sequencing

To configure the direction in which sequencing is enabled for data packets in a Layer 2 pseudowire, use the **sequencing** command in pseudowire class 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 the reset of packet sequencing after the disposition router receives a specified number of out-of-order packets.
number	The number of out-of-order packets that cause a reset of packet sequencing. The range is 5 to 65535.

Command Default

Sequencing is disabled.

Command Modes

Pseudowire class configuration

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	The **resync** keyword was added.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(27)SBC	L2TPv3 support for this command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(28)SB	AToM support for this command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

When you enable sequencing using any of the available options, the sending of sequence numbers is automatically enabled and the remote provider edge (PE) peer is requested to send sequence numbers. Out-of-order packets 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 you issue the **ip cef distributed** command, all traffic on the pseudowires is switched through the line cards.

It is useful to specify the **resync** keyword for situations when the disposition router receives many out-of-order packets. It allows the router to recover from situations where 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" so that the Sequence Number field is updated in tunneled packet headers for data packets that are both sent and received over the pseudowire:

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

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

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

Related Commands

Command	Description
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 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.

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.

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 extmap:

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

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

Defaults

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.

Defaults

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.

Release	Modification
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-nam*e 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 in conjunction 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.

Defaults

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.

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** keyword and the applicable *table-map-nam*e 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.
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)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.2(33)SB	This command was integrated into Cisco IOS Release 12.2(33)SB.

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.

Defaults

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.

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 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 IOS 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.
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.

Command	Description
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 IOS 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.

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
```

Table 6 describes the significant fields shown in the display.

Table 6 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 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.
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
```

Table 7 describes the fields shown in the displays.

Table 7 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.

Table 7 show atm vc Field Descriptions (continued)

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.

Table 7 show atm vc Field Descriptions (continued)

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.
Broadcasts	Number of process-switched broadcast packets.
InFast	Number of fast-switched input packets.

Table 7 show atm vc Field Descriptions (continued)

Field	Description
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.
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 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.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.4(2)T	The command output was changed to add Segment 1 and Segment 2 fields for Segment state and channel ID.
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.
12.2(33)SB	This command was updated to display HDLC local switching connections.
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.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.

Examples

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

```
Router# 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
```

Table 8 describes the significant fields shown in the display.

Table 8 show connection Field Descriptions

Field	Description
ID	Arbitrary connection identifier assigned by the operating system.
Name	Name of the connection.
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 controller vsi control-interface** command 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	Thsi 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.

Table 10 describes the significant fields shown in the display.

Table 9 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
```

Table 10 describes the significant fields shown in the display.

Table 10 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.
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*]]

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.

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.)

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
```

Table 11 describes the significant fields shown in the display.

Table 11 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
```

Table 12 describes the significant fields shown in the display.

Table 12 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.
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.

Table 12 show controllers vsi session Field Descriptions

Field	Description
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.
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 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
```

Table 13 describes the significant fields shown in the display.

Table 13 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.”

Table 13 show controllers vsi status Field Descriptions (continued)

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
```

Table 14 describes the significant fields shown in the display.

Table 14 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.
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
```

Table 15 describes the significant fields shown in the display.

Table 15 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.
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.

Table 15 show controllers vsi traffic session Field Descriptions (continued)

Field	Description
Conn Del Rsp	Number of VSI “delete connection response” messages received.
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
```

Table 16 describes the significant fields shown in the display.

Table 16 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.
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.

Table 16 show controllers xtagatm Field Descriptions (continued)

Field	Description
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.
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 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
```

Table 17 describes the significant fields shown in the display.

Table 17 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.

Table 17 show interface tunnel configuration Field Descriptions (continued)

Field	Description
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.

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 1/255, rxload 1/255
 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 1/255, rxload 1/255
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
```

Table 18 describes the significant fields shown in the displays.

Table 18 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.

Table 18 show interface xtagatm Field Descriptions

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.
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.

Table 18 show interface xtagatm Field Descriptions

Field	Description
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.
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.

Table 18 show interface xtagatm Field Descriptions

Field	Description
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.
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 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.

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
```

Table 19 describes the significant fields shown in the display.

Table 19 show ip bgp labels Field Descriptions

Field	Description
Network	Displays the network address from the eGBP 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 neighbors [*ip-address* [**advertised-routes | dampened-routes | flap-statistics | paths** [*reg-exp*] **| received prefix-filter | received-routes | routes | policy** [**detail**]]]

Syntax Description

ip-address	(Optional) IP address of a neighbor. If this argument is omitted, all neighbors are displayed.
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.
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.
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.

Command Default

The output of this command displays information for all neighbors.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
10.0	This command was introduced.
11.2	The **received-routes** keyword was added.
12.0(18)S	The output was modified to display the no-prepend configuration option and this command was integrated into Cisco IOS Release 12.0(18)S.
12.2(4)T	The **received** and **prefix-filter** keywords were added, and this command was integrated into Cisco IOS Release 12.2(4)T.

Release	Modification
12.0(21)ST	The output was modified to display Multiprotocol Label Switching (MPLS) label information.
12.0(22)S	Support for the BGP graceful restart capability was integrated into the output. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was also added.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(15)T	Support for the BGP graceful restart capability was integrated into the output.
12.0(25)S	The **policy** and **detail** keywords were added.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.0(27)S	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.3(7)T	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.0(31)S	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(18)SXE	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.4(4)T	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support BGP TCP path MTU discovery.
12.4(11)T	Support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.4(11)T.
12.2(33)SRB	Support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SXH	The output was modified to support BGP dynamic neighbors.

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.

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

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

In Cisco IOS Releases 12.0(25)S, 12.4(11)T, 12.2(33)SRB 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:

show ip bgp neighbors: Example

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.

```
Router# 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 and received
    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 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
```

Table 20 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 20 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 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 hhmmss, that the underlying TCP connection has been in existence.

Table 20 show ip bgp neighbors Field Descriptions (continued)

Field	Description
Last read	Time, in hhmmss, since BGP last received a message from this neighbor.
last write	Time, in hhmmss, since BGP last sent a message to this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving a 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.
Received	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.
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.
...update-group	Number of 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.

Table 20 show ip bgp neighbors Field Descriptions (continued)

Field	Description
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as bestpaths.
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.
* 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 AS 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound non-local 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 bestpath came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.

Table 20 show ip bgp neighbors Field Descriptions (continued)

Field	Description
* Bestpath from iBGP peer	Deploys inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a CE router.
* BGP distribute-list	Displays inbound denials due to a distribute list.
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 since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be... (not shown in the display)	Indicates that the BGP 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.
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 times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.
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.

Table 20 show ip bgp neighbors Field Descriptions (continued)

Field	Description
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 larger 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:	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Largest 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.
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.
out of order:	Number of packets received out of sequence.
with data	Number of update packets received with data.
Last reset	Elapsed time since this peering session was last reset.
unread input bytes	Number of bytes of packets still to be processed.
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 acknowledgements (transmissions before or without subsequent acknowledgements).
Second Congestion	Number of second retransmissions sent due to congestion.

show ip bgp neighbors advertised-routes: Example

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

```
Router# 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
```

Table 21 describes the significant fields shown in the display.

Table 21 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.
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 router 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.

Table 21 show ip bgp neighbors advertised-routes Field Descriptions (continued)

Field	Description
Metric	If shown, this is the value of the inter-autonomous 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.

show ip bgp neighbors paths: Example

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

```
Router# show ip bgp neighbors 172.29.232.178 paths ^10

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

Table 22 describes the significant fields shown in the display.

Table 22 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.

show ip bgp neighbors received prefix-filter: Example

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:

```
Router# 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
```

Table 23 describes the significant fields shown in the display.

Table 23 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.

show ip bgp neighbors policy: Example

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.

```
Router# 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
```

Cisco IOS Release 12.0(31)S and 12.4(4)T: Example

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.

```
Router# show ip bgp neighbors

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

Cisco IOS Release 12.2(33)SRA: Example

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.

```
Router# 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
```

Cisco IOS Release 12.2(33)SXH: Example

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.

```
Router# 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
.
.
.
```

Related Commands

Command	Description
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.

show ip bgp vpnv4

To display Virtual Private Network 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*} [**rib-failure**] [*ip-prefix/length* [**longer-prefixes**]] [*network-address* [*mask*] [**longer-prefixes**]] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**labels**]

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.
rib-failure	(Optional) Displays BGP routes that failed to install in the VRF table.
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 net masks.
community	(Optional) Displays routes that match this community.
community-list	(Optional) Displays routes that match this community list.
dampened-paths	(Optional) Displays paths suppressed because of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes that conform to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
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.
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.

summary	(Optional) Displays BGP neighbor status.
labels	(Optional) Displays incoming and outgoing BGP labels for each NLRI prefix.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	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	The **tags** keyword was replaced by the **labels** keyword to conform to the MPLS guidelines. This command was integrated into Cisco IOS Release 12.0(21)ST.
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	The output of the **show ip bgp vpnv4 all labels** command was enhanced to display explicit-null label information.
12.3	The **rib-failure** keyword was added for VRFs.
12.2(22)S	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 integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support per-VRF assignment of the BGP router ID.
12.2(31)SB2	The output was modified to support per-VRF assignment of the BGP router ID.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH, and 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	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 BGP Event-Based VPN Import.
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.

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 ?
```

Table 24 describes the significant fields shown in the display.

Table 24 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
```

Table 25 describes the significant fields shown in the display.

Table 25 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.

Table 25 show ip bgp vpnv4 rd labels Field Descriptions (continued)

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
```

Table 26 describes the significant fields shown in the display.

Table 26 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
```

Table 27 describes the significant fields shown in the display.

Table 27 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.
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
```

Table 28 describes the significant fields shown in the display.

Table 28 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.
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.

Active Route Processor

```
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)
```

Standby Route Processor

```
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
```

Table 29 describes the significant fields shown in the display.

Table 29 show ip bgp vpn4 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
```

Table 30 describes the significant fields shown in the display.

Table 30 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.

Table 30 show ip bgp vpnv4 all labels Field Descriptions (continued)

Field	Description
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 ?
```

Table 31 describes the significant fields shown in the display.

Table 31 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 this 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 bestpath 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 this 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 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 RTs 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 bestpath as 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
```

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.
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. Valid values are from 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-releated 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.

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
```

Table 32 describes the significant fields shown in the display.

Table 32 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
```

Table 33 describes the significant fields shown in the display.

Table 33 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
```

Table 34 describes the significant fields shown in the display.

Table 34 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.
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.

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
```

Table 35 describes the significant fields shown in the display.

Table 35 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.

Defaults

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
```

Table 36 describes the significant fields shown in the display.

Table 36 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.
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
```

Table 37 describes the significant fields shown in the display.

Table 37 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 the current state of the routing table, use the **show ip route** command in user EXEC or privileged EXEC mode.

show ip route [*ip-address* [*mask*] [**longer-prefixes**] | *protocol* [*process-id*] | **list** [*access-list-number* | *access-list-name*] | **static download**]

Syntax Description

ip-address	(Optional) Address about which routing information should be displayed.
mask	(Optional) Argument for a subnet mask.
longer-prefixes	(Optional) Specifies that only routes matching the *ip-address* and *mask* pair should be displayed.
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**, **hello**, **eigrp**, **isis**, **odr**, **ospf**, and **rip**.
process-id	(Optional) The number used to identify a process of the specified protocol.
list	(Optional) The list keyword is required to filter output by an access list name or number.
access-list-number	(Optional) Filters the displayed output from the routing table based on the specified access list name.
access-list-name	(Optional) Filters the displayed output from the routing table based on the specified access list number.
static	(Optional) All static routes.
download	(Optional) The route installed using the AAA route download function. This keyword is used only when AAA is configured.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
9.2	This command was introduced.
10.0	The “D—EIGRP, EX—EIGRP, N1—OSPF NSSA external type 1 route” and “N2—OSPF NSSA external type 2 route” codes were added to the command output.
10.3	The *process-id* argument was added.
11.0	The **longer-prefixes** keyword was added.
11.1	The “U—per-user static route” code was added to the command output.
11.2	The “o—on-demand routing” code was added to the command output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Release	Modification
11.3	The output from the **show ip route** *ip-address* command was enhanced to display the origination of an IP route in Intermediate System-to-Intermediate System (IS-IS) networks.
12.0(1)T	The "M—mobile" code was added to the command output.
12.0(3)T	The "P—periodic downloaded static route" code was added to the command output.
12.0(4)T	The "ia—IS-IS" code was added to the command output.
12.2(2)T	The output from the **show ip route** *ip-address* command was enhanced to display information on the multipaths to the specified network.
12.2(13)T	The *egp* and *igrp* arguments were removed because the exterior gateway protocol (EGP) and the Interior Gateway Routing Protocol (IGRP) are no longer available in Cisco IOS 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	The output was enhanced to display route tag information.
12.3(8)T	The 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(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.

Usage Guidelines

The **show ip route static download** command provides a way to display all dynamic static routes with name and distance information, including active and inactive ones. You can display all active dynamic static routes with both the **show ip route** and **show ip route static** commands after these active routes are added in the main routing table.

Examples

Routing Table Examples

The following examples show the standard routing tables displayed by the **show ip route** command. Use the codes displayed at the beginning of each report and the information in Table 38 to understand the type of route.

The following is sample output from the **show ip route** command when entered without an address:

```
Router# 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 is sample output that includes IS-IS Level 2 routes learned:

```
Router# 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 0.0.0.0, 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 using the **longer-prefixes** keyword. When the **longer-prefixes** keyword is included, the address and mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed.

In the following example, 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 to that result of 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Router# 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 examples display all downloaded static routes. A P designates which route was installed using AAA route download.

```
Router# 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 172.21.17.1 to network 0.0.0.0

        172.31.0.0/32 is subnetted, 1 subnets
P       172.31.229.41 is directly connected, Dialer1 20.0.0.0/24 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

Router# show ip route static

     172.27.4.0/8 is variably subnetted, 2 subnets, 2 masks
P       172.16.1.1/32 is directly connected, BRI0
P       172.27.4.0/8 [1/0] via 10.1.1.1, BRI0
S    172.31.0.0/16 [1/0] via 172.21.114.65, Ethernet0
S    10.0.0.0/8 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
     172.21.0.0/16 is variably subnetted, 5 subnets, 2 masks
S       172.21.114.201/32 is directly connected, BRI0
S       172.21.114.205/32 is directly connected, BRI0
S       172.21.114.174/32 is directly connected, BRI0
S       172.21.114.12/32 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
P    10.1.0.0/8 is directly connected, BRI0
P    10.2.2.0/8 is directly connected, BRI0
S*   0.0.0.0/0 [1/0] via 172.21.114.65, Ethernet0
S    172.29.0.0/16 [1/0] via 172.21.114.65, Ethernet0
```

The following example shows how to use the **show ip route static download** command to display all active and inactive routes installed using AAA route download:

```
Router# 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
```

Table 38 show ip route Field Descriptions

Field	Description
O	Indicates the protocol that derived the route. It can be one of the following values: R—Routing Information Protocol (RIP) derived O—Open Shortest Path First (OSPF) derived C—connected S—static B—Border Gateway Protocol (BGP) derived D—Enhanced Interior Gateway Routing Protocol (EIGRP) EX—EIGRP external i—IS-IS derived ia—IS-IS M—mobile P—periodic downloaded static route U—per-user static route o—on-demand routing
E2	Type of route. It can be one of the following values: *—Indicates the last path used when a packet was forwarded. It pertains only to the nonfast-switched packets. However, it does not indicate which path will be used next when forwarding a nonfast-switched packet, except when the paths are equal cost. IA—OSPF interarea route E1—OSPF external type 1 route E2—OSPF external type 2 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
10.110.0.0	Indicates the address of the remote network.
[160/5]	The first number in the 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 router to the remote network.

Table 38 show ip route Field Descriptions (continued)

Field	Description
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.

Specific Route Information

When you specify that you want information about a specific network displayed, more detailed statistics are shown. The following is sample output from the **show ip route** command when entered with the IP address 10.0.0.1:

```
Router# 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, it includes one of its own IP addresses to be used as the originator IP address. When other routers calculate IP routes, they can store the originator IP address with each route in the routing table.

The example above shows the output from the **show ip route** command when looking at 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 where a particular IP route has originated in your network. In the example the route to 10.0.0.1/32 was originated by a router with IP address 10.191.255.247.

Table 39 describes the significant fields shown when using the **show ip route** command with an IP address.

Table 39 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.
Tag	Integer that is used to implement the route.
type	Indicates the IS-IS route type (Level 1 or Level 2).
Redistributing via...	Indicates the redistribution protocol.
Last update from 10.191.255.251	Indicates the IP address of a router that is the next hop to the remote network and the router 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	Number of uses for this routing descriptor block.

The following is sample output using the **longer-prefixes** keyword. When the **longer-prefixes** keyword is included, the address and mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed.

In the following example, 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 to that result of 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Router# 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 output includes the tag 120 applied to the route 10.22.0.0/16. You must specify an IP prefix in order to see the tag value.

```
Router# 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
```

Static Routes Using a DHCP Gateway 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 172.2.2.2 [1/0] are static, and route 10.0.0.0/0 is a default route candidate.

```
Router# 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 172.2.2.2 [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
```

Related Commands

Command	Description
show dialer	Displays general diagnostic information for interfaces configured for DDR.
show interfaces tunnel	Displays a list of 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 Virtual Private Network (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*] [*tag*] [*output-modifiers*]] [*ip-prefix*] [**list** *number* [*output-modifiers*]] [**profile**] [**static** [*output-modifiers*]] [**summary** [*output-modifiers*]] [**supernets-only** [*output-modifiers*]]

Syntax Description

vrf-name	Name assigned to the VRF.
connected	(Optional) Displays all connected routes in a VRF.
protocol	(Optional) 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.
tag	(Optional) Cisco IOS routing area label.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
ip-prefix	(Optional) Specifies a network to display.
list *number*	(Optional) Specifies the IP access list to display.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
supernets-only	(Optional) Displays supernet entries only.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The *ip-prefix* argument was added. The output from the **show ip route vrf** *vrf-name ip-prefix* command was enhanced to display information on the 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	Enhanced Interior Gateway Routing Protocol (EIGRP) VRF support was added.
12.2(15)T	EIGRP VRF support was integrated into Cisco IOS Release 12.2(15)T.
12.2(18)S	EIGRP VRF support 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.
12.2(33)SXH	The output was enhanced to display remote label information and corresponding MPLS flags for prefixes that have remote labels stored in the Routing Information Base (RIB).

Usage Guidelines

This command displays specified information from the IP routing table of a VRF.

Examples

This example shows the IP routing table associated with the VRF named vrf1:

```
Router# 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 example shows BGP entries in the IP routing table associated with the VRF named vrf1:

```
Router# 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
```

This example shows the IP routing table associated with a VRF named PATH and network 10.22.22.0:

```
Router# 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
```

Table 40 describes the significant fields shown when the **show ip route vrf** *vrf-name ip-prefix* command is used.

Table 40 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	The metric to reach the destination network.
Tag	Integer that is used to implement the route.
type	Indicates that the route is an L1 type or L2 type route.
Last update from 10.22.5.10	Indicates the IP address of a router that is the next hop to the remote network and the router interface on which the last update arrived.
00:01:07 ago	Specifies the last time the route was updated (in hours:minutes:seconds).
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	Number of uses for this routing descriptor block.
AS Hops	Number of hops to the destination or to the router where the route first enters internal BGP (iBGP).

Example of Output Using the Cisco IOS Software Modularity for Layer 3 VPNs Feature

The following is sample output from the **show ip route vrf** command on routers 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:

```
Router# 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
```

Table 41 describes the significant fields shown in the display.

Table 41 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 does not display if any of the following conditions is true: • BGP is not the LDP. However, OSPF prefixes learned via sham link display an MPLS label. • MPLS is not supported. • The prefix was imported from another VRF, where the prefix was an IGP prefix and LDP provided the remote label for it.
MPLS Flags	The name of one of the following MPLS flags is displayed if any is set: • MPLS Required—Packets are forwarded to this prefix because the MPLS label stack is present. If MPLS is disabled in the outgoing interface, the packets are dropped. • No Global—MPLS packets for this prefix are forwarded from the VRF interface, not from the interface in global table. Using the VRF interface prevents loops in scenarios that use ieBGP multipath. • NSF—The prefix is from an NSF-aware neighbor. If the routing information temporarily disappears due to a disruption in the control plane, packets for this prefix are preserved.

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 **show ip rsvp fast bw-protect** command in user EXEC or privileged EXEC mode.

show ip rsvp fast bw-protect

Syntax Description

This command has no arguments or keywords.

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.

Examples

The following is sample output from the **show ip rsvp fast bw-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
```

Table 42 describes the significant fields shown in the display.

Table 42 show ip rsvp fast bw-protect Field Descriptions

Field	Description
Primary Tunnel	Identification of the tunnel being protected.
Protect I/F	Interface name.

Table 42 show ip rsvp fast bw-protect Field Descriptions (continued)

Field	Description
BW BPS:Type	Bandwidth, in bits per second, and type of bandwidth. Possible values are: • S—Subpool • G—Global pool
Backup Tunnel:Label	Identification of the backup tunnel.
State	Status of backup tunnel. Valid values are: • 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.
Type	Type of backup tunnel. Possible values are: • 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 **show ip rsvp fast detail** command in user EXEC or privileged EXEC mode.

show ip rsvp fast detail

Syntax Description

This command has no arguments or keywords.

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.

Examples

The following is sample output from the **show ip rsvp fast detail** 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)
    555.5.6.6 (Strict IPv4 Prefix, 8 bytes, /32)
    555.6.7.7 (Strict IPv4 Prefix, 8 bytes, /32)
    10.0.0.7 (Strict IPv4 Prefix, 8 bytes, /32)
  ERO: (outgoing)
    555.5.6.6 (Strict IPv4 Prefix, 8 bytes, /32)
    555.6.7.7 (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: 555.5.6.5  LSP ID: 8
```

```
    Bkup FilerSpec:
      Tun Sender: 555.5.6.5, 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
```

Table 43 describes the significant fields shown in the display.

Table 43 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).
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.

Table 43 show ip rsvp fast detail Field Descriptions (continued)

Field	Description
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).
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: • 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 **show ip rsvp hello** 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.

Examples

The following is sample output from the **show ip rsvp hello** 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
```

Table 44 describes the significant fields shown in the display. The fields describe the processes for which hello is enabled or disabled.

Table 44 show ip rsvp hello Field Descriptions

Field	Description
RSVP Hello for Fast-Reroute/Reroute	Status of Fast-Reroute/Reroute: • Enabled—Fast reroute and reroute (hello for state timer) are activated (enabled). • Disabled—Fast reroute and reroute (hello for state timer) are not activated (disabled).
Statistics	Status of hello statistics: • 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. • Disabled—Hello statistics are not configured. • 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: • Enabled—BFD is configured. • Disabled—BFD is not configured.
Graceful Restart	Restart capability: • Enabled—Restart capability is activated for a router (full mode) or its neighbor (help-neighbor). • Disabled—Restart capability is not activated.

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.

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
```

Table 45 describes the significant fields shown in the display.

Table 45 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).
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.

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
```

Table 46 describes the significant fields shown in the display.

Table 46 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).
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 EXE C

Command History

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

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
```

Table 47 describes the significant fields shown in the display.

Table 47 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.
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 **show ip rsvp hello instance detail** 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.

Usage Guidelines

Use the **show ip rsvp hello instance detail** command to display information about the processes (clients) currently configured.

Examples

The following is sample output from the **show ip rsvp hello instance detail** 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
```

Table 48 describes the significant fields shown in the display.

Table 48 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).
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.

Table 48 show ip rsvp hello instance detail Field Descriptions (continued)

Field	Description
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.
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 **show ip rsvp hello instance summary** 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.

Examples

The following is sample output from the **show ip rsvp hello instance summary** 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
```

Table 49 describes the significant fields shown in the display.

Table 49 show ip rsvp hello instance summary Field Descriptions

Field	Description
Active Instances	Active nodes that are sending hello requests.
Client	Clients on behalf of which hellos are sent; they include GR (graceful restart), RR (reroute = hello state timer), and FRR (Fast Reroute).

Table 49 show ip rsvp hello instance summary Field Descriptions (continued)

Field	Description
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.
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 **show ip rsvp hello statistics** 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.

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 **show ip rsvp hello statistics** 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
```

Table 50 describes the significant fields shown in the display.

Table 50 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.

Table 50 show ip rsvp hello statistics Field Descriptions (continued)

Field	Description
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** | **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 the 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 results of a loose hop expansion performed on the router.

Release	Modification
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 of point-to-point (P2P) tunnels from that of 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.

Usage Guidelines

Use the **show ip rsvp high-availability database** command to display information about entries in 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 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

Hello Example on an Active RP

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
```

Table 51 describes the significant fields shown in the display.

Table 51 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.
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 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 acks and negative acknowledgments (nacks) to messages sent.
Flags	Attribute used to identify or track data.
Data	Information about the last transmission.
Last sent Src_instance	Last source instance identifier sent.
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.

Hello Example on Standby RP

The following is sample output from the **show ip rsvp high-availability database hello** command 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 Table 51 except they are now in the read database for the standby RP.

Link-Management Interfaces Example on an Active RP

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
```

Table 52 describes the significant fields shown in the display.

Table 52 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 messages sent.
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.

Table 52 show ip rsvp high-availability database link-management interfaces—Active RP Field Descriptions (continued)

Field	Description
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.
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 Table 52 except they are now in the TE link read database instead of the TE link write database that is used by an active RP.

Link-Management System Example on an Active RP

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
```

Table 53 describes the significant fields shown in the display.

Table 53 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.
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.

Table 53 show ip rsvp high-availability database link-management system—Active RP Field Descriptions (continued)

Field	Description
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.
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 Table 53 except they are now in the TE system read database instead of the TE system write database that is used by an active RP.

LSP Example on an Active RP for a P2P Tunnel

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)
```

LSP Example on an Active RP for a P2MP Tunnel

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)
```

Table 54 describes the significant fields shown in the display.

Table 54 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.
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.

Table 54 show ip rsvp high-availability database lsp—Active RP Field Descriptions (continued)

Field	Description
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 Table 54 except they are now in the LSP read database instead of the LSP write database that is used by an active RP.

LSP-Head Example on an Active RP for a P2P Tunnel

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
```

LSP-Head Example on an Active RP for a P2MP Tunnel

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
```

Table 55 describes the significant fields shown in the display.

Table 55 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.
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.
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.

Table 55 show ip rsvp high-availability database lsp-head—Active RP Field Descriptions

Field	Description
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.
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 Table 55 except they are now in the LSP_head read database instead of the LSP_head write database that is used by an active RP.

Summary Example on an Active RP

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
```

Table 56 describes the significant fields shown in the display.

Table 56 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.
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.

Summary Example on a Standby RP

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
```

Table 57 describes the significant fields shown in the display.

Table 57 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 **show ip rsvp host** command in user EXEC or privileged EXEC mode.

show ip rsvp host {**senders** | **receivers**} [*group-name* | *group-address*]

Syntax Description

senders	RSVP-related sender information currently in the database.
receivers	RSVP-related receiver information currently in the database.
group-name	(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	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.

Usage Guidelines

Use the **show ip rsvp host** 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 **show ip rsvp host senders** 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
```

Table 58 describes the significant fields shown in the display.

Table 58 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.

Table 58 show ip rsvp host senders (No RSVP Identities Configured) Field Descriptions

Field	Description
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: • 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 SNMP SET directive from a remote management station. • CLI—The reservation was created via a local RSVP CLI command. • Host CLI—A combination of the host and CLI strings meaning that the static sender being displayed was created by the **ip rsvp sender-host** CLI command.

In the following example from the **show ip rsvp host senders** command, an RSVP identity is configured for the local sender and more information displays:

```
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
```

Table 59 describes the significant fields shown in the display.

Table 59 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.

Table 59 show ip rsvp host senders (RSVP Identity Configured) Field Descriptions (continued)

Field	Description
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 SNMP SET directive from a remote management station. • CLI—The reservation was created via a local RSVP CLI command. • Host CLI—A combination of the host and CLI strings meaning that the static sender being displayed was created by the **ip rsvp sender-host** CLI command.
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.
ID Type	Types of identities. RSVP defines two types: application IDs (Application) and user IDs (User). Cisco IOS software currently supports Application 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 interface configuration for Hello, use the **show ip rsvp interface detail** command in privileged EXEC mode.

show ip rsvp interface detail [*interface*]

Syntax Description

interface	(Optional) Interface for which you want to show the Hello configuration.

Command Default

The interface configuration for Hello 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.
12.2(33)SRC	This command was integrated into Cisco IOS Release 12.2(33)SRC.

Examples

The following is sample output from the **show ip rsvp interface detail** command:

```
Router# show ip rsvp interface detail GigabitEthernet 9/47

Gi9/47:
 RSVP: Enabled
 Interface State: Up
 Bandwidth:
  Curr allocated: 0 bits/sec
  Max. allowed (total): 0 bits/sec
  Max. allowed (per flow): 0 bits/sec
  Max. allowed for LSP tunnels using sub-pools (pool 1): 0 bits/sec
  Set aside by policy (total): 0 bits/sec
 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
 FRR Extension:
  Backup Path: Configured (or "Not Configured")
 BFD Extension:
  State: Disabled
  Interval: Not Configured
 RSVP Hello Extension:
  State: Disabled
```

```
Refresh Interval: FRR: 200  , Reroute: 2000
Missed Acks:      FRR: 4    , Reroute: 4
DSCP in HELLOs:   FRR: 0x30 , Reroute: 0x30
```

Table 60 describes the significant fields shown in the display.

Table 60 show ip rsvp interface detail Field Descriptions

Field	Description
RSVP	Status of the Resource Reservation Protocol (RSVP) protocol (Enabled or Disabled).
Interface State	Status of the interface (Up or Down).
Curr allocated	Amount of bandwidth (in bits per second [bps]) currently allocated.
Max. allowed (total)	Total maximum amount of bandwidth (in bps) allowed.
Max. allowed (per flow)	Maximum amount of bandwidth (in bps) allowed per flow.
Max. allowed for LSP tunnels using sub-pools	Maximum amount of bandwidth permitted for label-switched path (LSP) tunnels that obtain their bandwidth from subpools.
DSCP value used in RSVP msgs	The differentiated services code point (DSCP) value that is in RSVP messages.
BFD Extension State	State (Enabled or Disabled) of BFD extension.
RSVP Hello Extension State	State (Enabled or Disabled) of Hello extension.
Missed Acks	Number of sequential acknowledgments that the node did not receive.
DSCP in HELLOs	The DSCP value that is in hello messages.

Related Commands

Command	Description
ip rsvp signalling hello (interface)	Enables Hello on an interface where you need Fast Reroute protection.
ip rsvp signalling hello dscp	Sets the DSCP value that is in the IP header of the hello message sent out from an interface.
ip rsvp signalling hello refresh interval	Configures the Hello request interval.

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
```

Table 61 describes the significant fields shown in the display.

Table 61 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.
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."

Table 61 show ip traffic-engineering metrics detail Field Descriptions (continued)

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
```

Table 62 describes the significant fields shown in the display.

Table 62 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.