



CHAPTER 1

Configuring NetFlow



Note

Supervisor Engine 6-E and Catalyst 4900M chassis do not support Netflow; it is only supported on Supervisor Engine IV, Supervisor Engine V, Supervisor Engine V-10GE, or WS-F4531.



Note

A potential workaround to the lack of Netflow support on the Supervisor Engine 6-E is to install a Cisco Netflow Generation Appliance 3140. You can learn more about this product here:
http://www.cisco.com/en/US/prod/collateral/netmgts/ps6504/ps12508/ps12512/data_sheet_c78-706041.html

This chapter describes how to configure NetFlow Statistics on the Catalyst 4500 series switches. It also provides guidelines, procedures, and configuration examples.



Note

To use the NetFlow feature, you must have the Supervisor Engine V-10GE (the functionality is embedded in the supervisor engine), or the NetFlow Services Card (WS-F4531) and either a Supervisor Engine IV or a Supervisor Engine V.



Note

For complete syntax and usage information for the switch commands used in this chapter, first look at the *Cisco Catalyst 4500 Series Switch Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products//hw/switches/ps4324/index.html>

If the command is not found in the *Catalyst 4500 Command Reference*, it will be found in the larger Cisco IOS library. Refer to the *Cisco IOS Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/ps6350/index.html>



Note

Refer to the *NetFlow Solutions Guide* for more detailed information on NetFlow usage and management.

The following topics are included:

- [About NetFlow Statistics Collection, page 1-2](#)
- [Configuring NetFlow Statistics Collection, page 1-6](#)

- [NetFlow Statistics Collection Configuration Example, page 1-13](#)
- [NetFlow Configuration Examples, page 1-14](#)

About NetFlow Statistics Collection

A network flow is defined as a unidirectional stream of packets between a given source and destination—both defined by a network-layer IP address and transport-layer port number. Specifically, a flow is identified as the combination of the following fields: source IP address, destination IP address, source port number, destination port number, protocol type, type of service, and input interface.

NetFlow Statistics is a global traffic monitoring feature that allows flow-level monitoring of all IPv4-routed traffic through the switch using NetFlow Data Export (NDE). Collected statistics can be exported to an external device (NetFlow Collector/Analyzer) for further processing. Network planners can selectively enable NetFlow Statistics (and NDE) on a per-device basis to gain traffic performance, control, or accounting benefits in specific network locations.

NetFlow exports flow information in UDP datagrams in one of two formats. The version 1 format was the initial released version, and version 5 is a later enhancement to add Border Gateway Protocol (BGP) autonomous system (AS) information and flow sequence numbers. In version 1 and version 5 format, the datagram consists of a header and one or more flow records. The first field of the header contains the version number of the export datagram.

This section contains the following subsections:

- [Information Derived from Hardware, page 1-4](#)
- [Information Derived from Software, page 1-4](#)
- [Assigning the Input and Output Interface and AS Numbers, page 1-4](#)
- [Feature Interaction of NetFlow Statistics with UBRL and Microflow Policing, page 1-5](#)
- [VLAN Statistics, page 1-6](#)

NDE Versions

The Catalyst 4500 series switch supports NDE versions 1 and 5 for the captured statistics. NetFlow aggregation requires NDE version 8.

Depending on the current flow mask, some fields in the flow records might not have values. Unsupported fields contain a zero (0).

The following tables describe the supported fields for NDE version 5:

- [Table 1-1](#)—Version 5 header format
- [Table 1-2](#)—Version 5 flow record format

Table 1-1 NDE Version 5 Header Format

| Bytes | Content | Description |
|-------|-----------|--|
| 0–1 | version | NetFlow export format version number |
| 2–3 | count | Number of flows exported in this packet (1–30) |
| 4–7 | SysUptime | Current time in milliseconds since the switch booted |
| 8–11 | unix_secs | Current seconds since 0000 UTC 1970 |

Table 1-1 NDE Version 5 Header Format (continued)

| Bytes | Content | Description |
|-------|---------------|--|
| 12–15 | unix_nsecs | Residual nanoseconds since 0000 UTC 1970 |
| 16–19 | flow_sequence | Sequence counter of total flows seen |
| 20–21 | engine_type | Type of flow switching engine |
| 21–23 | engine_id | Slot number of the flow switching engine |

Table 1-2 NDE Version 5 Flow Record Format

| Bytes | Content | Description | Flow masks: • X=Populated • A=Additional field | | | | | |
|-------|-----------|--|--|----------------|--------------------|------------------------------|----------------|----------------|
| | | | Source | Destination | Destination Source | Destination Source Interface | Full | Full Interface |
| 0–3 | srcaddr | Source IP address | X | | X | X | X | X |
| 4–7 | dstaddr | Destination IP address | | X | X | X | X | X |
| 8–11 | nexthop | Next hop switch's IP address | | A ¹ | A | A | A | A |
| 12–13 | input | Ingress interface SNMP ifIndex | | | | X | | X |
| 14–15 | output | Egress interface SNMP ifIndex | | A ¹ | A | A | A | A |
| 16–19 | dPkts | Packets in the flow | X | X | X | X | X | X |
| 20–23 | dOctets | Octets (bytes) in the flow | X | X | X | X | X | X |
| 24–27 | first | SysUptime at start of the flow | X | X | X | X | X | X |
| 28–31 | last | SysUptime at the time the last packet of the flow was received | X | X | X | X | X | X |
| 32–33 | srcport | Layer 4 source port number or equivalent | | | | | X ² | X ² |
| 34–35 | dstport | Layer 4 destination port number or equivalent | | | | | X | X |
| 36 | pad1 | Unused (zero) byte | | | | | | |
| 37 | tcp_flags | Cumulative OR of TCP flags | | | | | | |
| 38 | prot | Layer 4 protocol (for example, 6=TCP, 17=UDP) | | | | | X | X |
| 39 | tos | IP type-of-service byte | | | | | | |
| 40–41 | src_as | Autonomous system number of the source, either origin or peer | X | | X | X | X | X |
| 42–43 | dst_as | Autonomous system number of the destination, either origin or peer | | X | X | X | X | X |
| 44–45 | src_mask | Source address prefix mask bits | X | | X | X | X | X |
| 46–47 | dst_mask | Destination address prefix mask bits | | X | X | X | X | X |
| 48 | pad2 | Pad 2 is unused (zero) bytes | | | | | | |

1. With the destination flow mask, the “Next hop switch’s IP address” field and the “Output interface’s SNMP ifIndex” field might not contain information that is accurate for all flows.
2. In PFC3BXL or PFC3B mode, ICMP traffic contains the ICMP code and type values.

Information Derived from Hardware

Information available in a typical NetFlow record from hardware includes the following:

- Packet and byte counts
- Start and end timestamps
- Source and destination IP addresses
- IP protocol
- Source and destination port numbers

Information Derived from Software

Information available in a typical NetFlow record from software includes the following:

- Input and output identifiers
- Routing information, including next-hop address, origin and peer AS, source and destination prefix mask

Assigning the Input and Output Interface and AS Numbers

The following topics are discussed:

- [Assigning the Inferred Fields, page 1-4](#)
- [Assigning the Output Interface and Output-Related Inferred Fields, page 1-4](#)
- [Assigning the Input Interface and Input-Related Inferred Fields, page 1-5](#)

Assigning the Inferred Fields

The Catalyst 4500 series switch collects netflow flows in hardware. The hardware collects a sub-set of all the NetFlow flow fields. The rest of the fields are then filled in by the software when the software examines the routing state.

The NetFlow Services Card does not provide enough information to accurately and consistently determine the input interface, output interface and other routing information associated with NetFlow flows. The Catalyst 4500 series switch has a software mechanism to compensate for this. The mechanism is described in the next paragraph.

Assigning the Output Interface and Output-Related Inferred Fields

Software determines the output interface information by looking up the Forwarding Information Base (FIB) entry in the default FIB table (based on the destination IP address). From this FIB entry, the software gains access to the destination AS number for this destination IP address, as well as the appropriate adjacency that stores the interface information. The output interface is based solely on the

destination IP address. If load balancing is enabled on the switch, the load balancing hash, instead of looking at the adjacency in the FIB entry, is applied to access the appropriate FIB path and access the appropriate adjacency. Although this process typically yields correct results, an inaccuracy can occur when using a PBR that shares IP addresses with the default FIB table. Under these circumstances, there would then be multiple FIB table entries and associated adjacencies for the same destination IP address.

Assigning the Input Interface and Input-Related Inferred Fields

Similarly, the input interface and the source AS number for the source IP address are determined by looking up the FIB entry in the default FIB table based on the source IP address. The input interface is based solely on the source IP address and a reverse lookup is done to determine to which interface a packet with this IP destination address needs to be routed. This process assumes that the forwarding paths are symmetrical. However, if this process yields multiple input interfaces, a deterministic algorithm is applied to pick one of them the one with the lowest IP address. Although this process typically yields correct values, there are scenarios where the values are inaccurate:

- If load balancing is being applied by an upstream adjacent switch, one input interface must be chosen arbitrarily out of the multiple input interfaces available. This action is necessary because the input interface that would be used depends on the type of load balancing algorithm being deployed by the adjacent upstream switch. It is not always feasible to know the algorithm. All flow statistics are attributed to one input interface. Software selects the interface with the lowest IP subnet number.
- In an asymmetric routing scheme in which the traffic for an IP subnet might be received on one interface and sent on another, the inferences noted previously for selecting an input interface, based on a reverse lookup, would be incorrect and cannot be verified.
- If PBR or VRF is enabled on the switch and the flow is destined to an address that resides in the PBR or VRF range or is sourced from an address that resides in the PBR or VRF range, the information is incorrect. The input and output interface most likely points to the default route (if configured) or have no value at all (NULL)
- If VRF is enabled on the switch on some interfaces and the flow comes from a VRF interface, the information is incorrect. The input and output interface most likely points to the default route (if configured) or have no value (NULL).



Note

The Supervisor Engine V-10GE provides the input interface information by using hardware, improving the accuracy of NetFlow information.

Feature Interaction of NetFlow Statistics with UBRL and Microflow Policing

On systems with Supervisor Engine V-10GE, there is a feature interaction between NetFlow Statistics and UBRL (User Based Rate Limiting). As part of correctly configuring UBRL on a given interface, the class map must specify a flow mask. In turn, this flow mask is used to create hardware-based NetFlow statistics for the flow. By default, for traditional full-flow NetFlow statistics, the full-flow mask is used. With UBRL, however, the masks can differ. If UBRL is configured on a given interface, the statistics are collected based on the mask configured for UBRL. Consequently, the system does not collect full-flow statistics for traffic transiting an interface configured with UBRL. For more details, refer to the [“Configuring User-Based Rate-Limiting”](#) section on page 1-37.

VLAN Statistics

With NetFlow support, you can report Layer 2 output VLAN statistics, as well as VLAN statistics for routed traffic in and out of a VLAN.

The following example shows the CLI output for a specific VLAN:

```
Switch# show vlan counters or show vlan id 22 count
* Multicast counters include broadcast packets
Vlan Id :22
L2 Unicast Packets :38
L2 Unicast Octets :2432
L3 Input Unicast Packets :14344621
L3 Input Unicast Octets :659852566
L3 Output Unicast Packets :8983050
L3 Output Unicast Octets :413220300
L3 Output Multicast Packets :0
L3 Output Multicast Octets :0
L3 Input Multicast Packets :0
L3 Input Multicast Octets :0
L2 Multicast Packets :340
L2 Multicast Octets :21760
```


Note

NetFlow support has hardware limitations that restrict the platform support to a subset of all NetFlow fields. Specifically, TCP Flags and the ToS byte (DSCP) are not supported.

Configuring NetFlow Statistics Collection

To configure NetFlow switching, complete the tasks in these sections:

- [Checking for Required Hardware, page 1-6](#)
- [Enabling NetFlow Statistics Collection, page 1-7](#)
- [Configuring Switched/Bridged IP Flows, page 1-8](#)
- [Exporting NetFlow Statistics, page 1-9](#)
- [Managing NetFlow Statistics Collection, page 1-9](#)
- [Configuring an Aggregation Cache, page 1-10](#)
- [Configuring a NetFlow Minimum Prefix Mask for Router-Based Aggregation, page 1-11](#)
- [Configuring NetFlow Aging Parameters, page 1-12](#)

Checking for Required Hardware

To ensure that the necessary hardware is enabled, enter the **show module** command, as follows:

```
Switch# show module all
Chassis Type : WS-C4507R

Power consumed by backplane : 40 Watts

Mod Ports Card Type                               Model
Serial No.
-----+-----+-----+-----+-----+-----+-----+-----
```

```

1      2  1000BaseX (GBIC) Supervisor(active)   WS-X4515
JAB062604KB
2      2  1000BaseX (GBIC) Supervisor(standby) WS-X4515
JAB062408CB
6     48 10/100BaseTX (RJ45)                   WS-X4148
JAB032305UH

```

```

M MAC addresses                Hw  Fw          Sw          Status
-----+-----+-----+-----+-----
1 0001.6442.2c00 to 0001.6442.2c01 0.4 12.1(14r)EW( 12.1(20030513:00 Ok
2 0001.6442.2c02 to 0001.6442.2c03 0.4 12.1(14r)EW( 12.1(20030513:00 Ok
6 0050.3ed8.6780 to 0050.3ed8.67af 1.6 12.1(14r)EW( 12.1(20030513:00 Ok

```

```

Mod  Submodule                Model          Serial No.    Hw  Status
-----+-----+-----+-----+-----
1   Netflow Services Card    WS-F4531      JAB062209CG  0.2  Ok
2   Netflow Services Card    WS-F4531      JAB062209AG  0.2  Ok

```

Switch#

**Note**

Enabling this feature does not impact the hardware-forwarding performance of the switch.

The effective size of the hardware flow cache table is 65,000 flows. (The hardware flow cache for the Supervisor Engine V-10GE is 85,000 flows.) If more than 85,000 flows are active simultaneously, statistics may be lost for some of the flows.

The effective size of the software flow table is 256, 000 flows. The NetFlow software manages the consistency between the hardware and software tables, keeping the hardware table open by purging inactive hardware flows to the software table.

User-configured timeout settings dictate when the flows are purged and exported through NDE from the software cache. Hardware flow management ensures consistency between hardware flow purging and the user-configured timeout settings.

Software-forwarded flows are also monitored. Moreover, statistics overflow if any flow receives traffic at a sustained rate exceeding 2 gigabits per second. Generally, this situation should not occur because a port cannot transmit at a rate higher than 1 gigabit per second.

**Note**

By design, even if the timeout settings are high, flows automatically “age out” as they approach their statistics limit.

**Note**

First and last-seen (start and end times) flow timestamp accuracy is within 3 seconds.

Enabling NetFlow Statistics Collection

**Note**

NetFlow Flow Statistics are disabled by default.

To enable NetFlow switching, first configure the switch for IP routing as described in the IP configuration chapters in the *Cisco IOS IP and IP Routing Configuration Guide*. After you configure IP routing, perform one of these tasks:

| Command | Purpose |
|---|--|
| Switch(config)# ip flow ingress | Enables NetFlow for IP routing. |
| Switch(config)# ip flow ingress infer-fields | Enables NetFlow with inferred input/output interfaces and source/destination BGP as information. The inter-fields option must be configured for AS information to be determined. |

Configuring Switched/Bridged IP Flows

NetFlow is defined as a collection of routed IP flows created and tracked for all routed IP traffic. In switching environments, considerable IP traffic is switched within a VLAN and is not routed. This traffic is termed *switched/bridged IP traffic*; the associated flow is termed *switched/bridged IP flows*. NetFlow hardware is capable of creating and tracking this type of flow. The NetFlow Switched IP Flows feature allows you to create, track, and export switched IP flows (that is, it creates and tracks flows for IP traffic that is being switched and not routed).

Be aware of the following:

- Switched IP flow collection cannot be enabled in isolation on Catalyst 4500 series switches. You need to enable both routed flow and switched flow collection to start collecting switched IP flows.
- Generally, the input and output interface information are NULL. If the traffic is being switched on a VLAN that is associated with an SVI, the input and output interface information points to the same Layer 3 interface.
- Switched flows are exported according to regular export configurations; a separate export CLI does not exist.
- In the main cache, because of a hardware limitation, IP flows and routed IP flows are indistinguishable.



Note

To enable switched IP flow collection on all interfaces, you need to enter both the **ip flow ingress** and **ip flow ingress layer2-switched** commands.



Note

To enable a user-based rate limiting policy on the switched IP flow traffic, you need to enter the **ip flow ingress layer2-switched** command, but not the **ip flow ingress** command. See “Configuring User-Based Rate-Limiting” on page 37.

To configure the NetFlow cache and enable switched IP flow collection, perform this task:

| | Command | Purpose |
|---------------|--|-----------------------------------|
| Step 1 | Switch# conf terminal | Enters configuration mode. |
| Step 2 | Switch(config)# ip flow ingress | Enables routed flow collection. |
| Step 3 | Switch(config)# ip flow ingress layer2-switched | Enables switched flow collection. |

This example shows how to display the contents of an IP flow cache that contains switch IP flows:

```
Switch# show ip cache flow
IP Flow Switching Cache, 17826816 bytes
 2 active, 262142 inactive, 2 added
 6 ager polls, 0 flow alloc failures
 Active flows timeout in 30 minutes
 Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 1081480 bytes
 2 active, 65534 inactive, 2 added, 2 added to flow
 0 alloc failures, 0 force free
 1 chunk, 1 chunk added
 last clearing of statistics never
Protocol          Total      Flows    Packets Bytes  Packets Active(Sec) Idle(Sec)
-----          Flows      /Sec     /Flow  /Pkt   /Sec     /Flow     /Flow

SrcIf             SrcIPAddress  DstIf             DstIPAddress      Pr SrcP DstP  Pkts
Fal               150.1.1.1     Fal               13.1.1.1           11 003F 003F 425K
Fal               13.1.1.1     Fal               150.1.1.1          11 003F 003F 425K
Switch#
```

Exporting NetFlow Statistics

To configure the switch to export NetFlow Statistics to a workstation when a flow expires, perform one of these tasks:

| Command | Purpose |
|---|--|
| Switch(config)# ip flow-export destination { <i>hostname</i> <i>ip-address</i> } <i>udp-port</i> | (Required) Configures the switch to export NetFlow cache entries to a specific destination (for example, a workstation). Note You can specify multiple destinations. |
| Switch(config)# ip flow-export version {1 {5 [origin-as peer-as]}} | (Optional) Configures the switch to export NetFlow cache entries to a workstation if you are using receiving software that requires version 1 or 5. Version 1 is the default. origin-as causes NetFlow to determine the origin BGP autonomous system of both the source and the destination hosts of the flow. peer-as causes NetFlow to determine the peer BGP autonomous system of both the input and output interfaces of the flow. |
| Switch(config)# ip flow-export source <i>interface</i> | (Optional) Specifies an interface whose IP address is used as the source IP address in the IP header of the NetFlow Data Export (NDE) packet. Default is the NDE output interface. |

Managing NetFlow Statistics Collection

You can display and clear NetFlow Statistics, including IP flow switching cache information and flow information, such as the protocol, total flow, flows per second, and so forth. You can also use the resulting information to obtain information about your switch traffic.

To manage NetFlow switching statistics, perform one or both of following tasks:

| Command | Purpose |
|------------------------------------|--|
| Switch# show ip cache flow | Displays the NetFlow switching statistics. |
| Switch# clear ip flow stats | Clears the NetFlow switching statistics. |

Configuring an Aggregation Cache

Aggregation of NetFlow Statistics is typically performed by NetFlow collection tools on management workstations. By extending this support to the Catalyst 4500 series switch, you can do the following:

- Reduce the required bandwidth between the switch and workstations, because fewer NDE packets are exported.
- Reduce the number of collection workstations required.
- Provide visibility to aggregated flow statistics at the CLI.

To configure an aggregation cache, you must enter the aggregation cache configuration mode, and you must decide which type of aggregation scheme you would like to configure: autonomous system, destination prefix, protocol prefix, or source prefix aggregation cache. Once you define the aggregation scheme, define the operational parameters for that scheme. More than one aggregation cache can be configured concurrently.

To configure an aggregation cache, perform this task:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | Switch(config)# ip flow-aggregation cache as | Enters aggregation cache configuration mode and enables an aggregation cache scheme (autonomous system, destination-prefix, prefix, protocol-port, or source-prefix). |
| Step 2 | Switch(config-flow-cache)# cache timeout inactive 199 | Specifies the number of seconds (in this example, 199) in which an inactive entry is allowed to remain in the aggregation cache before it is deleted. |
| Step 3 | Switch(config-flow-cache)# cache timeout active 45 | Specifies the number of minutes (in this example, 45) in which an active entry is active. |
| Step 4 | Switch(config-flow-cache)# export destination 10.42.41.1 9991 | Enables the data export. |
| Step 5 | Switch(config-flow-cache)# enabled | Enables aggregation cache creation. |

Verifying Aggregation Cache Configuration and Data Export

To verify the aggregation cache information, perform this task:

| Command | Purpose |
|--|---|
| Switch# show ip cache flow aggregation destination-prefix | Displays the specified aggregation cache information. |

To confirm data export, perform the following task:

| Command | Purpose |
|--|--|
| Switch# <code>show ip flow export</code> | Displays the statistics for the data export including the main cache and all other enabled caches. |

Configuring a NetFlow Minimum Prefix Mask for Router-Based Aggregation

The minimum prefix mask specifies the shortest subnet mask that is used for aggregating flows within one of the IP-address based aggregation caches (source-prefix, destination-prefix, prefix). In these caches, flows are aggregated based upon the IP address (source, destination, or both, respectively) and masked by the longer of the Minimum Prefix mask and the subnet mask of the route to the source/destination host of the flow (as found in the switch routing table).



Note

The default value of the minimum mask is zero. The configurable range for the minimum mask is from 1 to 32. You should choose an appropriate value depending on the traffic. A higher value for the minimum mask provides more detailed network addresses, but it may also result in increased number of flows in the aggregation cache.

To configure a minimum prefix mask for the Router-Based Aggregation feature, perform the tasks described in the following sections. Each task is optional.

- [Configuring the Minimum Mask of a Prefix Aggregation Scheme, page 1-11](#)
- [Configuring the Minimum Mask of a Destination-Prefix Aggregation Scheme, page 1-11](#)
- [Configuring the Minimum Mask of a Source-Prefix Aggregation Scheme, page 1-12](#)
- [Monitoring and Maintaining Minimum Masks for Aggregation Schemes, page 1-12](#)

Configuring the Minimum Mask of a Prefix Aggregation Scheme

To configure the minimum mask of a prefix aggregation scheme, perform this task:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | Switch(config)# <code>ip flow-aggregation cache prefix</code> | Configures the prefix aggregation cache. |
| Step 2 | Switch(config-flow-cache)# <code>mask source minimum value</code> | Specifies the minimum value for the source mask. |
| Step 3 | Switch(config-flow-cache)# <code>mask destination minimum value</code> | Specifies minimum value for the destination mask. |

Configuring the Minimum Mask of a Destination-Prefix Aggregation Scheme

To configure the minimum mask of a destination-prefix aggregation scheme, perform this task:

| | Command | Purpose |
|--------|---|---|
| Step 1 | Switch(config)# ip flow-aggregation cache destination-prefix | Configures the destination aggregation cache. |
| Step 2 | Switch(config-flow-cache)# mask destination minimum value | Specifies the minimum value for the destination mask. |

Configuring the Minimum Mask of a Source-Prefix Aggregation Scheme

To configure the minimum mask of a source-prefix aggregation scheme, perform this task:

| | Command | Purpose |
|--------|--|--|
| Step 1 | Switch(config)# ip flow-aggregation cache source-prefix | Configures the source-prefix aggregation cache. |
| Step 2 | Switch(config-flow-cache)# mask source minimum value | Specifies the minimum value for the source mask. |

Monitoring and Maintaining Minimum Masks for Aggregation Schemes

To view the configured value of the minimum mask, use the following commands for each aggregation scheme, as needed:

| Command | Purpose |
|--|---|
| Switch# show ip cache flow aggregation prefix | Displays the configured value of the minimum mask in the prefix aggregation scheme. |
| Switch# show ip cache flow aggregation destination-prefix | Displays the configured value of the minimum mask in the destination-prefix aggregation scheme. |
| Switch# show ip cache flow aggregation source-prefix | Displays the configured value of the minimum mask in the source-prefix aggregation scheme. |

Configuring NetFlow Aging Parameters

You can control when flows are purged from the software flow cache (and, if configured, reported through NDE) with the configuration aging parameters, **Active** and **Inactive**, of the **ip flow-cache timeout** command.

Active Aging specifies the period of time in which a flow should be removed from the software flow cache after the flow is created. Generally, this parameter is used to periodically notify external collection devices about active flows. This parameter operates independently of existing traffic on the flow. Active timeout settings are usually measured in minutes (default is 30 minutes).

Inactive Aging specifies how long to wait before removing a flow after the last packet is seen. The inactive parameter clears the flow cache of stale flows, which prevents new flows from starving (due to lack of resources). Inactive timeout settings are generally on the order of seconds (default is 15 seconds).

NetFlow Statistics Collection Configuration Example

The following example shows how to modify the configuration to enable NetFlow switching. It also shows how to export the flow statistics for further processing to UDP port 9991 on a workstation with the IP address of 40.0.0.2. In this example, existing NetFlow Statistics are cleared, which ensures that the **show ip cache flow** command displays an accurate summary of the NetFlow switching statistics:

```
Switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip route-cache flow
Switch(config)# ip flow-export destination 40.0.0.2 9991
Switch(config)# ip flow-export version 5
Switch(config)# end
Switch# show ip flow export
Flow export is enabled
  Exporting flows to 40.0.0.2 (9991)
  Exporting using source IP address 40.0.0.1
  Version 5 flow records
  2 flows exported in 1 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
Switch#

Switch# show ip cache flow

IP Flow Switching Cache, 17826816 bytes
  69 active, 262075 inactive, 15087 added
  4293455 aged polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 1081480 bytes
  0 active, 65536 inactive, 0 added, 0 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
  last clearing of statistics never

Protocol      Total      Flows      Packets Bytes  Packets Active(Sec) Idle(Sec)
-----      -
              Flows      /Sec       /Flow /Pkt   /Sec     /Flow     /Flow
TCP-Telnet    28         0.0        167   40    0.0     20.9     11.9
TCP-other     185        0.0         2    48    0.0     6.2      15.4
UDP-DNS       4          0.0         1    61    0.0     0.0      15.5
UDP-other    13466      0.0    3396586   46   91831.3  139.3    15.9
ICMP         97         0.0         2    95    0.0     2.3      15.4
IGMP          1          0.0         2    40    0.0     0.9      15.1
IP-other     1120       0.0    38890838   46   87453.0  1354.5   24.0
Total:       14901      0.0    5992629   46  179284.3  227.8    16.5

SrcIf          SrcIPAddress  DstIf          DstIPAddress  Pr SrcP DstP  Pkts
-----
SrcIf          SrcIPAddress  DstIf          DstIPAddress  Pr SrcP DstP  Pkts
Gi6/2          30.20.1.18   Gi6/1          30.10.1.18    11 4001 4001  537K
Gi6/2          30.20.1.19   Gi6/1          30.10.1.19    11 4001 4001  537K
Gi6/2          30.20.1.16   Gi6/1          30.10.1.16    11 4001 4001  537K
Gi6/2          30.20.1.17   Gi6/1          30.10.1.17    11 4001 4001  537K
Gi6/2          30.20.1.20   Gi6/1          30.10.1.20    11 4001 4001  537K
Gi6/2          30.20.1.10   Gi6/1          30.10.1.10    11 4001 4001  539K
```

```

Gi6/2      30.20.1.11      Gi6/1      30.10.1.11      11 4001 4001 539K
Gi6/2      30.20.1.14      Gi6/1      30.10.1.14      11 4001 4001 539K
Gi6/2      30.20.1.15      Gi6/1      30.10.1.15      11 4001 4001 539K
Gi6/2      30.20.1.12      Gi6/1      30.10.1.12      11 4001 4001 539K
Gi6/2      30.20.1.13      Gi6/1      30.10.1.13      11 4001 4001 539K
Gi5/48     171.69.23.149   Local      172.20.64.200   06 8214 0017 759
Gi6/1      30.10.1.12      Gi6/2      30.20.1.12      11 4001 4001 539K
Gi6/1      30.10.1.13      Gi6/2      30.20.1.13      11 4001 4001 539K
Gi6/1      30.10.1.14      Gi6/2      30.20.1.14      11 4001 4001 539K
Gi6/1      30.10.1.15      Gi6/2      30.20.1.15      11 4001 4001 539K
Gi6/1      30.10.1.10      Gi6/2      30.20.1.10      11 4001 4001 539K
Gi6/1      30.10.1.11      Gi6/2      30.20.1.11      11 4001 4001 539K
Gi6/1      30.10.1.20      Gi6/2      30.20.1.20      11 4001 4001 537K
Gi6/1      30.10.1.16      Gi6/2      30.20.1.16      11 4001 4001 537K
Gi6/1      30.10.1.17      Gi6/2      30.20.1.17      11 4001 4001 537K
Gi6/1      30.10.1.18      Gi6/2      30.20.1.18      11 4001 4001 537K
Gi6/1      30.10.1.19      Gi6/2      30.20.1.19      11 4001 4001 537K
Switch#

```

NetFlow Configuration Examples

This section provides the following basic configuration examples:

- [NetFlow Enabling Scheme Examples, page 1-14](#)
- [NetFlow Aggregation Configuration Examples, page 1-14](#)
- [NetFlow Minimum Prefix Mask Router-Based Aggregation Scheme Examples, page 1-16](#)

NetFlow Enabling Scheme Examples



Note

Enabling NetFlow on a per-interface basis is not supported on a Catalyst 4500 switch.

This example shows how to enable NetFlow globally:

```

Switch# configure terminal
Switch(config)# ip flow ingress

```

This example shows how to enable NetFlow with support for inferred fields:

```

Switch# configure terminal
Switch(config)# ip flow ingress infer-fields

```

NetFlow Aggregation Configuration Examples

This section provides the following aggregation cache configuration examples:

- [Autonomous System Configuration, page 1-15](#)
- [Destination Prefix Configuration, page 1-15](#)
- [Prefix Configuration, page 1-15](#)
- [Protocol Port Configuration, page 1-15](#)
- [Source Prefix Configuration, page 1-16](#)

Autonomous System Configuration

This example shows how to configure an autonomous system aggregation cache with an inactive timeout of 200 seconds, a cache active timeout of 45 minutes, an export destination IP address of 10.42.42.1, and a destination port of 9992:

```
Switch(config)# ip flow-aggregation cache as
Switch(config-flow-cache)# cache timeout inactive 200
Switch(config-flow-cache)# cache timeout active 45
Switch(config-flow-cache)# export destination 10.42.42.1 9992
Switch(config-flow-cache)# enabled
```

Destination Prefix Configuration

This example shows how to configure a destination prefix aggregation cache with an inactive timeout of 200 seconds, a cache active timeout of 45 minutes, an export destination IP address of 10.42.42.1, and a destination port of 9992:

```
Switch(config)# ip flow-aggregation cache destination-prefix
Switch(config-flow-cache)# cache timeout inactive 200
Switch(config-flow-cache)# cache timeout active 45
Switch(config-flow-cache)# export destination 10.42.42.1 9992
Switch(config-flow-cache)# enabled
```

Prefix Configuration

This example shows how to configure a prefix aggregation cache with an inactive timeout of 200 seconds, a cache active timeout of 45 minutes, an export destination IP address of 10.42.42.1, and a destination port of 9992:

```
Switch(config)# ip flow-aggregation cache prefix
Switch(config-flow-cache)# cache timeout inactive 200
Switch(config-flow-cache)# cache timeout active 45
Switch(config-flow-cache)# export destination 10.42.42.1 9992
Switch(config-flow-cache)# enabled
```

Protocol Port Configuration

This example shows how to configure a protocol port aggregation cache with an inactive timeout of 200 seconds, a cache active timeout of 45 minutes, an export destination IP address of 10.42.42.1, and a destination port of 9992:

```
Switch(config)# ip flow-aggregation cache protocol-port
Switch(config-flow-cache)# cache timeout inactive 200
Switch(config-flow-cache)# cache timeout active 45
Switch(config-flow-cache)# export destination 10.42.42.1 9992
Switch(config-flow-cache)# enabled
```

Source Prefix Configuration

This example shows how to configure a source prefix aggregation cache with an inactive timeout of 200 seconds, a cache active timeout of 45 minutes, an export destination IP address of 10.42.42.1, and a destination port of 9992:

```
Switch(config)# ip flow-aggregation cache source-prefix
Switch(config-flow-cache)# cache timeout inactive 200
Switch(config-flow-cache)# cache timeout active 45
Switch(config-flow-cache)# export destination 10.42.42.1 9992
Switch(config-flow-cache)# enabled
```

NetFlow Minimum Prefix Mask Router-Based Aggregation Scheme Examples

This section provides examples for the NetFlow minimum prefix mask aggregation cache configuration:

- [Prefix Aggregation Scheme, page 1-16](#)
- [Destination-Prefix Aggregation Scheme, page 1-16](#)
- [Source-Prefix Aggregation Scheme, page 1-16](#)

Prefix Aggregation Scheme

it is an example of a prefix aggregation cache configuration:

```
!
ip flow-aggregation cache prefix
mask source minimum 24
mask destination minimum 28
```

In this example, assume the following configuration:

```
ip route 118.42.20.160 255.255.255.224 110.42.13.2
ip route 122.16.93.160 255.255.255.224 111.22.21.2
```

Both routes have a 27-bit subnet mask in the routing table on the switch.

Flows travelling from the 118.42.20.160 subnet to the 122.16.93.160 subnet whose source IP addresses match with a mask of 27 bits and whose destination IP addresses match with a mask of 28 bits are aggregated together in the cache statistics.

Destination-Prefix Aggregation Scheme

it is an example of a destination-prefix aggregation cache configuration:

```
!
ip flow-aggregation cache destination-prefix
mask destination minimum 32
!
```

Source-Prefix Aggregation Scheme

it is an example of a source-prefix aggregation cache configuration:

```
ip flow-aggregation cache source-prefix
mask source minimum 30
```