show interface XTagATM

To display information about an extended MPLS ATM interface, use the **show interface XTagATM** EXEC command.

show interface XTagATM *if-num*

Syntax Description

if-num	Specifies the MPLS ATM interface number.

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

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

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 output from the **show interface XTagATM** command:

```
Router# show interface XTagATM0

XTagATM0 is up, line protocol is up
  Hardware is Tag-Controlled Switch Port
  Interface is unnumbered.  Using address of Loopback0 (12.0.0.17)
  MTU 4470 bytes, BW 156250 Kbit, DLY 80 usec, rely 255/255, load 1/255
  Encapsulation ATM Tagswitching, loopback not set
  Encapsulation(s): AAL5
  Control interface: ATM1/0, switch port: bpx 10.2
  9 terminating VCs, 16 switch cross-connects
  Switch port traffic:
     129302 cells input, 127559 cells output
  Last input 00:00:04, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/0, 0 drops; input queue 0/75, 0 drops
  Terminating traffic:
  5 minute input rate 1000 bits/sec, 1 packets/sec
  5 minute output rate 0 bits/sec, 1 packets/sec
     61643 packets input, 4571695 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     53799 packets output, 4079127 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffers copied, 0 interrupts, 0 failures
```

Table 20 describes the significant fields in the sample command output shown above.

Table 20 show interface XTagATM Field Descriptions

Field	Description
XTagATM0 is up	Interface is currently active.
line protocol is up	Displays the line protocol as up.
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.
rely	Reliability of the interface as a fraction of 255 (255/255 is 100% reliability), calculated as an exponential average over 5 minutes.
load	Load on the interface as a fraction of 255 (255/255 is completely saturated), calculated as an exponential average over 5 minutes.
Encapsulation ATM Tagswitching	Encapsulation method.
loopback not set	Indicates that loopback is not set.
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.
9 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.
16 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 counts	Indicates that counters below this line apply only to packets sent or received on terminating VCs.
5-minute input rate, 5-minute output rate	Average number of bits and packets sent per second in the last 5 minutes.
packets input	Total number of error-free packets received by the system.
bytes	Total number of bytes, including data and MAC encapsulation, in the error-free packets received by the system.

Table 20 show interface XTagATM Field Descriptions (continued)

Field	Description
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.
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.

Table 20 show interface XTagATM Field Descriptions (continued)

Field	Description
collisions	Number of messages re-sent due to an Ethernet collision. This is usually the result of an overextended LAN (Ethernet or transceiver cable too long, more than two repeaters between stations, or too many cascaded multiport transceivers). A packet that collides is counted only one time in output packets.
interface resets	Number of times an interface has been completely reset. Resets occur if packets queued for transmission were not sent within several seconds. On a serial line, this can be caused by a malfunctioning modem that is not supplying the transmit clock signal, or by a cable problem. If the system notices that the carrier detect line of a serial interface is up, but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.
output buffers copied	Number of packets copied from a MEMD buffer into a system buffer before being placed on the output hold queue.
interrupts	Displays the value of hwidb to tx_restarts.
failures	Number of packets discarded because no MEMD buffer was available.

Related Commands

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

show ip bgp vpnv4

To display VPN address information from the BGP table, use the **show ip bgp vpnv4** command in EXEC mode.

show ip bgp vpnv4 {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} [*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]] [*network-address* [*mask*] [**longer-prefixes**] [*output-modifiers*]] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**][**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**tags**]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays NLRIs that have a matching route distinguisher.
vrf *vrf-name*	Displays NLRIs associated with the named VRF.
ip-prefix/length	(Optional) IP prefix address (in dotted decimal format) and length of mask (0 to 32).
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.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
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 nonnatural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed on account of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes conforming 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 AS paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the AS path "regular expression."
regexp	(Optional) Displays routes matching the AS path regular expression.
summary	(Optional) Displays BGP neighbor status.
tags	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

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

Usage Guidelines

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

Examples

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

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 14.14.14.14
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: 100:1 (vrf1)
*> 11.0.0.0         50.0.0.1 0 0 101 i
*>i12.0.0.0         13.13.13.13 0    100 0 102 i
*> 50.0.0.0         50.0.0.1 0 0 101 i
*>i51.0.0.0         13.13.13.13 0    100 0 102 i
```

Table 21 describes the significant fields shown in the output.

Table 21 show ip bgp vpnv4 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 NLRIs that have a route-distinguisher value of 100:1.

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

NetworkNext Hop      In tag/Out tag
Route Distinguisher: 100:1 (vrf1)
   2.0.0.0          10.20.0.60      34/notag
   10.0.0.0         10.20.0.60      35/notag
   12.0.0.0         10.20.0.60      26/notag
```

```
                  10.20.0.60      26/notag
   13.0.0.0       10.15.0.15      notag/26
```

Table 22 describes the significant fields shown in the output.

Table 22 show ip bgp vpnv4 rd tags Field Descriptions

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

The following example shows VPNv4 routing entries for the VRF called vrf1.

```
Router# show ip bgp vpnv4 vrf vrf1

BGP table version is 18, local router ID is 14.14.14.14
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: 100:1 (vrf1)
*> 11.0.0.0         50.0.0.1 0 0 101 i
*>i12.0.0.0         13.13.13.13 0    100 0 102 i
*> 50.0.0.0         50.0.0.1 0 0 101 i
*>i51.0.0.0         13.13.13.13 0    100 0 102 i
```

Table 23 describes the significant fields shown in the output.

Table 23 show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays network address from the BGP table.
Next Hop	Displays 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.

Related Commands

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

show ip cache

To display the routing table cache used to fast switch IP traffic, use the **show ip cache** EXEC command.

show ip cache [*prefix mask*] [*type number*]

Syntax Description

prefix	(Optional) Displays only the entries in the cache that match the prefix and mask combination.
mask	(Optional) Displays only the entries in the cache that match the prefix and mask combination.
type	(Optional) Displays only the entries in the cache that match the interface type and number combination.
number	(Optional) Displays only the entries in the cache that match the interface type and number combination.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.

Usage Guidelines

The **show ip cache** display shows MAC headers up to 92 bytes.

Examples

The following is sample output from the **show ip cache** command:

```
Router# show ip cache

IP routing cache version 4490, 141 entries, 20772 bytes, 0 hash overflows
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last 7 seconds, 0 in last 3 seconds
Last full cache invalidation occurred 0:06:31 ago

Prefix/Length       Age       Interface       MAC Header
131.108.1.1/32      0:01:09   Ethernet0/0     AA000400013400000C0357430800
131.108.1.7/32      0:04:32   Ethernet0/0     00000C01281200000C0357430800
131.108.1.12/32     0:02:53   Ethernet0/0     00000C029FD000000C0357430800
131.108.2.13/32     0:06:22   Fddi2/0         00000C05A3E000000C035753AAAA0300
                                              00000800
131.108.2.160/32    0:06:12   Fddi2/0         00000C05A3E000000C035753AAAA0300
                                              00000800
131.108.3.0/24      0:00:21   Ethernet1/2     00000C026BC600000C03574D0800
131.108.4.0/24      0:02:00   Ethernet1/2     00000C026BC600000C03574D0800
131.108.5.0/24      0:00:00   Ethernet1/2     00000C04520800000C03574D0800
131.108.10.15/32    0:05:17   Ethernet0/2     00000C025FF500000C0357450800
131.108.11.7/32     0:04:08   Ethernet1/2     00000C010E3A00000C03574D0800
131.108.11.12/32    0:05:10   Ethernet0/0     00000C01281200000C0357430800
131.108.11.57/32    0:06:29   Ethernet0/0     00000C01281200000C0357430800
```

Table 24 describes the significant fields shown in the output.

Table 24 show ip cache Field Descriptions

Field	Description
IP routing cache version	Version number of this table. This number is incremented any time the table is flushed.
entries	Number of valid entries.
bytes	Number of bytes of processor memory for valid entries.
hash overflows	Number of times autonomous switching cache overflowed.
Minimum invalidation interval	Minimum time delay between cache invalidation request and actual invalidation.
maximum interval	Maximum time delay between cache invalidation request and actual invalidation.
quiet interval	Length of time between cache flush requests before the cache will be flushed.
threshold <*n*> requests	Maximum number of requests that can occur while the cache is considered quiet.
Invalidation rate <*n*> in last <*m*> seconds	Number of cache invalidations during the last <*m*> seconds.
0 in last 3 seconds	Number of cache invalidation requests during the last quiet interval.
Last full cache invalidation occurred <*hh:mm:ss*> ago	Time since last full cache invalidation was performed.
Prefix/Length	Network reachability information for cache entry.
Age	Age of cache entry.
Interface	Output interface type and number.
MAC Header	Layer 2 encapsulation information for cache entry.

The following is sample output from the **show ip cache** command with a prefix and mask specified:

```
Router# show ip cache 131.108.5.0 255.255.255.0

IP routing cache version 4490, 119 entries, 17464 bytes, 0 hash overflows
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last second, 0 in last 3 seconds
Last full cache invalidation occurred 0:11:56 ago

Prefix/Length       Age       Interface       MAC Header
131.108.5.0/24      0:00:34   Ethernet1/2     00000C04520800000C03574D0800
```

The following is sample output from the **show ip cache** command with an interface specified:

```
Router# show ip cache e0/2

IP routing cache version 4490, 141 entries, 20772 bytes, 0 hash overflows
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last second, 0 in last 3 seconds
Last full cache invalidation occurred 0:06:31 ago

Prefix/Length       Age       Interface       MAC Header
131.108.10.15/32    0:05:17   Ethernet0/2     00000C025FF500000C0357450800
```

show ip cache flow

To display a summary of the NetFlow switching statistics, use the **show ip cache flow** command in EXEC mode.

show ip cache [*prefix mask*] [*type number*] [**verbose**] **flow**

Syntax Description

prefix mask	(Optional) Displays only the entries in the cache that match the prefix and mask combination.
type number	(Optional) Displays only the entries in the cache that match the interface type and number combination.
verbose	(Optional) Displays additional information

Command Modes

EXEC

Command History

Release	Modification
11.1	This command was introduced.
11.1 CA	The information display for the command was updated.

Usage Guidelines

Some of the content in the display of the **show ip cache flow** command uses multiline headings and multiline data fields. Figure 2 shows how to associate the headings with the correct data fields when there are two lines of headings and two lines of data fields. The first line of the headings is associated with the first line of data fields. The second line of the headings is associated with the second line of data fields.

When other features are configured, the number of lines in the headings and data fields increases. The method for associating the headings with the correct data fields remains the same.

Figure 2 How to Use the Multiline Headings and Multiline Data Fields in the Display Output of the show ip cache flow Command

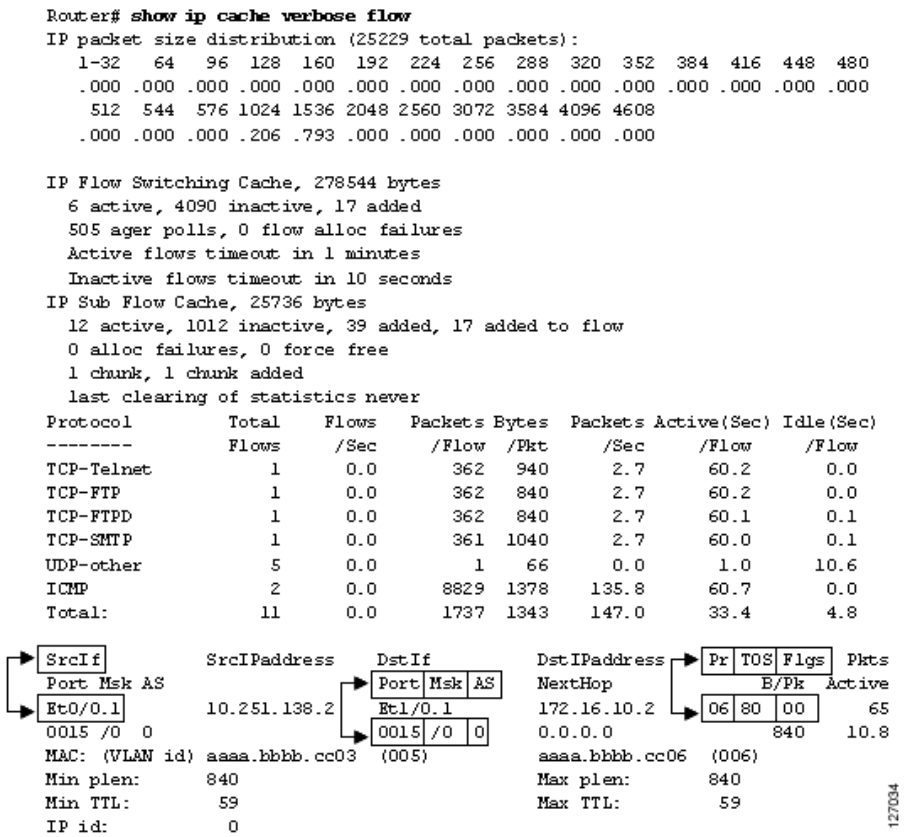

Displaying NetFlow Cache Information on a Distributed Cisco 7500 Series Platform

To display NetFlow cache information using the **show ip cache flow** command on a Cisco 7500 series router that is running dCEF, enter the following sequence of commands:

```
Router# if-con slot-number
LC-slot-number# show ip cache [prefix mask] [type number] [verbose] flow
```

Displaying NetFlow Cache Information on a Distributed Cisco 12000 Series Platform

To display NetFlow cache information using the **show ip cache flow** command on a Cisco 12000 Series Internet router, you enter the following sequence of commands:

```
Router# attach slot-number
LC-slot-number# show ip cache [prefix mask] [type number] [verbose] flow
```

Examples

The following is an example display of a main cache using the **show ip cache flow** command:

```
Router# show ip cache flow

IP packet size distribution (230151 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .999 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

   512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
```

```
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

The output above shows the percentage distribution of packets by size range. In this display, 99.9 percent of the packets fall in the size range from 1 to 32 bytes.

```
IP Flow Switching Cache, 4456448 bytes
  65509 active, 27 inactive, 820628747 added
  955454490 ager polls, 0 flow alloc failures
  Exporting flows to 1.1.15.1 (2057)
  820563238 flows exported in 34485239 udp datagrams, 0 failed
  last clearing of statistics 00:00:03

Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-BGP             71      0.0         1    49      0.0       2.5      15.8
UDP-other           17      0.0         1   328      0.0       0.0      15.7
ICMP             18966      6.7        10    28     72.9       0.1      22.9
Total:           19054      6.7        10    28     72.9       0.1      22.9

SrcIf          SrcIPaddress    DstIf          DstIPaddress    Pr TOS Flgs  Pkts
Port Msk AS                    Port Msk AS    NextHop              B/Pk  Active
Et1/1          52.52.52.1      Fd4/0          42.42.42.1      01 55  10    3748
0000 /8  50                    0000 /8  40    202.120.130.2          28    17.8
Et1/2          52.52.52.1      Fd4/0          42.42.42.1      01 CC  10    3568
0000 /8  50                    0000 /8  40    202.120.130.2          28    17.8
Et1/2          10.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1124
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          11.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1157
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.7
Et1/2          14.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1149
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          15.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1127
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.7
Et1/2          12.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1204
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          13.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1159
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          18.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1223
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          19.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1264
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          16.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1170
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          17.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1167
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          22.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1193
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.8
Et1/2          23.1.3.2        Fd4/0          42.42.42.1      01 C0  10    1212
0000 /0  0                     0000 /8  40    202.120.130.2          28    17.7
Et1/1          50.50.50.1      Local          31.31.31.1      06 C0  18       2
00B3 /32 0                     2AF8 /32 0     0.0.0.0                49    10.1
```

The following shows sample output from the **show ip cache** *prefix mask* **flow** command:

```
Router# show ip cache 10.0.0.1 256.0.0.0 flow

IP packet size distribution (25 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

The output above shows the percentage distribution of packets by size range. In this display, 100 percent of the packets fall in the128 byte range.

```
IP Flow Switching Cache, 4456704 bytes
  1 active, 65535 inactive, 5 added
  68 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
  last clearing of statistics never

Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
ICMP                 4      0.0         5   100      0.0       0.0      15.2
Total:               4      0.0         5   100      0.0       0.0      15.2

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Et1/2         10.0.0.2        Local         10.0.0.1        01 0000 0800     5
```

The following shows sample output from the **show ip cache** *type number* **flow** command:

```
Router# show ip cache e1/2 flow

IP packet size distribution (30 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456704 bytes
  1 active, 65535 inactive, 6 added
  85 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
  last clearing of statistics never

Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
ICMP                 5      0.0         5   100      0.0       0.0      15.1
Total:               5      0.0         5   100      0.0       0.0      15.1

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Et1/2         10.0.0.2        Local         10.0.0.1        01 0000 0800     5
```

Table 25 describes the significant fields shown in the flow switching cache lines of the display.

Table 25 show ip cache flow Field Descriptions in Flow Switching Cache Display

Field	Description
bytes	Number of bytes of memory used by the NetFlow cache.
active	Number of active flows in the NetFlow cache at the time this command was entered.
inactive	Number of flow buffers that are allocated in the NetFlow cache, but are not currently assigned to a specific flow at the time this command is entered.
added	Number of flows created since the start of the summary period.
ager polls	Number of times the NetFlow code looked at the cache to cause entries to expire (used by Cisco for diagnostics only).

Table 25 show ip cache flow Field Descriptions in Flow Switching Cache Display (continued)

Field	Description
flow alloc failures	Number of times the NetFlow code tried to allocate a flow but could not.
Exporting flows	IP address and User Datagram Protocol (UDP) port number of the workstation to which flows are exported.
flows exported in udp datagrams	Total number of flows exported and the total number of UDP datagrams used to export the flows to the workstation.
failed	Number of flows that could not be exported by the router because of output interface limitations.
last clearing of statistics	Standard time output (hh:mm:ss) since the **clear ip flow stats** EXEC command was executed. This time output changes to hours and days after the time exceeds 24 hours.

Table 26 describes the significant fields shown in the activity by protocol lines of the display.

Table 26 show ip cache flow Field Descriptions in Activity By Protocol Display

Field	Description
Protocol	IP protocol and the “well known” port number as described in RFC 1340.
Total Flows	Number of flows for this protocol since the last time statistics were cleared.
Flows/Sec	Average number of flows for this protocol seen per second; equal to total flows/number of seconds for this summary period.
Packets/Flow	Average number of packets observed for the flows seen for this protocol. Equal to total packets for this protocol or number of flows for this protocol for this summary period.
Bytes/Pkt	Average number of bytes observed for the packets seen for this protocol (total bytes for this protocol or the total number of packet for this protocol for this summary period).
Packets/Sec	Average number of packets for this protocol per second (total packets for this protocol) or the total number of seconds for this summary period.
Active(Sec)/Flow	Sum of all the seconds from the first packet to the last packet of an expired flow (for example, TCP FIN, timeout, and so on) in seconds or total flows for this protocol for this summary period.
Idle(Sec)/Flow	Sum of all the seconds from the last packet seen in each nonexpired flow for this protocol until the time at which this command was entered, in seconds or total flows for this protocol for this summary period.

The following sample output is for the **show ip cache flow** command when the **tunnel flow egress-records** command enables a generic routing encapsulation (GRE) tunnel with both Cisco Express Forwarding (CEF) and NetFlow configured. The last line is a NetFlow record that is created for packets that are encapsulated by a tunnel.

```
SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Se3/2         10.1.0.1        Tu0            40.1.1.1       01 0000 0800    5
Local         100.20.1.1      Fa0/0         100.20.1.2      2F 0000 0000    5
```

Table 27 describes the significant fields in the NetFlow record lines of the displays:

Table 27 show ip cache flow Field Descriptions in NetFlow Record Display

Field	Description
SrcIf	Interface on which the packet was received.
SrcIPaddress	IP address of the device which transmitted the packet.
DstIf	Interface from which the packet was transmitted.
DstIPaddress	IP address of the destination device.
Pr	IP protocol "well-known" port number as described in RFC 1340, displayed in hexadecimal format.
SrcP	IP port from which the packet is transmitted, displayed in hexadecimal format.
DstP	IP port where the packet is to be delivered, displayed in hexadecimal format.
Pkts	Number of packets switched through this flow.

The following shows sample output from the **show ip cache verbose flow** command for interface e1/2 on 10.0.0.1 255.0.0.0:

```
Router# show ip cache 10.0.0.1 255.0.0.0 e1/2 verbose flow

IP packet size distribution (35 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

The output above show the percentage distribution of packets by size range. In this display,100 percent of the packets fall in the 138 byte size range.

```
IP Flow Switching Cache, 4456704 bytes
  1 active, 65535 inactive, 7 added
  99 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
  last clearing of statistics never

Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
ICMP                 6      0.0         5   100      0.0       0.0      15.2
Total:               6      0.0         5   100      0.0       0.0      15.2

SrcIf          SrcIPaddress    DstIf          DstIPaddress    Pr TOS Flgs  Pkts
Port Msk AS                    Port Msk AS    NextHop              B/Pk  Active
Et1/2          10.0.0.2        Local          10.0.0.1        01 00  10       5
0000 /8  0                     0800 /8  0     0.0.0.0                100     0.0
```

Table 28 describes the significant fields in the NetFlow record lines of the display.

Table 28 show ip cache verbose flow Field Descriptions in NetFlow Record Display

Field	Description
SrcIf	Interface on which the packet was received.
Port Msk AS	Source Border Gateway Protocol (BGP) autonomous system. This is always set to 0 in MPLS flows.

Table 28 show ip cache verbose flow Field Descriptions in NetFlow Record Display (continued)

Field	Description
SrcIPaddress	IP address of the device which transmitted the packet.
DstIf	Interface from which the packet was transmitted.
Port Msk AS	Destination BGP autonomous system. This is always set to 0 in MPLS flows.
DstIPaddress	IP address of the destination device.
NextHop	Specifies the BGP next hop address. This is always set to 0 in MPLS flows.
Pr	IP protocol well-known port number as described in RFC 1340, displayed in hexadecimal format.
B/Pk	Average number of bytes observed for the packets seen for this protocol (total bytes for this protocol or the total number of flows for this protocol for this summary period).
Flgs	TCP flags (result of bitwise OR of TCP flags from all packets in the flow).
Active	Number of active flows in the NetFlow cache at the time this command was entered.
Pkts	Number of packets switched through this flow.

Related Commands

Command	Description
clear ip flow stats	Clears the NetFlow switching statistics.
ip route-cache	Controls the use of high-speed switching caches for IP routing.
tunnel flow egress-records	Creates a NetFlow record for packets that are encapsulated by the GRE tunnel.

show ip cache flow aggregation

To display the aggregation cache configuration, use the **show ip cache flow aggregation** command in EXEC mode.

show ip cache [*prefix mask*] [*type number*] [**verbose**] **flow aggregation** *type*

Syntax Description

prefix mask	(Optional) Displays only the entries in the cache that match the prefix and mask combination.
type number	(Optional) Displays only the entries in the cache that match the interface type and number combination.
verbose	(Optional) Displays additional information from the aggregation cache.
type	Displays the configuration of a particular aggregation cache as follows: • Autonomous system • Destination prefix • Prefix • Protocol-port • Source prefix

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(3)T	This command was introduced.

Examples

The following is an example display of an autonomous system aggregation cache using the **show ip cache flow aggregation as** command:

```
Router# show ip cache flow aggregation as

IP Flow Switching Cache, 278544 bytes
  2 active, 4094 inactive, 13 added
  178 ager polls, 0 flow alloc failures

Src If       Src AS  Dst If       Dst AS  Flows   Pkts  B/Pk  Active
Fa1/0         0      Null          0         1      2     49    10.2
Fa1/0         0      Se2/0         20        1      5    100     0.0
```

The following is a sample display of an autonomous system aggregation cache for the prefix mask 10.0.0.1 255.0.0.0 using the **show ip cache flow aggregation as** command:

```
Router# show ip cache 10.0.0.1 255.0.0.0 flow aggregation as

IP Flow Switching Cache, 278544 bytes
  2 active, 4094 inactive, 13 added
  178 ager polls, 0 flow alloc failures

Src If         Src AS  Dst If         Dst AS     Flows    Pkts  B/Pk  Active
e1/2            0      Null            0            1     2       49    10.2
e1/2            0      e1/2           20            1     5      100     0.0
```

The following is a sample display of an autonomous system aggregation cache for 10.0.0.1 255.0.0.0 Ethernet1/2 using the **show ip cache verbose flow aggregation as** command:

```
Router# show ip cache 10.0.0.1 255.0.0.0 e1/2 verbose flow aggregation as

IP Flow Switching Cache, 278544 bytes
  2 active, 4094 inactive, 13 added
  178 ager polls, 0 flow alloc failures

Src If         Src AS  Dst If         Dst AS     Flows    Pkts  B/Pk  Active
e1/2            0      Null            0            1     2       49    10.2
e1/2            0      e1/2           20            1     5      100     0.0
```

Table 29 describes the significant fields shown in these examples.

Table 29 show ip cache flow aggregation Field Descriptions

Field	Description
bytes	Number of bytes of memory used by the NetFlow cache.
active	Number of active flows in the NetFlow cache at the time this command was entered.
inactive	Number of flow buffers that are allocated in the NetFlow cache, but are not currently assigned to a specific flow at the time this command is entered.
added	Number of flows created since the start of the summary period.
ager polls	Number of times the NetFlow code looked at the cache to cause entries to expire (used by Cisco for diagnostics only).
flow alloc failures	Number of times the NetFlow code tried to allocate a flow but could not.
Src If	Specifies the source interface.
Src AS	Specifies the source autonomous system.
Dst If	Specifies the destination interface.
Dst AS	Specifies the destination autonomous system.
Flows	Number of flows.
Pkts	Number of packets.
B/Pk	Average number of bytes observed for the packets seen for this protocol (total bytes for this protocol or the total number of flows for this protocol for this summary period).
Active	Number of active flows in the NetFlow cache at the time this command was entered.

Related Commands

Command	Description
ip flow-aggregation cache	Enables aggregation cache configuration mode.

show ip cef

To display entries in the forwarding information base (FIB) or to display a summary of the FIB, use the **show ip cef** command in EXEC mode:

show ip cef [**vrf** *vrf-name*] [[**unresolved** [**detail**]] | [**detail** | **summary**]]

Specific FIB Entries Based on IP Address Information

show ip cef [**vrf** *vrf-name*] [*network* [*mask*]] [**longer-prefixes**] [**detail**]

Specific FIB Entries Based on Interface Information

show ip cef [**vrf** *vrf-name*] [*type number*] [**detail**]

Specific FIB Entries Based on Nonrecursive Routes

show ip cef [**vrf** *vrf-name*] **non-recursive** [**detail**]

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
unresolved	(Optional) Displays unresolved FIB entries.
detail	(Optional) Displays detailed FIB entry information.
summary	(Optional) Displays a summary of the FIB.
network	(Optional) Displays the FIB entry for the specified destination network.
mask	(Optional) Displays the FIB entry for the specified destination network and mask.
longer-prefixes	(Optional) Displays FIB entries for more specific destinations.
type number	(Optional) Interface type and number for which to display FIB entries.
non-recursive	Displays only nonrecursive routes.

Command Modes

EXEC

Command History

Release	Modification
11.2 GS	This command was introduced to support the Cisco 12012 Internet router.
11.1 CC	Multiple platform support was added.
12.0(5)T	The **vrf** keyword was added.

Usage Guidelines

The **show ip cef** command without any keywords or arguments shows a brief display of all FIB entries.

The **show ip cef detail** command shows detailed FIB entry information for all FIB entries.

Examples

The following is sample output from the **show ip cef unresolved** command:

```
Router# show ip cef unresolved

IP Distributed CEF with switching (Table Version 136632)
45776 routes, 13 unresolved routes (0 old, 13 new)
45776 leaves, 2868 nodes, 8441480 bytes, 136632 inserts, 90856 invalidations
1 load sharing elements, 208 bytes, 1 references
1 CEF resets, 1 revisions of existing leaves
refcounts: 527292 leaf, 465617 node

148.214.0.0/16, version 136622
0 packets, 0 bytes
  via 171.69.233.56, 0 dependencies, recursive
  unresolved
148.215.0.0/16, version 136623
0 packets, 0 bytes
  via 171.69.233.56, 0 dependencies, recursive
  unresolved
148.218.0.0/16, version 136624
0 packets, 0 bytes
```

Table 30 describes the significant fields shown in the display.

Table 30 show ip cef unresolved Field Descriptions

Field	Description
routes	Total number of entries in the CEF table
unresolved routes	Number of entries in the CEF table that do not have resolved recursions categorized by old and new routes
leaves, nodes, bytes	Number of elements in the CEF Trie and how much memory they use
inserts	Number of nodes inserted
invalidations	Number of entries that have been invalidated
load sharing elements, bytes, references	Information about load sharing elements: how many, number of associated bytes, and number of associated references
version	Version of the CEF table
packets, bytes	Number of packets and bytes switched through the name entry
dependencies	Number of table entries which point to the named entry
recursive	Indicates that the destination is reachable through another route
unresolved	Number of entries that do not have resolved recursions

The following is sample output from the **show ip cef summary** command:

```
Router# show ip cef summary

IP Distributed CEF with switching (Table Version 135165)
45788 routes, 0 reresolve, 4 unresolved routes (0 old, 4 new)
45788 leaves, 2868 nodes, 8442864 bytes, 135165 inserts, 89377 invalidations
0 load sharing elements, 0 bytes, 0 references
1 CEF resets, 0 revisions of existing leaves
refcounts: 527870 leaf, 466167 node
```

For a description of significant fields in this display, see Table 30.

The following is sample output from the **show ip cef detail** command for Ethernet interface 0. It shows all the prefixes resolving through adjacency pointing to next hop Ethernet interface 0/0 and next hop interface IP address 172.19.233.33.

```
Router# show ip cef e0/0 172.19.233.33 detail

IP Distributed CEF with switching (Table Version 136808)
45800 routes, 8 unresolved routes (0 old, 8 new) 45800 leaves, 2868 nodes, 8444360 bytes,
136808 inserts, 91008 invalidations 1 load sharing elements, 208 bytes, 1 references 1 CEF
resets, 1 revisions of existing leaves refcounts: 527343 leaf, 465638 node

172.19.233.33/32, version 7417, cached adjacency 172.19.233.33 0 packets, 0 bytes,
Adjacency-prefix
via 172.19.233.33, Ethernet0/0, 0 dependencies
next hop 172.19.233.33, Ethernet0/0
valid cached adjacency
```

Table 31 describes the significant fields in the display.

Table 31 show ip cef detail Field Descriptions

Field	Description
routes	Total number of entries in the CEF table
unresolved routes	Number of entries in the CEF table that do not have resolved recursions categorized by old and new routes
leaves, nodes, bytes	Number of elements in the CEF Trie and how much memory they use
inserts	Number of nodes inserted
invalidations	Number of entries that have been invalidated
load sharing elements, bytes, references	Information about load sharing elements: how many, number of associated bytes, and number of associated references
version	Version of the CEF table
cached adjacency	Type of adjacency to which this CEF table entry points
packets, bytes	Number of packets and bytes switched through the name entry
dependencies	Number of table entries which point to the named entry
next hop	Type of adjacency or the next hop toward the destination

The following example shows the forwarding table associated with the VRF named vrf1:

```
Router# show ip cef vrf vrf1

Prefix              Next Hop            Interface
0.0.0.0/32          receive
11.0.0.0/8          50.0.0.1            Ethernet1/3
12.0.0.0/8          52.0.0.2            POS6/0
50.0.0.0/8          attached            Ethernet1/3
50.0.0.0/32         receive
50.0.0.1/32         50.0.0.1            Ethernet1/3
50.0.0.2/32         receive
50.255.255.255/32   receive
51.0.0.0/8          52.0.0.2            POS6/0
224.0.0.0/24        receive
255.255.255.255/32  receive
```

Table 32 describes the significant fields shown in the display.

Table 32 show ip cef vrf Field Descriptions

Field	Description
Prefix	Specifies the network prefix.
Next Hop	Specifies the Border Gateway Protocol (BGP) next hop address.
Interface	Specifies the VRF interface.

Related Commands

Command	Description
show cef drop	Displays which packets the line cards dropped or displays which packets were not express-forwarded.
show cef interface	Displays CEF-related interface information.

show ip cef adjacency

To display Cisco Express Forwarding (CEF) recursive and direct prefixes resolved through an adjacency, use the **show ip cef adjacency** command in EXEC mode.

show ip cef [**vrf** *vrf-name*] **adjacency** *type number ip-prefix* [**detail**]

To display CEF recursive and direct prefixes resolved through special adjacency types representing nonstandard switching paths, use this form of the **show ip cef adjacency** EXEC command.

show ip cef [**vrf** *vrf-name*] **adjacency** {**discard** | **drop** | **glean** | **null** | **punt**} [**detail**]

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
type number	Interface type and number for which to display forwarding information base (FIB) entries.
ip-prefix	Next hop IP prefix, in dotted decimal format (A.B.C.D).
detail	(Optional) Displays detailed information for each CEF adjacency type entry.
discard	Discard adjacency. Sets up for loopback interfaces. Loopback IP addresses are receive entries in the FIB table.
drop	Drop adjacency. Packets forwarded to this adjacency are dropped.
glean	Glean adjacency. Represents destinations on a connected interface for which no ARP cache entry exists.
null	Null adjacency. Formed for the Null0 interface. Packets forwarded to this adjacency are dropped.
punt	Punt adjacency. Represents destinations that cannot be switched in the normal path and that are punted to the next fastest switching vector.

Command Modes

EXEC

Command History

Release	Modification
11.1 CC	This command was introduced.
12.0(5)T	The **vrf** keyword was added.

Usage Guidelines

An adjacency is a node that can be reached by one Layer 2 hop.

This command shows all prefixes resolved through a regular next hop adjacency or through a special adjacency type such as discard, drop, glean, null and punt.

The following sample output is from the **show ip cef adjacency** command when the **glean** type is specified.

```
Router# show ip cef adjacency glean

Prefix              Next Hop            Interface
```

```
9.2.61.0/24         attached            Ethernet1/0/0
172.17.250.252/32   9.2.61.1            Ethernet1/0/0
```

The following sample output is from the **show ip cef adjacency drop** command with **detail** specified:

```
Router# show ip cef adjacency drop detail

IP CEF with switching (Table Version 4), flags=0x0
  4 routes, 0 reresolve, 0 unresolved (0 old, 0 new), peak 0
  4 leaves, 8 nodes, 8832 bytes, 13 inserts, 9 invalidations
  0 load sharing elements, 0 bytes, 0 references
  universal per-destination load sharing algorithm, id 00B999CA
  3 CEF resets, 0 revisions of existing leaves
  Resolution Timer: Exponential (currently 1s, peak 1s)
  0 in-place modifications
  refcounts:  533 leaf, 536 node

224.0.0.0/4, version 3
0 packets, 0 bytes, Precedence routine (0)
  via 0.0.0.0, 0 dependencies
    next hop 0.0.0.0
    valid drop adjacency
```

The following sample output shows the direct IP prefix when the next hop Gigabit Ethernet interface 3/0 is specified:

```
Router# show ip cef adjacency GigabitEthernet 3/0 172.20.26.29

Prefix              Next Hop            Interface
34.1.1.0/24         172.20.26.29        GigabitEthernet3/0
```

Table 33 describes the significant fields shown in the display.

Table 33 show ip cef adjacency Field Descriptions

Field	Description
Prefix	Destination IP prefix.
Next Hop	Next hop IP address.
Interface	Next hop interface.

Related Commands

Command	Description
show adjacency	Displays CEF adjacency table information.

show ip cef events

To display all recorded Cisco Express Forwarding (CEF) forwarding information base (FIB) and adjacency events, use the **show ip cef events** command in EXEC mode.

show ip cef [**vrf** *vrf-name*] **events** [*ip-prefix*] [**new** | **within** *seconds*] [**detail**] [**summary**]

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
ip-prefix	(Optional) Next hop IP prefix, in dotted decimal format (A.B.C.D).
new	(Optional) Displays new CEF events not previously shown.
within *seconds*	(Optional) Displays CEF events that occurred within a specified number of seconds.
detail	(Optional) Displays detailed information for each CEF event entry.
summary	(Optional) Displays a summary of the CEF event log.

Command Modes

EXEC

Command History

Release	Modification
12.0(15)S	This command was introduced.
12.2(2)T	This command was integrated into Cisco IOS Release 12.2(2)T.

Usage Guidelines

This command shows the state of the table event log and must be enabled for events to record.

The **ip cef table event-log** command controls parameters such as event log size.

Examples

The following sample output is from the **show ip cef events** command with **summary** specified:

```
Router# show ip cef events summary

CEF table events summary:
     Storage for 10000 events (320000 bytes), 822/0 events recorded/ignored
     Matching all events, traceback depth 16
     Last event occurred 00:00:06.516 ago.
```

The following sample output is from the **show ip cef events** command displaying events that occurred within 1 second:

```
Router# show ip cef events within 1

CEF table events (storage for 10000 events, 14 events recorded)
+00:00:00.000:[Default-table] *.*.*.*/*          New FIB table          [OK]
+00:00:00.000:[Default-table] 9.1.80.194/32      FIB insert in mtrie    [OK]
+00:00:00.000:[Default-table] 9.1.80.0/32        FIB insert in mtrie    [OK]
+00:00:00.000:[Default-table] 9.1.80.255/32      FIB insert in mtrie    [OK]
+00:00:00.004:[Default-table] 9.1.80.0/24        FIB insert in mtrie    [OK]
```

```
+00:00:00.004:[Default-table] 9.1.80.0/24        NBD up                [OK]
+00:00:00.004:[Default-table] 224.0.0.0/4        FIB insert in mtrie   [OK]
+00:00:00.012:[Default-table] 9.1.80.0/24        NBD up                [Ignr]
+00:00:00.012:[Default-table] 224.0.0.0/4        FIB remove            [OK]
+00:00:00.016:[Default-table] 224.0.0.0/4        FIB insert in mtrie   [OK]
+00:00:05.012:[Default-table] 224.0.0.0/4        FIB remove            [OK]
+00:00:05.012:[Default-table] 224.0.0.0/4        FIB insert in mtrie   [OK]
+00:00:28.440:[Default-table] 224.0.0.0/4        FIB remove            [OK]
+00:00:28.440:[Default-table] 224.0.0.0/4        FIB insert in mtrie   [OK]
First event occured at 00:00:36.568 (00:04:40.756 ago)
Last event occured at 00:01:05.008 (00:04:12.316 ago)
```

Table 34 describes the significant fields shown in the display.

Table 34 show ip cef events Field Descriptions

Field	Description
+00:00:00.000	Time stamp of the IP CEF event.
[Default-table]	Type of VPN routing and forwarding (VRF) table for this event entry.
`*.*.*.*/*`	All IP prefixes.
9.1.80.194/32	IP prefix associated with the event.
FIB insert in mtrie	IP prefix insert in the FIB table event.
NBD up	IP prefix up event.
FIB remove	FIB entry remove event.
[Ignr]	CEF ignored event.
[OK]	CEF processed event.

Related Commands

Command	Description
IP cef table consistency-check	Enables CEF table consistency checker types and parameters.
ip cef table event-log	Controls CEF table event-log characteristics.

show ip cef exact-route

To display the exact route for a source-destination IP address pair, use the **show ip cef exact-route** command in EXEC mode.

show ip cef [**vrf** *vrf-name*] **exact-route** *source-address destination-address*

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
source-address	Specifies the network source address.
destination-address	Specifies the network destination address.

Command Modes

EXEC

Command History

Release	Modification
12.1(4)T	This command was introduced.

Usage Guidelines

When you are load balancing per destination, this command shows the exact next hop that is used for a given IP source-destination pair.

Examples

The following sample output is from the **show ip cef exact-route** command:

```
Router# show ip cef exact-route 1.1.1.1 172.17.249.252

1.1.1.1         -> 172.17.249.252 :Ethernet2/0/0 (next hop 9.1.104.1)
```

Table 35 describes the significant fields shown in the display.

Table 35 show ip cef exact-route Field Descriptions

Field	Description
1.1.1.1 -> 172.17.249.252	From source 1.1.1.1 to destination 172.17.249.252.
Ethernet2/0/0 (next hop 9.1.104.1)	Next hop is 9.1.104.1 on Ethernet 2/0/0.

show ip cef inconsistency

To display Cisco Express Forwarding (CEF) IP prefix inconsistencies, use the **show ip cef inconsistency** command in EXEC mode.

show ip cef [**vrf** *vrf-name*] **inconsistency** [**records** [**detail**]]

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
records	(Optional) Displays all recorded inconsistencies.
detail	(Optional) Displays detailed information for each CEF table entry.

Command Modes

EXEC

Command History

Release	Modification
12.0(15)S	This command was introduced.
12.2(2)T	This command was integrated into Cisco IOS Release 12.2(2)T.

Usage Guidelines

This command is available only on routers with line cards.

This command displays recorded CEF inconsistency records found by the lc-detect, scan-rp, scan-rib, and scan-lc detection mechanisms.

You can configure the CEF consistency detection mechanisms using the **ip cef table consistency-check** command.

Examples

The following sample output is from the **show ip cef inconsistency** command:

```
Router# show ip cef inconsistency

Table consistency checkers (settle time 65s)
 lc-detect:running
  0/0/0 queries sent/ignored/received
 scan-lc:running [100 prefixes checked every 60s]
  0/0/0 queries sent/ignored/received
 scan-rp:running [100 prefixes checked every 60s]
  0/0/0 queries sent/ignored/received
 scan-rib:running [1000 prefixes checked every 60s]
  0/0/0 queries sent/ignored/received
Inconsistencies:0 confirmed, 0/16 recorded
```

Table 36 describes the significant fields shown in the display.

Table 36 show ip cef inconsistency Field Descriptions

Field	Description
settle time	Time after a recorded inconsistency is confirmed.
lc-detect running	Consistency checker lc-detect is running.
0/0/0 queries	Number of queries sent, ignored, and received.
Inconsistencies:0 confirmed, 0/16 recorded	Number of inconsistencies confirmed, and recorded. Sixteen is the maximum number of inconsistency records to be recorded.

Related Commands

Command	Description
IP cef table consistency-check	Enables CEF table consistency checker types and parameters.

show ip cef traffic prefix-length

To display Cisco Express Forwarding (CEF) traffic statistics, use the **show ip cef traffic prefix-length** command in EXEC mode.

show ip cef [**vrf** *vrf-name*] **traffic prefix-length**

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.
vrf-name	(Optional) Name assigned to the VRF.
prefix-length	Displays traffic statistics by prefix size.

Command Modes

EXEC

Command History

Release	Modification
11.1 CC	This command was introduced.
12.0(5)T	The **vrf** keyword was added.

Usage Guidelines

This command is used to display CEF switched traffic statistics by destination prefix length. The **ip cef accounting prefix-length** command must be enabled for the counters to increment.

Examples

The following sample output is from the **show ip cef traffic prefix-length** command:

```
Router# show ip cef traffic prefix-length

IP prefix length switching statistics:
----------------------------------------
Prefix       Number of       Number of
Length       Packets         Bytes
----------------------------------------
     0               0               0
     1               0               0
     2               0               0
     3               0               0
     4               0               0
     5               0               0
.
.
.
    28               0               0
    29               0               0
    30               0               0
    31               0               0
    32               0               0
```

Table 37 describes the significant fields shown in the display.

Table 37 show ip cef traffic prefix-length Field Descriptions

Field	Description
Prefix Length	Destination IP prefix length for CEF switched traffic.
Number of packets	Number of packets forwarded for the specified IP prefix length.
Number of bytes	Number of bytes transmitted for the specified IP prefix length.

Related Commands

Command	Description
ip cef accounting	Enables network accounting of CEF.

show ip explicit-paths

To display the configured IP explicit paths, use the **show ip explicit-paths** EXEC command. An IP explicit path is a list of IP addresses, each representing a node or link in the explicit path.

show ip explicit-paths [{**name** *word* | **identifier** *number*}] [**detail**]

Syntax Description

name *word*	(Optional) Name of the explicit path.
identifier *number*	(Optional) 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.

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)S	This command was introduced.

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 3.3.28.3
    2: next-address 3.3.27.3
```

Table 38 describes the significant fields displayed in the output.

Table 38 show ip explicit-paths Field Descriptions

Field	Description
PATH	Path name 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. Commands might be renumbered as a result.
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.

Command	Description
list	Displays all or part of the explicit paths.
next-address	Specifies the next IP address in the explicit path.

show ip flow export

To display the statistics for the data export, including the main cache and all other enabled caches, use the **show ip flow export** command in user EXEC or privileged EXEC mode.

show ip flow export [**template**]

Syntax Description

template	(Optional) Shows the data export statistics (such as template timeout and refresh rate) for the template-specific configurations.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
11.1CC	This command was introduced.
12.2(2)T	This command was modified to display multiple NetFlow export destinations.
12.0(24)S	The **template** keyword was added.
12.3(1)	This command was integrated into Cisco IOS Release 12.3(1).

Examples

The following is sample output from the **show ip flow export** command:

```
Router# show ip flow export

Flow export v5 is enabled for main cache
  Exporting flows to 10.51.12.4 (9991) 10.1.97.50 (9111)
  Exporting using source IP address 9.1.97.17
  Version 5 flow records
  11 flows exported in 8 udp datagrams
  0 flows failed due to lack of export packet
  0 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup failures
  0 export packets were dropped enqueuing for the RP
  0 export packets were dropped due to IPC rate limiting
  0 export packets were dropped due to output drops
```

Table 39 describes the significant fields shown in the display.

Table 39 show ip flow export Field Descriptions

Field	Description
Exporting flows to 10.51.12.4 (9991) 10.1.97.50 (9111)	Specifies the export destinations and ports. The ports are in parentheses.
Exporting using source IP address 9.1.97.17	Specifies the source address or interface.
Version 5 flow records	Specifies the version of the flow.
11 flows exported in 8 udp datagrams	The total number of export packets sent, and the total number of flows contained within them.
0 flows failed due to lack of export packet	No memory was available to create an export packet.
0 export packets were sent up to process level	The packet could not be processed by CEF or by fast switching, possibly because another feature requires running on the packet.
0 export packets were dropped due to no fib 0 export packets were dropped due to adjacency issues	Indicates that CEF was unable to switch the packet or forward it up to the process level.
0 export packets were dropped due to fragmentation failures 0 export packets were dropped due to encapsulation fixup failures	Indicates that the packet was dropped because of problems constructing the IP packet.
0 export packets were dropped enqueuing for the RP 0 export packets were dropped due to IPC rate limiting	Indicates that there was a problem transferring the export packet between the RP and the line card.
0 export packets were dropped due to output drops	Indicates that the send queue was full while the packet was being transmitted.

Related Commands

Command	Description
clear adjacency	Configures aggregation cache operational parameters.
exit	Leaves aggregation cache mode.
ip flow-aggregation cache	Enables aggregation cache configuration mode.

show ip mcache

To display the contents of the IP multicast fast-switching cache, use the **show ip mcache** command in EXEC mode.

show ip mcache [*group* [*source*]]

Syntax Description

group	(Optional) Displays the fast-switching cache for the single group. The *group* argument can be either a Class D IP address or a DNS name.
source	(Optional) If the *source* argument is also specified, displays a single multicast cache entry. The *source* argument can be either a unicast IP address or a DNS name.

Command Modes

EXEC

Command History

Release	Modification
11.0	This command was introduced.

Usage Guidelines

Use this command on the RP.

Examples

The following is sample output from the **show ip mcache** command. This entry shows a specific source (wrn-source 204.62.246.73) sending to the World Radio Network group (224.2.143.24).

```
show ip mcache wrn wrn-source

IP Multicast Fast-Switching Cache
(204.62.246.73/32, 224.2.143.24), Fddi0, Last used: 00:00:00
  Ethernet0       MAC Header: 01005E028F1800000C1883D30800
  Ethernet1       MAC Header: 01005E028F1800000C1883D60800
  Ethernet2       MAC Header: 01005E028F1800000C1883D40800
  Ethernet3       MAC Header: 01005E028F1800000C1883D70800
```

Table 40 describes the significant fields shown in the output.

Table 40 show ip mcache Field Descriptions

Field	Description
204.62.246.73	Source address.
224.2.143.24	Destination address.
Fddi0	Incoming or expected interface on which the packet should be received.

Table 40 show ip mcache Field Descriptions (continued)

Field	Description
Last used:	Latest time the entry was accessed for a packet that was successfully fast switched: • "semi-fast" indicates that the first part of the outgoing interface list is fast switched and the rest of the list is process-level switched. • "mds" indicates that multicast distributed switching is being used instead of the fast cache. • "never" indicates that the fast cache entry is not used (it is process switched).
Ethernet0 MAC Header:	Outgoing interface list and respective MAC header that is used when rewriting the packet for output. If the interface is a tunnel, the MAC header will show the real next hop MAC header and then, in parentheses, the real interface name.

The following is sample output from the **show ip mcache** command when MDS is in effect.

```
Router# show ip mcache

IP Multicast Fast-Switching Cache
(*, 224.2.170.73), Fddi3/0/0, Last used: mds
  Tunnel3         MAC Header: 5000602F9C150000603E473F60AAAA030000000800 (Fddi3/0/0)
  Tunnel0         MAC Header: 5000602F9C150000603E473F60AAAA030000000800 (Fddi3/0/0)
  Tunnel1         MAC Header: 5000602F9C150000603E473F60AAAA030000000800 (Fddi3/0/0)
```

show ip mds forwarding

On a line card, to display the MFIB table and forwarding information for multicast distributed switching (MDS), use the **show ip mds forwarding** command in EXEC mode.

show ip mds forwarding [*group-address*] [*source-address*]

Syntax Description

group-address	(Optional) Address of the IP multicast group for which to display the MFIB table.
source-address	(Optional) Address of the source of IP multicast packets for which to display the MFIB table.

Command Modes

EXEC

Command History

Release	Modification
11.2(11)GS	This command was introduced.

Usage Guidelines

Use this command on the line card. This command displays the MFIB table, forwarding information, and related flags and counts.

Note To reach the console for a line card, enter **attach** *slot#* (slot number where the line card resides).

On a GSR only, line card commands can be executed from the RP using the following syntax: **execute** [**slot** *slot-number* | **all**] *command.*

The *command* argument is any of the line card **show** commands, such as **show ip mds summary** and **show ip mds forward**.

Examples

The following is sample output from the **show ip mds forwarding** command:

```
Router# show ip mds forwarding

IP multicast MDFS forwarding information and statistics:
Flags: N - Not MDFS switchable, F - Not all MDFS switchable, O - OIF Null
       R - In-ratelimit, A - In-access, M - MTU mismatch, P - Register set

Interface state: Interface, Next-Hop, Mac header

(*, 224.2.170.73),
  Incoming interface: Null
  Pkts: 0, last used: never, Kbps: 0, fast-flags: N
  Outgoing interface list: Null

(128.97.62.86, 224.2.170.73) [31]
  Incoming interface: Fddi3/0/0
  Pkts: 3034, last used: 00:00:00, Kbps: 0, fast-flags: M
  Outgoing interface list:
```

Table 41 describes the significant fields shown in the output.

Table 41 show ip mds forwarding Field Descriptions

Field	Description
(128.97.62.86, 224.2.170.73) [31])	Source and group addresses. The number in brackets is the hash bucket for the route.
Incoming interface:	Expected interface for a multicast packet from the source. If the packet is not received on this interface, it is discarded.
Pkts	Total number of packets switched by that entry.
last used:	Time when this MFIB entry was used to switch a packet.
Kbps:	Kilobits per second of the switched traffic.
Outgoing interface list:	Interfaces through which packets will be forwarded.

show ip mds interface

To display the status of multicast distributed switching (MDS) interfaces, use the **show ip mds interface** command in EXEC mode.

show ip mds interface

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
11.2(11)GS	This command was introduced.

Usage Guidelines

Use this command on the RP.

Examples

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

```
Router# show ip mds interface

Ethernet1/0/0 is up, line protocol is up
Ethernet1/0/1 is up, line protocol is up
Fddi3/0/0 is up, line protocol is up
FastEthernet3/1/0 is up, line protocol is up
```

Table 42 describes the significant fields in the output.

Table 42 show ip mds interface Field Descriptions

Field	Description
Ethernet1/0/0 is up	Status of interface.
line protocol is up	Status of line protocol.

show ip mds stats

To display switching statistics or line card statistics for multicast distributed switching (MDS), use the **show ip mds stats** command in EXEC mode.

show ip mds stats [**switching** | **linecard**]

Syntax Description

switching	(Optional) Displays switching statistics.
linecard	(Optional) Displays line card statistics.

Command Modes

EXEC

Command History

Release	Modification
11.2(11)GS	This command was introduced.

Usage Guidelines

Use this command on the RP.

Examples

The following is sample output from the **show ip mds stats** command used with the **switching** keyword:

```
Router# show ip mds stats switching

Slot Total       Switched    Drops       RPF         Punts       Failures
                                                                 (switch/clone)
 1   0           0           0           0           4           0/0
 3   20260925    18014717    253         93          2247454     1/0
```

Table 43 describes the significant fields in the output.

Table 43 show ip mds stats switching Field Descriptions

Field	Description
Slot	Slot number for the line card.
Total	Total number of packets received.
Switched	Total number of packets switched.
Drops	Total number of packets dropped.
RPF	Total number of packets that failed RPF lookup.
Punts	Total number of packets sent to the RP because the line card could not switch them.
Failures (switch/clone)	Times that the RP tried to switch but failed because of lack of resources or clone for the RSP only; failed to get a packet clone.

The following is sample output from the **show ip mds stats** command with the **linecard** keyword:

```
Router# show ip mds stats linecard

Slot      Status    IPC(seq/max) Q(high/route)  Reloads
 1        active    10560/10596      0/0            9
 3        active    11055/11091      0/0            9
```

show ip mds summary

To display a summary of the MFIB table for multicast distributed switching (MDS), use the **show ip mds summary** command in EXEC mode.

show ip mds summary

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
11.2(11)GS	This command was introduced.

Usage Guidelines

Use this command on a line card. On a GSR only, line card commands can be executed from the RP using the following syntax:

execute [**slot** *slot-number* | **all**] *command*

The *command* argument is any of the line card **show** commands, such as **show ip mds summary** and **show ip mds forward**.

Examples

The following is sample output from the **show ip mds summary** command:

```
Router# show ip mds summary

IP multicast MDFS forwarding information and statistics:
Flags: N - Not MDFS switchable, F - Not all MDFS switchable, O - OIF Null
       R - In-ratelimit, A - In-access, M - MTU mismatch, P - Register set

Interface state: Interface, Next-Hop, Mac header

(*, 224.2.170.73),
  Incoming interface: Null
  Pkts: 0, last used: never, Kbps: 0, fast-flags: N
(128.97.62.86, 224.2.170.73) [31]
  Incoming interface: Fddi3/0/0
  Pkts: 3045, last used: 00:00:03, Kbps: 0, fast-flags: M
(128.223.3.7, 224.2.170.73) [334]
  Incoming interface: Fddi3/0/0
  Pkts: 0, last used: never, Kbps: 0, fast-flags: M
```

Table 44 describes the significant fields in the output.

Table 44 show ip mds summary Field Descriptions

Field	Description
(128.97.62.86, 224.2.170.73) [31]	Source and group addresses. The number in brackets is the hash bucket for the route.
Incoming interface	Expected interface for a multicast packet from the source. If the packet is not received on this interface, it is discarded.
Pkts	Total number of packets switched by that entry.
last used	Time when this MFIB entry was used to switch a packet.
Kbps	Kilobits per second of the switched traffic.