Edited
M. Amyx
4/20/2009

ip flow-capture

(This notation says that one and only one option must be chosen. Is that correct?--no combinations? So if you want to capture values in more than one field, you must issue the command once for each field? or you can capture only one field at any one time? Perhaps those restrictions ought to be mentioned in the usage guidelines.)

To enable the capture of values from Layer 2 or additional Layer 3 fields in NetFlow traffic, use the **ip flow-capture** command in global configuration mode. To disable capturing Layer 2 or Layer 3 fields from NetFlow traffic, use the **no** form of this command.

ip flow-capture {**fragment-offset** | **icmp** | **ip-id** | **mac-addresses** | **packet-length** | **ttl** | **vlan-id** | **nbar**}

no ip flow-capture {**fragment-offset** | **icmp** | **ip-id** | **mac-addresses** | **packet-length** | **ttl** | **vlan-id** | **nbar**}

- (hyphen)

Syntax Description

fragment-offset	Captures the value of the 13 bit IP fragment offset field from the first fragmented IP datagram in a flow.
icmp	Captures the value of the ICMP type and code fields from the first ICMP datagram in a flow.
ip-id	Captures the value of the IP-ID field from the first IP datagram in a flow.
mac-addresses	Captures the values of the source MAC addresses from ingress packets and the destination MAC addresses from egress packets from ~~from~~ the first packet in a flow. **Note** This command only applies to traffic that is received or transmitted over Ethernet interfaces.
packet-length	Captures the value of the packet length field from IP datagrams in a flow.
ttl	Captures the value of the Time-to-Live (TTL) field from IP datagrams in a flow.
vlan-id	Captures the value of the 802.1q or ISL VLAN-ID field from VLAN-encapsulated frames in a flow when the frames are received or transmitted on trunk ports.
nbar	Exports NBAR information along with the NetFlow version 9 record.

. (period)

(spell out acronym)

(spell out acronym)

(Throughout these commands and the fm, this is most often spelled with cap. V: Version. Should be globally consistent.)

Command Default

The **ip flow-capture** command is not enabled by default.

(A command-level default must say what the state or behavior of the system is without the command. In this case it will be something like this: "Values from Layer 2 and Layer 3 fields are not captured.")

Command Modes

Global configuration (config)

Command History

Release	Modification
12.3(14)T	This command was introduced.
12.4(2)T	The **fragment-offset** keyword was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(18)ZYA2	This command was modified. The **nbar** keyword was added.

Usage Guidelines

You must enable NetFlow accounting on an interface or a subinterface using the **ip flow** {**ingress** | **egress**} command for the **ip flow-capture** command to take effect. You can enable NetFlow accounting before or after you have entered the **ip flow-capture** command in global configuration mode.

If you want to export the information captured by the **ip flow-capture** command, you must configure NetFlow export using the **ip flow-export destination** command, and you must configure NetFlow to use the Version 9 export format. Use the **ip flow-export version 9** command to configure the NetFlow Version 9 export format.

(Note capitalization here and elsewhere.)

The fields captured by the **ip flow-capture** command are currently not available in the NetFlow MIB.

ip flow-capture fragment-offset

IP fragmentation occurs when the size of an IP datagram exceeds the maximum transmission unit (MTU) of the Layer 2 frame type used by the next-hop network. For example, the IP MTU size of an ATM network is 4470 bytes. When a host needs to transmit an IP datagram that exceeds 4470 bytes on an ATM network, it must first fragment the datagram into two or more smaller IP datagrams.

, (comma)

An IP datagram sent by a host system such as a web server can also be fragmented by a router in the network if the router needs to transmit the IP datagram on a next-hop network that has an MTU that is smaller than the current size of the IP datagram. For example if a router receives a 4470-byte IP datagram on an ATM interface and the next hop network is a 100-Mbps Fast Ethernet network with an MTU of 1514, the router must fragment the IP datagram into three smaller IP datagrams (4470/1514). It is possible for an IP datagram to be fragmented two or more times on its path from the sending host to the destination host.

- (hyphen)

A fragmented IP datagram is reassembled by the destination host. The last fragment of an IP datagram is identified when the “more fragments” flag is set to 0. The length of a complete IP datagram is calculated by the receiving host by means of the fragment offset field and the length of the last fragment.

The **ip flow-capture fragment-offset** command captures the value of the IP fragment offset field from the first fragmented IP packet in the flow. If you are seeing several flows with the same value for the IP fragment offset field, it is possible that your network is being attacked by a host that is sending the same IP packets ~~over and over.~~

(GLOBAL) again and again (means the same thing as "over and over," but is more literal and thus easier to translate)

ip flow-capture icmp

ICMP is used for several purposes. One of the most common is the ping command. ICMP echo requests are sent by a host to a destination to verify that the destination is reachable by IP. If the destination is reachable, it should respond by sending an ICMP echo reply. Refer to RFC 792 (http://www.ietf.org/rfc/rfc0792.txt) for more information on ICMP.

(expand acronym)

(bold)

ICMP packets have been used in many types of attacks on networks. Two of the most common attacks are denial-of-service (DoS) attacks and the “ping of death” attack.

the

(also give the title)

- DoS attack—Any action or actions that prevent any part of a system from functioning in accordance with its intended purpose. This includes any action that causes unauthorized delay of service. Generally, DoS attacks do not destroy data or resources, but prevent access or use. In network operations, flooding a device with ping packets when the device has not been configured to block or ignore them might effect a denial of service.
- “ping of death”—An attack that sends an improperly large ping echo request packet with the intent of overflowing the input buffers of the destination machine and causing it to crash.

Finding out the types of ICMP traffic in your network can help you decide if your network is being attacked by ICMP packets.

The **ip flow-capture icmp** command captures the value of the ICMP type field and the ICMP code field from the first ICMP packet detected in a flow.

ip flow-capture ip-id

It is possible for a host to receive IP datagrams from two or more senders concurrently. It is also possible for a host to receive multiple IP datagrams from the same host for different applications concurrently. For example, a server might be transferring email and HTTP traffic from the same host concurrently. When a host is receiving multiple IP datagrams concurrently it must be able to identify the fragments from each of the incoming datagrams to ensure that they do not get mixed up during the datagram reassembly process. The receiving host uses the IP header identification field and the source IP address of the IP datagram fragment to ensure that it rebuilds the IP datagrams correctly.

The **ip flow-capture ip-id** command captures the value of the IP header identification field from the first packet in the flow. The value in the IP header identification field is a sequence number assigned by the host that originally transmitted the IP datagram. All of the fragments of an IP datagram have the same identifier value. This ensures that the destination host can match the IP datagram to the fragment during the IP datagram reassembly process. The sending host is responsible for ensuring that each subsequent IP datagram it sends to the same destination host has a unique value for the IP header identification field.

If you are seeing several flows with the same value for the IP header identification field, it is possible that your network is being attacked by a host that is sending the same IP packets ~~over and over.~~

ip flow-capture packet-length

The value in the packet length field in an IP datagram indicates the length of the IP datagram, excluding the IP header.

Use the **ip flow-capture packet-length** command to capture the value of the IP header packet length field for packets in the flow. The **ip flow-capture packet-length** command keeps track of the minimum and maximum values captured from the flow. The minimum and maximum packet length values are stored in separate fields. This data is updated when a packet with a packet length that is lower or higher than the currently stored value is received. For example if the currently stored value for the minimum packet length is 1024 bytes and the next packet received has a packet length of 512 bytes, the 1024 is replaced ~~with~~ 512.

If you are seeing several IP datagrams in the flow with the same value for the packet-length field, it is possible that your network is being attacked by a host that is constantly sending the same IP packets over-and-over.

ip flow-capture ttl

The TTL field is used to prevent the indefinite forwarding of IP datagrams. The TTL field contains a counter value set by the source host. Each router that processes this datagram decreases the TTL value by 1. When the TTL value reaches 0, the datagram is discarded.

There are two scenarios where an IP packet without a TTL field could live indefinitely in a network:

- The first scenario occurs when a host sends an IP datagram to an IP network that ~~doesn't~~ exist and all of the routers in the network have a gateway of last resort configured—that is, a gateway to which they forward IP datagrams for unknown destinations. Each router in the network receives the datagram and attempts to determine the best interface to use to forward it. Because the destination network is unknown, the best interface for the router to use to forward the datagram to the next hop is always the interface to which the gateway of last resort is assigned.
- The second scenario occurs when there is a mis-configuration in the network that results in a routing loop. For example, suppose that one router forwards an IP datagram to another router because it appears to be the correct next-hop router. The receiving router sends it back because it believes that the correct next-hop router is the router that it received the IP datagram from in the first place.

The **ip flow-capture ttl** command keeps track of the TTL values captured from packets in the flow. The minimum and maximum TTL values are stored in separate fields. This data is updated when a packet with a TTL that is lower or higher than the currently stored value is received. For example if the currently stored value for the minimum TTL is 64 and the next packet received has a TTL of 12, the 64 is replaced by 12.

If you are seeing several flows with the same value for the TTL, it is possible that your network is being attacked by a host that is constantly sending the same IP packets ~~over and over.~~ Under normal circumstances, flows come from many sources, each a different distance away. Therefore you should see a variety of TTLs across all the flows that NetFlow is capturing.

(Ambiguous placement. Does this refer to the command above or below it? Would be better to place it directly beneath the heading for the command it pertains to.)

again and again

MAC address

ip flow-capture mac-addresses

The **ip flow-capture mac-addresses** command captures the incoming source ~~mac-address~~ and the outgoing destination ~~mac-address~~ from the first Layer 2 frame in the flow. If you discover that your network is being attacked by Layer 3 traffic, you can use these addresses to identify the device that is transmitting the traffic that is being received by the router and the next hop or final destination device to which the router is forwarding the traffic.

Note This command only applies to traffic that is received or transmitted over Ethernet interfaces.

(expand)

ip flow-capture vlan-id

A VLAN is a broadcast domain within a switched network. A broadcast domain is defined by the network boundaries within which a network propagates a broadcast frame generated by a station. Some switches can be configured to support single or multiple VLANs. Whenever a switch supports multiple VLANs, broadcasts within one VLAN never appear in another VLAN.

Each VLAN is also a separate Layer 3 network. A router or a multilayer switch must be used to interconnect the Layer 3 networks that are assigned to the VLANs. For example, in order for a device on VLAN 2 with an IP address of 172.16.0.76 to communicate with a device on VLAN 3 with an IP address of 172.17.0.34, the two devices must use a router as an intermediary device, because they are on different Class B IP networks. This is typically accomplished by connecting a switch to a router and configuring the link between them as a VLAN trunk. In order for the link to be used as a VLAN trunk, the interfaces on the router and the switch must be configured for the same VLAN encapsulation type.

Note When a router is configured to route traffic between VLANs, it is often referred to as an inter-VLAN router.

When a router or a switch needs to send traffic on a VLAN trunk, it must either tag the frames using the IEEE 802.1q protocol or encapsulate the frames using the Cisco Inter-Switch Link (ISL) protocol. The VLAN tag or encapsulation header must contain the correct VLAN ID to ensure that the device receiving the frames can process them properly. The device that receives the VLAN traffic examines the VLAN ID from each frame to find out how it should process the frame. For example, when a switch receives an IP broadcast datagram such as an Address Resolution Protocol (ARP) datagram with an 802.1q tagged VLAN ID of 6 from a router, it forwards the datagram to every interface that is assigned to VLAN 6 and any interfaces that are configured as VLAN trunks.

The **ip flow-capture vlan-id** command captures the VLAN ID number from the first frame in the flow it receives that has an 802.1q tag or that is encapsulated with ISL. When the received traffic in the flow is transmitted over an interface that is configured with either 802.1q or ISL trunking, the **ip flow-capture vlan-id** command captures the destination VLAN ID number from the 802.1q or ISL VLAN header from the first frame in the flow.

Note The **ip flow-capture vlan-id** command does not capture the type of VLAN encapsulation in use. The receiving and transmitting interfaces can use different VLAN protocols. If only one of the interfaces is configured as a VLAN trunk, the VLAN ID field is blank for the other interface.

Your router configuration must meet the following criteria before NetFlow can capture the value in the VLAN-ID field:

- It must have at least one LAN interface that is configured with one or more subinterfaces.
- The subinterfaces where you want to receive VLAN traffic must have either 802.1q or ISL enabled.
- The subinterfaces that are configured to receive VLAN traffic must have the **ip flow ingress** command configured on them.

If you discover that your network is being attacked by Layer 3 traffic, you can use the VLAN-ID information to help you find out which VLAN the device that is sending the traffic is on. The information can also help you identify the VLAN to which the router is forwarding the traffic.

ip flow-capture nbar

The **ip flow-capture nbar** command captures the application IDs and sub application IDs exported as part of the NetFlow version 9 record. The application IDs are mapped to applications. Using the **ip flow-export template options nbar** command, this mapping information is exported to the NetFlow data collector. To capture Network Based Application Recognition (NBAR) information, you must enable the NetFlow version 9.

Examples

The following example shows how to configure NetFlow to capture the value of the IP fragment-offset field from the IP datagrams in the flow:

```
Router(config)# ip flow-capture fragment-offset
```

The following example shows how to configure NetFlow to capture the value of the ICMP Type field and the value of the Code field from the IP datagrams in the flow:

```
Router(config)# ip flow-capture icmp
```

The following example shows how to configure NetFlow to capture the value of the IP-ID field from the IP datagrams in the flow:

```
Router(config)# ip flow-capture ip-id
```

The following example shows how to configure NetFlow to capture the value of the packet length field from the IP datagrams in the flow:

```
Router(config)# ip flow-capture packet-length
```

The following example shows how to configure NetFlow to capture the TTL field from the IP datagrams in the flow:

```
Router(config)# ip flow-capture ttl
```

The following example shows how to configure NetFlow to capture the MAC addresses from the IP datagrams in the flow:

```
Router(config)# ip flow-capture mac-addresses
```

The following example shows how to configure NetFlow to capture the vlan-id from the IP datagrams in the flow:

```
Router(config)# ip flow-capture vlan-id
```

NBAR

The following example shows how to configure NetFlow to capture ~~nbar~~ information:

```
Router(config)# ip flow-capture nbar
```

(These must be in alphabetical order. Please place in order as shown by the numbers.)

Related Commands

Command	Description
ip flow ingress	Enables NetFlow (ingress) accounting for traffic arriving on an interface.
ip flow egress	Enables NetFlow egress accounting for traffic that the router is forwarding.
ip flow-egress input-interface	Removes the NetFlow egress accounting flow key that specifies an output interface and adds a flow key that specifies an input interface for NetFlow egress accounting.
ip flow-cache timeout	Specifies NetFlow accounting flow cache parameters.
ip flow-cache entries	Changes the number of entries maintained in the NetFlow accounting cache.
ip flow-export template options nbar	Exports application mapping information to NetFlow data collector.
show ip cache flow	Displays a summary of ~~the~~ NetFlow accounting statistics.
show ip cache verbose flow	Displays a detailed summary of the NetFlow accounting statistics.
show ip flow export	Displays the status and the statistics for NetFlow accounting data export.
show ip flow interface	Displays NetFlow accounting configuration for interfaces.

the

the

(delete extra space)

ip flow-export template

To configure template options for the export of NetFlow accounting information in NetFlow cache entries, use the **ip flow-export template** command in global configuration mode. To return to the default behavior, use the **no** form of this command.

(What state or behavior does the 'no' form return it to? Does it remove configuration opinions? That's what goes here.)

Configure template only

ip flow-export template {**refresh-rate** *packets* | **timeout-rate** *minutes*}

no ip flow-export template {**refresh-rate** | **timeout-rate**}

Configure template options

ip flow-export template options {**export-stats** | **refresh-rate** *packets* | **timeout-rate** *minutes* | **sampler** | **nbar**}

(Make sure operators are not bold.)

no ip flow-export template options {**export-stats** | **refresh-rate** | **timeout-rate** | **sampler** | **nbar**}

Syntax Description

template	Enables the **refresh-rate** and **timeout-rate** keywords for the configuring of Version 9 export templates.
refresh-rate *packets*	(Optional) Specifies the number of export packets that are sent before the options and flow templates are resent. Range: 1 to 600 ~~packets.~~ The default is 20 ~~packets~~. **Note** This applies to the **ip flow-export template refresh-rate** *packets* command.
timeout-rate *minutes*	(Optional) Specifies the interval (in minutes) that the router waits after sending the templates (flow and options) before sending them again. Range: 1 to 3600 ~~minutes~~. The default is 30 minutes. **Note** This applies to the **ip flow-export template timeout-rate** *minutes* command.
options	(Optional) Enables the **export-stats**, **refresh-rate**, **sampler** and **timeout-rate** keywords for configuring Version 9 export options.
export-stats	(Optional) Enables the export of statistics including the total number of flows exported and the total number of packets exported.
sampler	(Optional) When Version 9 export is configured, this keyword enables the export of an option containing a random-sampler configuration, including the sampler ID, sampling mode, and sampling interval for each configured random sampler. **Note** You must have a flow sampler map configured before you can configure the sampler keyword for the **ip flow-export template options** command.
refresh-rate *packets*	(Optional) Specifies the number of packets that are sent before the configured options records are resent. Range: 1 to 600 packets. The default is 20 packets. **Note** This applies to the **ip flow-export template options refresh-rate** *packets* command.

(space)

.

(This treatment is a mixture of conventions. With the shorthand form used for range, the default should be "Default: 20." Unit is not part of the value of the argument.)

(What does this mean? Note should be unnecessary. Never seen a note like this before.)

Default: 30.

(None of these appears to be optional as they appear in the syntax above. They are shown as a required choice among mutually exclusive elements.)

(Delete--already documented above.)

timeout-rate *minutes*	(Optional) Specifies the interval (in minutes) that the router will wait after sending the options records before they are sent again. Range: 1 to 3600 minutes. The default is 30 minutes. **Note** This applies to the **ip flow-export template options timeout-rate** *minutes* command.
nbar	Exports application mapping information to the NetFlow data collector.

(Delete--already documented above.)

(See the IWG. What goes here is the command-level default. Syntax defaults belong in the syntax description table only. This will be something like "Template options are not configured." SME must ok.)

Command Default

~~The default parameters as noted in the Syntax Description table are used.~~

Command Modes

Global configuration (config)

Command History

Release	Modification
12.3(2)T	This command was introduced.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(18)SXF	This command was integrated into Cisco IOS Release 12.2(18)SXF.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(18)ZYA2	This command was modified. The **nbar** keyword was added.

Usage Guidelines

The **ip flow-export template options export-stats** command requires that the NetFlow Version 9 export format be already configured on the router.

The **ip flow-export template options sampler** command requires that the NetFlow Version 9 export format and a flow sampler map be already configured on the router.

(Ok to expand again here, but it must be expanded on page 1.)

The **ip flow-export template options nbar** command exports application IDs to string mapping as options. It displays string values for application IDs to which they are mapped. To export the application mapping information, you must enable NetFlow Export version 9 export format and have Network Based Application Recognition (NBAR) configured on the device.

Examples

The following example shows how to configure NetFlow so that the networking device sends the export statistics (total flows and packets exported) as options data:

```
Router(config)# ip flow-export template options export-stats
```

The following example shows how to configure NetFlow to wait until 100 export packets have been sent, or 60 minutes have passed since the last time the templates were sent (whichever comes first) before the templates are resent to the destination host:

```
Router(config)# ip flow-export template refresh-rate 100
Router(config)# ip flow-export template timeout-rate 60
```

The following example shows how to configure NetFlow to enable the export of information about NetFlow random samplers:

```
Router(config)# ip flow-export template option sampler
```

Tip You must have a **flow-sampler** map configured before you can configure the sampler keyword for the **ip flow-export template options** command.

The following example shows how to configure NetFlow to enable the export of application mapping information:

```
Router(config)# ip flow-export template option nbar
```

s(?) (Doesn't command say 'options'?)

Related Commands

Command	Description
ip flow-export destination	Enables the export of NetFlow accounting information in NetFlow cache entries to a remote device such as a server running an application that analyzes NetFlow data.
ip flow-export interface-names	Enables the inclusion of the interface names for the flows during the export of NetFlow accounting information in NetFlow cache entries.
ip flow-export source	Specifies the interface from which NetFlow will derive the source IP address for the NetFlow export datagrams containing NetFlow accounting information from NetFlow cache entries.
ip flow-export version	Specifies the export version format for the exporting of NetFlow accounting information in NetFlow cache entries
show ip flow export	Displays the status and the statistics for NetFlow accounting data export.

show ip flow export

To display the status and the statistics for NetFlow accounting 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 [**sctp**] [**verbose**] [**template** | **nbar**]

(no bold for operators)

Syntax Description

sctp	(Optional) Displays the status and statistics for export destinations that are configured to use the Stream Control Transmission Protocol (SCTP).
verbose	(Optional) Displays the current values for the SCTP fail-over and restore-time timers in addition to the status and statistics that are displayed by the **show ip flow export sctp** command. For a Multiprotocol Label Switching (MPLS) Prefix/Application/Label (PAL) record, displays additional export information, such as the number of MPLS PAL records exported to a NetFlow collector.
template	(Optional) Displays the data export statistics (such as template timeout and refresh rate) for the template-specific configurations.
nbar	(Optional) Displays cumulative Network Based Application Recognition (NBAR) statistics.

Command Modes

User EXEC
Privileged EXEC

(#)
(Insert prompt: pound sign within parentheses)

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)	Support for the NetFlow v9 Export Format feature was added.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(18)S	Support for the NetFlow v9 Export Format, and Multiple Export Destination features was added.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(14)SX	Support for this command was introduced on the Supervisor Engine 720.
12.2(18)SXD	The output was changed to include information about NDE for hardware-switched flows.
12.2(18)SXF	This command was integrated into Cisco IOS Release 12.2(18)SXF.
12.4(4)T	The **sctp** and **verbose** keywords were added.
12.2(28)SB	The number of MPLS PAL records exported by NetFlow was added to the verbose keyword output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Release	Modification
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.2(18)ZYA2	This command was modified. The **nbar** keyword was added.

Examples

The following is sample output from the **show ip flow export** command with NetFlow export over User Datagram Protocol (UDP) (the default NetFlow export transport protocol) configured on the networking device:

Note No NetFlow export over SCTP destinations are configured:

(Is this note complete? It should end with a period, not a colon.)

```
Router# show ip flow export

Flow export v9 is enabled for main cache
  Exporting flows to 172.17.10.2 (100)
  Exporting using source interface Loopback0
  Version 9 flow records
  62 flows exported in 17 udp datagrams
  0 flows failed due to lack of export packet
  8 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
```

The following is sample output from the **show ip flow export** command with NetFlow export over UDP and NetFlow SCTP export destinations configured:

```
Router# show ip flow export

Flow export v9 is enabled for main cache
  Exporting flows to 172.17.10.2 (100)
  Exporting flows to 172.16.45.57 (100) via SCTP
  Exporting using source interface Loopback0
  Version 9 flow records
  Cache for destination-prefix aggregation:
    Exporting flows to 192.168.247.198 (200) via SCTP
    Exporting using source IP address 172.16.254.254
  479 flows exported in 318 udp datagrams
  467 flows exported in 315 sctp messages
  0 flows failed due to lack of export packet
  159 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
```

Table 1 describes the significant fields shown in the display of the **show ip flow export** command.

Table 1 show ip flow export Field Descriptions

Field	Description
Exporting flows to	Indicates the export destinations and ports. The ports are in parentheses. **Note** When the export destination is configured with the NetFlow Reliable Transport Using SCTP feature the port number is followed by the text "via SCTP" in the display output.
Exporting using source IP address or Exporting using source interface	Indicates the source IP address or source interface. **Note** The source interface is used when you have configured the **ip flow-export source** *interface-type interface-number* command.
Version flow records	Displays the version of the flow records.
Cache for destination-prefix aggregation	Indicates the type of NetFlow aggregation caches that are configured. **Note** The indented lines below the name of the NetFlow aggregation cache indicate the export parameters that are configured for this cache.
flows exported in udp datagrams	Indicates the total number of export packets (datagrams) sent over UDP, and the total number of flows contained within them.
flows exported in sctp messages	Displays the total number of export packets (messages) sent over SCTP, and the total number of flows contained within them. **Note** SCTP is a message-oriented transport protocol. Therefore SCTP traffic is referred to as messages instead of datagrams.
flows failed due to lack of export packet	Indicates the number of flows that failed because no memory was available to create an export packet.
159 export packets were sent up to process level	The packet could not be processed by Cisco Express Forwarding or by fast switching.
export packets were dropped due to no fib export packets were dropped due to adjacency issues	Indicates the number of packets that Cisco Express Forwarding was unable to switch, or forward to the process level.
0 export packets were dropped due to fragmentation failures 0 export packets were dropped due to encapsulation fixup failures	Indicates the number of packets that were dropped because of problems constructing the IP packet.

(Specific values are not included with other fields in this table)

Table 1 show ip flow export Field Descriptions (continued)

Field	Description
0 export packets were dropped enqueuing for the RP	Indicates the number of times that there was a problem transferring the export packet between the RP and the line card.
0 export packets were dropped due to IPC rate limiting	(What does this have to do with why packets were dropped? Does not sound like it goes with the field description.)
0 export packets were dropped due to output drops	Indicates the number of times that the send queue was full while the packet was being sent.

The following is sample output from the **show ip flow export sctp** command with NetFlow SCTP export primary and backup SCTP export destinations configured for the NetFlow main cache and the NetFlow destination-prefix aggregation cache. The primary SCTP export destinations are active:

```
Router# show ip flow export sctp

IPv4 main cache exporting to 172.16.45.57, port 100, none
status: connected
backup mode: fail-over
912 flows exported in 619 sctp messages.
0 packets dropped due to lack of SCTP resources
fail-over time: 25 milli-seconds
restore time:   25 seconds
backup: 192.168.247.198, port 200
   status: not connected
   fail-overs: 2
   9 flows exported in 3 sctp messages.
   0 packets dropped due to lack of SCTP resources
destination-prefix cache exporting to 172.16.12.200, port 100, full
status: connected
backup mode: redundant
682 flows exported in 611 sctp messages.
0 packets dropped due to lack of SCTP resources
fail-over time: 25 milli-seconds
restore time:   25 seconds
backup: 192.168.247.198, port 200
   status: connected
   fail-overs: 8
   2 flows exported in 2 sctp messages.
   0 packets dropped due to lack of SCTP resources
```

The following is sample output from the **show ip flow export sctp** command with NetFlow SCTP export primary and backup SCTP export destinations configured for the NetFlow main cache and the NetFlow destination-prefix aggregation cache. The backup SCTP export destinations are active because the primary SCTP export destinations are unavailable.

```
Router# show ip flow export sctp

IPv4 main cache exporting to 172.16.45.57, port 100, none
status: fail-over
backup mode: fail-over
922 flows exported in 625 sctp messages.
0 packets dropped due to lack of SCTP resources
fail-over time: 25 milli-seconds
restore time:   25 seconds
backup: 192.168.247.198, port 200
   status: connected, active for 00:00:24
```

```
   fail-overs: 3
   11 flows exported in 4 sctp messages.
   0 packets dropped due to lack of SCTP resources
destination-prefix cache exporting to 172.16.12.200, port 100, full
status: fail-over
backup mode: redundant
688 flows exported in 617 sctp messages.
0 packets dropped due to lack of SCTP resources
fail-over time: 25 milli-seconds
restore time:   25 seconds
backup: 192.168.247.198, port 200
   status: connected, active for 00:00:00
   fail-overs: 13
   2 flows exported in 2 sctp messages.
   0 packets dropped due to lack of SCTP resources
Router#
```

Table 2 describes the significant fields shown in the display of the **show ip flow export sctp** and the **show ip flow export sctp verbose** commands.

Table 2 show ip flow export sctp Field Descriptions

Field	Description
IPv4 main cache exporting to 172.16.45.57, port 100, none	Indicates the type of cache, the IP address and port number used to reach the destination, and the level of reliability for the association: • IPv4 main cache—The type of NetFlow cache to which the display output applies. • 172.16.45.57—The IP address used for the SCTP export destination. • port 100—The SCTP port used for the SCTP export destination. • none—The level of reliability for this association. **Note** The reliability options are full and none.
status	The current state of each association. The states are: • initializing—The association is being established. • connected—The association is established. **Note** If this is a backup SCTP export destination configured for fail-over mode, you see an additional message indicating how long the association has been active. For example, active for 00:00:01. • not connected—The association will be established when the primary SCTP export backup destination is no longer available. • fail-over—The primary SCTP export destination is no longer available. The backup SCTP export destination is being used. • re-establishing—An association that has been active before is being reestablished.

Table 2 show ip flow export sctp Field Descriptions (continued)

Field	Description
backup mode	The backup mode of each association. The modes are: • redundant—The association is established (connected). **Note** The fact that the association is established does not mean that it is being used to export NetFlow data. • fail-over—The association will be established after the primary association fails.
flows exported in sctp messages	Indicates the total number of export packets (messages) sent over SCTP, and the total number of flows contained within them. **Note** SCTP is a message-oriented transport protocol. Therefore, SCTP traffic is referred to as messages instead of datagrams.
packets dropped due to lack of SCTP resources	The number of packets that were dropped due to lack of SCTP resources.
fail-over time: milli-seconds	The period of time that the networking device waits after losing connectivity to the primary SCTP export destination before attempting to use a backup SCTP export destination. **Note** This field is displayed when you use the **verbose** keyword after the **show ip flow export sctp** command.
restore time: seconds	The period of time that the networking device waits before reverting to the primary SCTP export destination after connectivity to it has been restored. **Note** This field is displayed when you use the **verbose** keyword after the **show ip flow export sctp** command.
backup: 192.168.247.198 port 200	The IP address and SCTP port used for the SCTP export backup destination. • 192.168.247.198—The IP address of the SCTP backup association. • port 200—The SCTP port used for the SCTP backup association.
fail-overs	The number of times that fail-over has occurred.
destination-prefix cache exporting to 172.16.12.200, port 100, full	Indicates the type of cache configures, the destination address and port number for the SCTP export, and the level of reliability for the association: • destination-prefix cache—The type of NetFlow aggregation cache configured. • 172.16.12.200—The IP address used for the SCTP export destination. • port 100—Indicates the SCTP port used for the SCTP export destination. • full—The level of reliability for this association,

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

```
Router# show ip flow export template
```

```
   Template Options Flag = 1
   Total number of Templates added = 4
   Total active Templates = 4
   Flow Templates active = 3
   Flow Templates added = 3
   Option Templates active = 1
   Option  Templates added = 1
   Template ager polls = 2344
   Option Template ager polls = 34
Main cache version 9 export is enabled
 Template export information
   Template timeout = 30
   Template refresh rate = 20
 Option export information
   Option timeout = 800
   Option refresh rate = 300
Aggregation cache destination-prefix version 9 export is enabled
 Template export information
   Template timeout = 30
   Template refresh rate = 20
 Option export information
   Option timeout = 30
   Option refresh rate = 20
```

Table 3 describes the significant fields shown in the display of the **show ip flow export template** command.

Table 3 show ip flow export template Field Descriptions

Field	Description
Template Options Flag	Identifies which options are enabled. The values are: • *0—No option* template configured. • 1—Version 9 option export statistics configured. • 2—Random sampler option template configured. • 4—Version 9 option export statistics for IPv6 configured.
Total number of Templates added	Indicates the number of Flow Templates and Option Templates that have been added since Version 9 export was first configured. This value in this field is the sum of the "Flow Templates added" and the "Option Templates added" fields. The value is incremented when a new template is created, because each template requires a unique ID.
Total active Templates	This is the sum of the values in the "Flow Templates active" and "Option Templates" active fields. The value in this field is incremented when a new data template or option template is created.

Table 3 show ip flow export template Field Descriptions (continued)

Field	Description
Flow Templates active	Indicates the number of (data) templates in use for Version 9 data export. When a new data template is created, this count, the "Total active Templates," the "Flow Templates added," and the "Total number of Templates added" counts are all incremented. **Note** When a data template is removed, only the "Flow Templates active" count and the "Total active Templates" count are decremented.
Flow Templates added	Indicates the number of Flow Templates and Option Templates that have been added since Version 9 export was first configured. The value is incremented when a new flow template is created, because each template requires a unique ID.
Option Templates active	Indicates the number of option templates which are currently in use for Version 9 options export. Configuring a new option increments this count and also the "Total active Templates," the "Option Templates added," and the "Total number of Templates added" counts. Removing (unconfiguring) an option decrements only the "Option Templates active" count and the "Total active Templates" count.
Option Templates added	Indicates the number of Option Templates that have been added since Version 9 export was first configured. The count is incremented when a new option template is created, because each template requires a unique ID.
Template ager polls	The number of times, since Version 9 export was configured, that the (data) template ager has run. The template ager checks up to 20 templates per invocation, resending any that need refreshed.
Option Template ager polls	The number of times, since Version 9 export was configured, that the option template ager has run. The template ager checks up to 20 templates per invocation, resending any that need refreshed.
Main cache version 9 export is enabled	NetFlow export Version 9 is enabled for the main NetFlow cache.
Template export information	Template timeout—The interval (in minutes) that the router waits after sending the templates (flow and options) before they are sent again. You can specify from 1 to 3600 minutes. The default is 30 minutes. • Template refresh rate—The number of export packets that are sent before the options and flow templates are sent again. You can specify from 1 to 600 packets. The default is 20 packets.

Table 3 show ip flow export template Field Descriptions (continued)

Field	Description
Option export information	• Option timeout—The interval (in minutes) that the router will wait after sending the options records before they are sent again. You can specify from 1 to 3600 minutes. The default is 30 minutes. • Option refresh rate—The number of packets that are sent before the configured options records are sent again. You can specify from 1 to 600 packets. The default is 20 packets.
Aggregation cache destination-prefix version 9 export is enabled	NetFlow export Version 9 is enabled for the NetFlow destination-prefix aggregation cache.

The following example displays the additional line in the **show ip flow export** command output when the **verbose** keyword is specified and MPLS PAL records are being exported to a NetFlow collector:

```
Router# show ip flow export verbose

Flow export v9 is enabled for main cache
  Exporting flows to 10.23.0.5 (4200)
  Exporting using source IP address 10.2.72.35
  Version 9 flow records
  Cache for destination-prefix aggregation:
    Exporting flows to 10.2.0.1 (4200)
    Exporting using source IP address 10.2.72.35
    182128 MPLS PAL records exported
  189305 flows exported in 6823 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 swat72f3#
```

The line of output added for the MPLS PAL records precedes the "*x* flows exported in *y* UDP datagrams" line. In this example, the additional line of output precedes "189305 flows exported in 6823 UDP datagrams."

The following example shows the sample output of the **show ip flow export** nbar command:

```
Router# show ip flow export nbar
  Nbar netflow is enabled
  10 nbar flows exported
  0 nbar flows failed to export due to lack of internal buffers
```

Related Commands

Command	Description
ip flow-export	Enables export of NetFlow accounting information in NetFlow cache entries.
show ip cache flow	Displays a summary of the NetFlow accounting statistics.
show ip cache verbose flow	Displays a detailed summary of the NetFlow accounting statistics.
show ip flow interface	Displays the NetFlow accounting configuration on interfaces.
show mpls flow mappings	Displays the full MPLS PAL table.

clear ip flow stats

To clear the NetFlow accounting statistics, use the **clear ip flow stats** command in privileged EXEC mode.

clear ip flow stats [**nbar**]

Syntax Description

nbar	(Optional) Clears Network Based Application Recognition (NBAR) NetFlow statistics.

Defaults

This command has no default settings.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
11.1CA	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(14)SX	Support for this command was introduced on the Supervisor Engine 720.
12.2(17d)SXB	Support for this command on the Supervisor Engine 2 was extended to the 12.2(17d)SXB release.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(18)SXF	This command was integrated into Cisco IOS Release 12.2(18)SXF.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(18)ZYA2	This command was modified. The **nbar** keyword was added.

Usage Guidelines

You must have NetFlow accounting configured on your router before you can use this command.

The **show ip cache flow** command displays the NetFlow accounting statistics. Use the **clear ip flow stats** command to clear the NetFlow accounting statistics.

Examples

The following example shows how to clear the NetFlow accounting statistics on the router:

```
Router# clear ip flow stats
```

Related Commands

Command	Description
show ip cache flow	Displays a summary of the NetFlow accounting statistics.
show ip cache verbose flow	Displays a detailed summary of the NetFlow accounting statistics.

(Delete extra space)

Command	Description
show ip flow interface	Displays NetFlow accounting configuration for interfaces.
show ip interface	Displays the usability status of interfaces configured for IP.