



Configuring NetFlow Statistics Collection

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



Note

This feature is only available if the NetFlow Services Card (WS-F4531) is present.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference* and related publications at <http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/index.htm>. Refer to the “NetFlow Solutions Guide” for more detailed information on NetFlow usage and management.

The following topics are included:

- [Overview of NetFlow Statistics Collection, page 30-1](#)
- [Checking for Required Hardware, page 30-4](#)
- [Caveat for the NetFlow Feature, page 30-3](#)
- [Configuring NetFlow Statistics Collection, page 30-4](#)
- [Configuring Netflow Aging Parameters, page 30-9](#)
- [NetFlow Statistics Collection Configuration Example, page 30-9](#)
- [NetFlow Configuration Examples, page 30-10](#)

Overview of NetFlow Statistics Collection

A network flow is defined as a unidirectional sequence of packets between a given source and destination endpoints. Network flows are highly granular; flow endpoints are identified both by an IP address and transport layer application port numbers. NetFlow also utilizes the IP type and the input interface identifier to uniquely identify flows.

NetFlow statistics is a global traffic monitoring feature that allows flow-level monitoring of all IPv4-routed traffic through the switch. Collected statistics can be exported to an external device (NetFlow Collector/Analyzer) through the NetFlow Data Export (NDE). 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. Traffic monitoring does not need to be operating on each device in the network.

Information Derived from Hardware

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

- the packet and byte counts
- start and end timestamps
- source and destination IP addresses
- IP protocol
- source and destination port numbers.

Information Derived from Software

The software infers the following fields:

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

Determining the Input and Output interface and AS Numbers

The input and output interface values are correct when Policy Based Routing (PBR) is not used. There are simple symmetric routing scenarios, and no load balancing schemes are applied by an adjacent upstream router. The input and output interface values determined by the software, however, are not guaranteed to be accurate in the specific situations explained below.

Software determines the output interface information by looking up the FIB (Forwarding Information Base) 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. Therefore, the output interface is based solely on the destination IP address. If load balancing is enabled on the switch, instead of looking at the adjacency in the FIB entry, the load balancing hash will be applied to access the appropriate FIB path and access the appropriate adjacency. Although this process will typically yield correct results, a potential 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.

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. Therefore, 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 will be 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 router, 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 router. It is not always feasible to know the algorithm. Therefore, all flow statistics will be attributed to one input interface. Software selects the interface with the lowest IP subnet number.

- In an asymmetric routing scheme, where the traffic for an IP subnet might be received on an interface that is different from the interface where packets are sent to this IP subnet, the inferences noted previously for selecting an input interface, based on a reverse lookup, would be incorrect and cannot be verified. The reason is that RPF-based forwarding is not supported in the Catalyst 4500 series switch Cisco IOS-based supervisor engines.
- If PBR (Policy Based Routing) is enabled on the switch and the flow is destined to an address that resides in the PBR range or is sourced from an address that resides in the PBR range, the information will be incorrect. In this case, the input and output interface will most likely point to the default route (if configured) or will have no value at all (NULL).
- In the case of detecting DOS attacks, even if the route is symmetrical, it might not yield the correct results because there may be multiple paths to the destination, which is the source of the attacks.


Note

NetFlow consumes additional memory and CPU resources; therefore, it is important that you understand the resources required on your switch before enabling NetFlow.

VLAN Statistics

The NetFlow Services module, in combination with the Catalyst 4500 series switch Supervisor Engine IV, provides the capability of reporting VLAN statistics for routed traffic in and out of a VLAN, as well as Layer 2 output VLAN statistics. The CLI output for a specific VLAN is shown below:

```
cat4k-sup4-2# sh 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
```

Caveat for the NetFlow Feature

The NetFlow Services module has hardware limitations that restrict the platform support to a subset of all NetFlow fields. Specifically, the following fields will not be supported:

- TCP Flags
- ToS

The effective size of the software flow table is 256 kilobytes. 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 will 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 will automatically “age out” as they approach their statistics limit.

Enabling NetFlow Statistics Collection

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 route-cache flow	Enables Netflow switching for IP routing.
Switch(config)# ip route-cache flow inter-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.

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 {hostname ip-address} udp-port	(Required) Configures the router to export NetFlow cache entries to a specific destination (for example, a workstation).

Command	Purpose
Switch(config)# ip flow-export version {1 {5 [origin-as peer-as]}}	(Optional) Configures the router 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 Border Gateway Protocol (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 <interface>	(Optional) Specifies an interface whose IP address will be 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 are now able to do the following:

- Reduce the required bandwidth between the router and workstations, as 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	Router(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	Router(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	Router(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	Router(config-flow-cache)# export destination 10.42.41.1 9991	Enables the data export.
Step 5	Router(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
Router# show ip cache flow aggregation destination-prefix	Displays the specified aggregation cache information.

To confirm data export, perform the following task:

Command	Purpose
Router# show ip flow export	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 will be used for aggregating flows within one of the IP-address based aggregation caches (e.g. 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 will provide 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](#)
- [Configuring the Minimum Mask of a Destination-Prefix Aggregation Scheme](#)
- [Configuring the Minimum Mask of a Source-Prefix Aggregation Scheme](#)
- [Monitoring and Maintaining Minimum Masks for Aggregation Schemes](#)

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	Router(config)# ip flow-aggregation cache prefix	Configures the prefix aggregation cache.
Step 2	Router(config-flow-cache)# mask source minimum value	Specifies the minimum value for the source mask.
Step 3	Router(config-flow-cache)# mask destination minimum value	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	Router(config)# ip flow-aggregation cache destination-prefix	Configures the destination aggregation cache.
Step 2	Router(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	Router(config)# ip flow-aggregation cache source-prefix	Configures the source-prefix aggregation cache.
Step 2	Router(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
Router# show ip cache flow aggregation prefix	Displays the configured value of the minimum mask in the prefix aggregation scheme.
Router# show ip cache flow aggregation destination-prefix	Displays the configured value of the minimum mask in the destination-prefix aggregation scheme.
Router# 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 tend to be on the order of minutes (default is 30min).

Inactive Aging specifies how long after the last packet is seen a flow is removed. The Inactive parameter clears the flow cache of “stale” flows thereby preventing new flows from starving (due to lack of resources). Inactive timeout settings tend to be on the order of seconds (default is 15sec).

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, thereby ensuring 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
 0 active, 262144 inactive, 4 added
14 ager polls, 0 flow alloc failures
Active flows timeout in 1 minutes
Inactive flows timeout in 10 seconds
last clearing of statistics 15:48:37

```

Protocol	Total	Flows	Packets	Bytes	Packets	Active(Sec)	Idle(Sec)
-----	Flows	/Sec	/Flow	/Pkt	/Sec	/Flow	/Flow
UDP-other	1	0.0	3	46	0.0	0.0	10.3
IP-other	1	0.0	100	38	0.0	0.0	10.2
Total:	2	0.0	51	38	0.0	0.0	10.2

```

SrcIf          SrcIPAddress  DstIf          DstIPAddress   Pr SrcP DstP  Pkts
Switch#

```

NetFlow Configuration Examples

This section provides the following basic configuration examples:

- [Sample Netflow Enabling Schemes, page 30-10](#)
- [Sample NetFlow Aggregation Configurations, page 30-10](#)
- [Sample NetFlow Minimum Prefix Mask Router-Based Aggregation Schemes, page 30-12](#)

Sample Netflow Enabling Schemes



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 route-cache flow

```

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

```

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

```

Sample NetFlow Aggregation Configurations

This section provides the following aggregation cache configuration examples:

- [Autonomous System Configuration, page 30-11](#)
- [Destination Prefix Configuration, page 30-11](#)
- [Prefix Configuration, page 30-11](#)
- [Protocol Port Configuration, page 30-11](#)

- [Source Prefix Configuration, page 30-12](#)

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
```

Sample NetFlow Minimum Prefix Mask Router-Based Aggregation Schemes

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

- [Prefix Aggregation Scheme](#)
- [Destination-Prefix Aggregation Scheme](#)
- [Source-Prefix Aggregation Scheme](#)

Prefix Aggregation Scheme

This 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

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

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

Source-Prefix Aggregation Scheme

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

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