC	This command was integrated into Cisco IOS Release 12.2(33)SRC.
12.4(20)T	This command was integrated into Cisco IOS Release 12.4(20)T.
12.2(33)SRE	This command was modified. The output was updated to display the source address used in the PHOP address field.
15.1(2)T	This command was modified. The output was updated to display the overhead percent.
15.1(1)S	This command was integrated into Cisco IOS Release 15.1(1)S.
15.2(2)SNG	This command was implemented on the Cisco ASR 901 Series Aggregation Services Routers.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Usage Guidelines

To display the hello configuration for a specific interface, use the **show ip rsvp interface detail** command with the *type* and *number* arguments.

Examples

The following is sample output from the **show ip rsvp interface detail** command:

```
Router# show ip rsvp interface detail GigabitEthernet 9/47
Tu0:
   RSVP: Enabled
   Interface State: Up
   Bandwidth:
     Curr allocated: 10K bits/sec
     Max. allowed (total): 75K bits/sec
     Max. allowed (per flow): 75K bits/sec
     Max. allowed for LSP tunnels using sub-pools: 0 bits/sec
     Set aside by policy (total): 0 bits/sec
   Admission Control:
     Header Compression methods supported:
       rtp (36 bytes-saved), udp (20 bytes-saved)
     Tunnel IP Overhead percent:
       4
     Tunnel Bandwidth considered:
       Yes
   Traffic Control:
     RSVP Data Packet Classification is ON via CEF callbacks
   Signalling:
     DSCP value used in RSVP msgs: 0x3F
     Number of refresh intervals to enforce blockade state: 4
   Authentication: disabled
     Key chain:   <none>
     Type:        md5
     Window size: 1
     Challenge:   disabled
   Hello Extension:
     State: Disabled
```

The table below describes the significant fields shown in the display.

Table 59: show ip rsvp interface detail Field Descriptions

Field	Description
RSVP	Status of the Resource Reservation Protocol (RSVP) (Enabled or Disabled).
Interface State	Status of the interface (Up or Down).
Curr allocated	Amount of bandwidth (in bits per second [b/s]) currently allocated.
Max. allowed (total)	Total maximum amount of bandwidth (in b/s) allowed.
Max. allowed (per flow)	Maximum amount of bandwidth (in b/s) allowed per flow.
Max. allowed for LSP tunnels using sub-pools	Maximum amount of bandwidth permitted for the label switched path (LSP) tunnels that obtain their bandwidth from subpools.

Field	Description
Tunnel IP Overhead percent	Overhead percent to override the RSVP bandwidth manually.
Tunnel Bandwidth considered	Indicates if the tunnel bandwidth is considered.
DSCP value used in RSVP msgs	Differentiated services code point (DSCP) value in the RSVP messages.

show ip traffic-engineering

To display information about the traffic engineering configuration and metric information associated with it, use the **show ip traffic-engineering** command in privileged EXEC mode.

show ip traffic-engineering [**metrics [detail]**]

Syntax Description

metrics	(Optional) Displays metric information associated with traffic engineering.
detail	(Optional) Displays information in long form.

Command Modes

Privileged EXEC

Command History

Release	Modification
11.1CT	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The goal of the loop prevention algorithm is that traffic should not be sent down the tunnel if there is a possibility that, after leaving the tunnel, steady state routing will route the traffic back to the head of the tunnel.

The strategy of the loop prevention algorithm is to compare the Layer 3 routing distance to the egress from the tunnel tailend and tunnel headend. The loop check passes only if the tunnel tail is closer to the egress than the tunnel head is.

The loop prevention algorithm allows you to use the tunnel for a route if one the following cases applies:

- Given that the two ends of the tunnel are routing to the egress using the same dynamic protocol in the same area, the Layer 3 routing distance from the tailend to the egress is less than the Layer 3 routing distance from the headend to the egress.
- The route to the egress is directly connected at the tunnel tailend router, but not at the tunnel headend router.
- The egress is unreachable from the tunnel headend router, but is reachable from the tunnel tailend router.

The loop prevention algorithm prevents you from using the tunnel for a given egress in all other cases, in particular, the following cases:

- The routers at the ends of the tunnel get their route to the egress from different dynamic routing protocols.
- The routing protocols at the two ends of the tunnel route to the egress through different areas.
- The two ends each use a static route to the egress.
- The tunnel headend router's route to the egress is a connected route.
- The egress is unreachable from the tunnel tailend router.

Devices request metrics via an LDP adjacency. The display output shows detailed metric information.

The metric information includes a metric type (shown as routing_protocol/routing_protocol_subtype) and a metric value.

The routing protocol is as follows:

- Open Shortest Path First (OSPF)
- Intermediate System to Intermediate System (IS-IS)
- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Connected
- Static
- Other (some other routing protocol)

The routing protocol subtype is specific to each routing protocol.

Examples

The following is sample output from the **show ip traffic-engineering metrics detail** command:

```
Router# show ip traffic-engineering metrics detail
Metrics requested BY this device
 Prefix 43.0.0.1/32
  TDP id 2.2.2.2:0, metric: connected/0
    type request, flags metric-received, rev 6, refcnt 1
  TDP id 4.4.4.4:0, metric: ospf-300/2
    type request, flags metric-received, rev 7, refcnt 1
 Prefix 44.0.0.0/8
  TDP id 18.18.18.18:0, metric: connected/0
    type request, flags metric-received, rev 1, refcnt 1
Metrics requested FROM this device
 Prefix 36.0.0.0/8
  TDP id 18.18.18.18:0, metric: connected/0
    type advertise, flags none, rev 1, refcnt 1
```

The table below describes the significant fields shown in the display.

Table 60: show ip traffic-engineering metrics detail Field Descriptions

Field	Description
Prefix	Destination network and mask.
TDP id	The LDP identifier of the LDP peer device at the other end of the tunnel. The LDP peer device advertises these metrics to this neighbor.

Field	Description
metric	The routing protocol and metric within that protocol for the prefix in question.
type	For metrics being requested by this device, the type is either “request” or “release.” For metrics being requested from this device, the type is “advertise.”
flags	For metrics being requested by this device, “metric-received” indicates that the other end has responded with a metric value. For metrics being requested from this device, response-pending indicates that the metric value has not yet been sent to the requester.
rev	An internal identifier for the metric request or advertisement. The rev number is assigned when the request/advertisement is created. The rev number is updated if the local information for the metric changes.
refcnt	For a metric of type request, the number of traffic engineering routes interested in this metric value. Otherwise, refcnt is 1.

Related Commands

Command	Description
traffic-engineering filter	Specifies a filter with a given number and properties.
traffic-engineering route	Configures a route for a specified filter, through a specified tunnel.

show ip traffic-engineering configuration

To display information about configured traffic engineering filters and routes, use the **show ip traffic-engineering configuration** command in privileged EXEC mode.

show ip traffic-engineering configuration [*interface*] [*filter-number*] **[detail]**

Syntax Description

interface	(Optional) Specifies an interface for which to display traffic engineering information.
filter-number	(Optional) A decimal value representing the number of the filter to display.
detail	(Optional) Displays command output in long form.

Command Modes

Privileged EXEC

Command History

Release	Modification
11.1CT	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The sample output can show all filters or can be limited by interface, filter number, or both.

Examples

The following is sample output from the **show ip traffic-engineering configuration detail** command:

```
Router# show ip traffic-engineering configuration detail
Traffic Engineering Configuration
    Filter 5: egress 44.0.0.0/8, local metric: ospf-0/1
        Tunnel5 route installed
          interface up, preference 1
          loop check on, passing, remote metric: connected/0
    Filter 6: egress 43.0.0.1/32, local metric: ospf-300/3
        Tunnel7 route installed
          interface up, preference 50
          loop check on, passing, remote metric: ospf-300/2
        Tunnel6 route not installed
          interface up, preference 75
          loop check on, passing, remote metric: connected/0
```

The table below describes the significant fields shown in the display.

Table 61: show ip traffic-engineering configuration detail Field Descriptions

Field	Description
Filter	The configured filter identifier for the traffic engineering route.
egress	The prefix/mask configured with the filter local metric.
local metric	The routing protocol and metric value of the local LSR for the egress prefix/mask.
Tunnel5	The tunnel for the traffic engineering route.
route installed/not installed	Indicates whether the route is installed in the forwarding tables (typically CEF and label interface up/down).
interface	Indicates whether the tunnel interface for the traffic engineering route is up or down. The traffic engineering route is not installed if the tunnel interface is down.
preference	The configured administrative preference for the traffic engineering route.
loop check	Indicates whether the loop check has been configured on or off.
passing/failing	If the loop check is configured on, indicates whether the check is passing. The traffic engineering route is not installed if the loop check is configured on and is failing.
remote metric	The routing protocol and the metric within that protocol for the prefix in question, as seen by the LSR that is advertising the metric. As part of the loop check, a comparison is made between the remote metric and the local metric.

Related Commands

Command	Description
show ip traffic-engineering routes	Displays information about the requested filters configured for traffic engineering.