



CHAPTER 47

Configuring NetFlow

This chapter describes how to configure NetFlow statistics collection on the Catalyst 6500 series switches.



Note

For complete syntax and usage information for the commands used in this chapter, refer to this publication:

http://www.cisco.com/en/US/docs/ios/netflow/command/reference/nf_book.html

This chapter contains the following sections:

- [Understanding NetFlow, page 47-1](#)
- [Default NetFlow Configuration, page 47-5](#)
- [NetFlow Configuration Guidelines and Restrictions, page 47-5](#)
- [Configuring NetFlow, page 47-6](#)

Understanding NetFlow

These sections describe how NetFlow works:

- [NetFlow Overview, page 47-1](#)
- [NetFlow on the PISA, page 47-2](#)
- [NetFlow on the PFC3B, page 47-2](#)

NetFlow Overview

The NetFlow feature collects traffic statistics about the packets that flow through the switch and stores the statistics in the NetFlow table. The NetFlow table on the PISA captures statistics for flows routed in software and the NetFlow table on the PFC3B captures statistics for flows routed in hardware.

Several features use the NetFlow table: features such as network address translation (NAT) use NetFlow to modify the forwarding result; other features (such as QOS microflow policing) use the statistics from the NetFlow table to apply QOS policies. The NetFlow Data Export (NDE) feature provides the ability to export the statistics to an external device (called a NetFlow collector).

You can configure NetFlow to collect statistics for both routed and bridged traffic.

Collecting and exporting a large volume of statistics can significantly impact supervisor engine and PISA processor usage, so NetFlow provides configuration options to control the volume of statistics. These options include the following:

- NetFlow flow masks determine the granularity of the flows to be measured. Very specific flow masks generate a large number of NetFlow table entries and a large volume of statistics to export. Less specific flow masks aggregate the traffic statistics into fewer NetFlow table entries and generate a lower volume of statistics.
- Sampled NetFlow exports data for a subset of traffic in a flow, which can greatly reduce the volume of statistics exported. Sampled NetFlow does not reduce the volume of statistics collected.
- NetFlow aggregation merges the collected statistics prior to export. Aggregation reduces the volume of records exported, but does not reduce the volume of statistics collected. Note that NetFlow aggregation increases switch CPU utilization and reduces the data available at the collector. NetFlow aggregation uses NetFlow version 8.

NetFlow defines three configurable timers to identify stale flows that can be deleted from the table. NetFlow deletes the stale entries to free up table space for new entries.

NetFlow on the PISA

The NetFlow table on the PISA captures statistics for flows routed in software. NetFlow on the PISA supports NetFlow aggregation. For information about the NetFlow aggregation schemes, refer to the following document:

Cisco IOS NetFlow Configuration Guide.

For information about configuring NetFlow aggregation on the PISA, refer to the following document:

Cisco IOS NetFlow Configuration Guide.

NetFlow on the PISA supports ToS-based router aggregation, described in this document:

Cisco IOS NetFlow Configuration Guide.

For information about NetFlow for multicast IP, refer to the NetFlow Multicast Support documentation, available in the following document:

Cisco IOS NetFlow Configuration Guide.

The NetFlow Multicast Support document contains a prerequisite specifying that you need to configure multicast fast switching or multicast distributed fast switching (MDFS). However, this prerequisite does not apply when configuring NetFlow multicast support on the Supervisor Engine 32 PISA.

NetFlow on the PFC3B

The NetFlow table on the PFC3B captures statistics for flows routed in hardware. The PFC3B supports sampled NetFlow and NetFlow aggregation. The PFC3B does not support NetFlow ToS-based router aggregation.

These sections describe NetFlow on the PFC3B in more detail:

- [Flow Masks, page 47-3](#)
- [Flow Mask Conflicts, page 47-4](#)

Flow Masks

A flow is a unidirectional stream of packets between a given source and a given destination. A flow mask specifies the fields in the incoming packet that NetFlow uses to identify the flow. NetFlow gathers statistics for each flow defined by the flow mask.

The PFC3B supports the following flow masks:

- **source-only**—A less-specific flow mask. The PFC3B maintains one entry for each source IP address. Statistics for all flows from a given source IP address aggregate into this entry.
- **destination**—A less-specific flow mask. The PFC3B maintains one entry for each destination IP address. Statistics for all flows to a given destination IP address aggregate into this entry.
- **destination-source**—A more-specific flow mask. The PFC3B maintains one entry for each source and destination IP address pair. Statistics for all flows between the same source IP address and destination IP address aggregate into this entry.
- **destination-source-interface**—A more-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the destination-source flow mask.
- **full**—A more-specific flow mask. The PFC3B creates and maintains a separate table entry for each IP flow. A full entry includes the source IP address, destination IP address, protocol, and protocol ports.
- **full-interface**—The most-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the full-flow mask.

The flow mask determines the granularity of the statistics gathered, which controls the size of the NetFlow table. The less-specific flow masks result in fewer entries in the NetFlow table and the most-specific flow masks result in the most NetFlow entries.

For example, if the flow mask is set to **source-only**, the NetFlow table contains only one entry per source IP address. The statistics for all flows from a given source are accumulated in the one entry. However, if the flow mask is configured as **full**, the NetFlow table contains one entry per full flow. Many entries may exist per source IP address, so the NetFlow table can become very large. See the [“NetFlow Configuration Guidelines and Restrictions”](#) section on page 47-5 for information about NetFlow table capacity.

Flow Mask Conflicts

Several features use the NetFlow table. [Table 47-1](#) lists the flow mask requirements for each feature.

Table 47-1 Feature Requirements for Flow Masks

Feature	Source	Destination	Destination Source	Destination Source Interface	Full Flow	Interface Full Flow	Non-interface Full Flow
Reflexive ACL						X	
TCP Intercept					X	X	
Context Based Access Control (CBAC)					X		
Web Cache Redirect (WCCP)					X	X	
Server Load Balancing (SLB)					X	X	
Network Address Translation (NAT)						X	X
NetFlow Data Export (NDE)	X	X	X	X	X	X	
Sampled NetFlow						X	
NetFlow Aggregation		X		X	X	X	
Microflow Policing	X	X			X	X	

Because of the variety of feature requirements, potential flow mask conflicts can occur. Note the following flow mask constraints:

- All features must share the same limited set of flow masks.
- The PFC3B can apply only one flow mask to each packet lookup.

The Feature Manager software in the PISA is responsible for resolving feature conflicts. The Feature Manager's main strategy is to select a common flow mask that satisfies all the configured NetFlow features.

However, the Feature Manager may not find a common flow mask for the configured features, because some features have very specific requirements for the flow mask. To resolve the feature conflict, Feature Manager software may direct one of the features to be processed in software on the PISA.

In the extreme case, Feature Manager software gives priority to the feature that is configured first and rejects configuration requests for subsequent features. When you attempt to configure a subsequent feature that the Feature Manager cannot accommodate, you receive a failure message at the CLI.

Follow these guidelines to avoid problems with feature conflicts:

- Configure your highest priority features first. If an unresolvable conflict occurs, your lower priority features may be blocked.
- If possible, configure features only on the interfaces where the feature is required.
- Pay attention to response messages. If the Feature Manager turns off hardware assist for a feature, you need to ensure that feature processing does not overload the RP processor.

Note the following specific feature conflicts:

- CBAC requires the full flow mask, and is given priority over other flow-based features. If a flow mask conflict occurs, the other flow-based features are processed in the PISA.
- In general, NDE is flexible because you configure the minimum flow mask. If you have configured other flow-based features, Feature Manager software may set a more specific flow mask to meet all the feature requirements.
- Sampled NetFlow requires the full-interface flow mask. This may cause conflict with other flow-based features on the same interface.
- NDE conflicts with QoS. NDE and QoS microflow policing cannot be configured on the same interface.
- If NAT is configured on a Layer 3 interface with any feature that uses dynamic ACEs (for example, Web Proxy Authentication or NAC Layer 3 IP validation), trailing fragments may not be NAT translated correctly if NAT is configured for overload. You can use the **mls ip nat netflow-frag-14-zero** command to ensure that NAT functions correctly in this case.

Default NetFlow Configuration

Table 47-2 shows the default NetFlow configuration.

Table 47-2 Default NetFlow Configuration

Feature	Default Value
NetFlow of routed IP traffic	Disabled
NetFlow of ingress bridged IP traffic	Disabled
Sampled NetFlow	Disabled
NetFlow Aggregation	Disabled

NetFlow Configuration Guidelines and Restrictions

When configuring NetFlow, follow these guidelines and restrictions:

- The CEF table (and not the NetFlow table) implements Layer 3 switching in hardware.
- NetFlow supports bridged IP traffic.
- NetFlow supports multicast IP traffic.
- No statistics are available for flows that are switched when the NetFlow table is full.
- If the NetFlow table utilization exceeds the recommended utilization levels, there is an increased probability that there will be insufficient room to store statistics. Table 47-3 lists the recommended maximum utilization levels.

Table 47-3 NetFlow Table Utilization

PFC	Recommended NetFlow Table Utilization	Total NetFlow Table Capacity
PFC3B	117,760 (115 K) entries	131,072 (128 K) entries

Configuring NetFlow

These sections describe how to configure NetFlow:

- [Configuring NetFlow on the PFC3B, page 47-6](#)
- [Configuring NetFlow on the PISA, page 47-10](#)



Note

When you configure NAT on an interface, the PFC3B sends all fragmented packets to the PISA to be processed in software. (CSCdz51590)

Configuring NetFlow on the PFC3B

These sections describe how to configure NetFlow statistics collection on the PFC3B:

- [NetFlow PFC3B Commands Summary, page 47-6](#)
- [Enabling NetFlow on the PFC3B, page 47-7](#)
- [Setting the Minimum IP MLS Flow Mask, page 47-7](#)
- [Configuring the MLS Aging Time, page 47-7](#)
- [Configuring NetFlow Aggregation on the PFC3B, page 47-9](#)
- [Enabling NetFlow for Ingress-Bridged IP Traffic, page 47-9](#)
- [Enabling NetFlow for Multicast IP Traffic, page 47-10](#)
- [Displaying PFC3B NetFlow Information, page 47-10](#)

NetFlow PFC3B Commands Summary

Table 47-4 shows a summary of the NetFlow commands available on the PFC3B.

Table 47-4 Summary of PFC3B NetFlow commands

Command	Purpose
<code>mls netflow</code>	Enables NetFlow on the PFC3B.
<code>mls flow ip</code>	Sets the minimum flow mask.
<code>mls aging</code>	Sets the configurable aging parameters.
<code>show mls netflow {...}</code>	Displays NetFlow PFC3B information for unicast and multicast traffic.
<code>show mls netflow aggregation flowmask</code>	Displays the NetFlow aggregation flow mask.



Note

- When you configure NetFlow aggregation on the PISA, it is enabled automatically on the PFC3B.
- When you configure NetFlow for Layer 2 traffic on the PISA, it is enabled automatically on the PFC3B.
- When you configure multicast NetFlow on the PISA, it is enabled automatically on the PFC3B.

Enabling NetFlow on the PFC3B

To enable NetFlow statistics collection on the PFC3B, perform this task:

Command	Purpose
Router(config)# mls netflow	Enables NetFlow on the PFC3B.
Router(config)# no mls netflow	Disables NetFlow on the PFC3B.

This example shows how to disable NetFlow statistics collection on the PFC3B (the default setting is enabled):

```
Router(config)# no mls netflow
```

Setting the Minimum IP MLS Flow Mask

You can set the minimum specificity of the flow mask for the NetFlow table on the PFC3B. The actual flow mask may be more specific than the level configured in the **mls flow ip** command, if other configured features need a more specific flow mask (see the “[Flow Mask Conflicts](#)” section on page 47-4).

To set the minimum IP MLS flow mask, perform this task:

Command	Purpose
Router(config)# mls flow ip { source destination destination-source interface-destination-source full interface-full }	Sets the minimum IP MLS flow mask for the protocol.
Router(config)# no mls flow ip	Reverts to the default IP MLS flow mask (null).

This example shows how to set the minimum IP MLS flow mask:

```
Router(config)# mls flow ip destination
```

To display the IP MLS flow mask configuration, perform this task:

Command	Purpose
Router# show mls netflow flowmask	Displays the flow mask configuration.

This example shows how to display the MLS flow mask configuration:

```
Router# show mls netflow flowmask
current ip flowmask for unicast: destination address
Router#
```

Configuring the MLS Aging Time

The MLS aging time (default 300 seconds) applies to all NetFlow table entries. You can configure the normal aging time in the range of 32 to 4092 seconds. Flows can age as much as 4 seconds sooner or later than the configured interval. On average, flows age within 2 seconds of the configured value.

Other events might cause MLS entries to be purged, such as routing changes or a change in link state.

**Note**

If the number of MLS entries exceeds the recommended utilization (see the “[NetFlow Configuration Guidelines and Restrictions](#)” section on page 47-5), only adjacency statistics might be available for some flows.

To keep the NetFlow table size below the recommended utilization, enable the following parameters when using the **mls aging** command:

- **normal**—Configures an inactivity timer. If no packets are received on a flow within the duration of the timer, the flow entry is deleted from the table.
- **fast aging**—Configures an efficient process to age out entries created for flows that only switch a few packets, and then are never used again. The **fast aging** parameter uses the **time** keyword value to check if at least the **threshold** keyword value of packets have been switched for each flow. If a flow has not switched the threshold number of packets during the time interval, then the entry is aged out.
- **long**—Configures entries for deletion that have been active for the specified value even if the entry is still in use. Long aging is used to prevent counter wraparound, which can cause inaccurate statistics.

A typical table entry that is removed by fast aging is the entry for flows to and from a Domain Name Server (DNS) or TFTP server.

If you need to enable MLS fast aging time, initially set the value to 128 seconds. If the size of the NetFlow table continues to grow over the recommended utilization, decrease the setting until the table size stays below the recommended utilization. If the table continues to grow over the recommended utilization, decrease the normal MLS aging time.

To configure the MLS aging time, perform this task:

Command	Purpose
Router(config)# mls aging { fast [threshold {1-128} time {1-128}] long 64-1920 normal 32-4092}	Configures the MLS aging time for a NetFlow table entry.
Router(config)# no mls aging fast	Disables fast aging.
Router(config)# no mls aging { long normal }	Reverts to the default MLS aging time.

This example displays how to configure the MLS aging time:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# mls aging fast threshold 64 time 30
```

To display the MLS aging-time configuration, perform this task:

Command	Purpose
Router# show mls netflow aging	Displays the MLS aging-time configuration.

This example shows how to display the MLS aging-time configuration:

```
Router# show mls netflow aging
enable timeout packet threshold
-----
```

```
normal aging true 300 N/A
fast aging true 32 100
long aging true 900 N/A
```

Configuring NetFlow Aggregation on the PFC3B

NetFlow Aggregation is configured automatically on the PFC3B when you configure NetFlow Aggregation on the PISA (see the [“Configuring NetFlow Aggregation on the PISA”](#) section on page 47-11).

To display NetFlow Aggregation information for the PFC3B, perform this task:

Command	Purpose
Router # show ip cache flow aggregation {as destination-prefix prefix protocol-port source-prefix) module slot_num	Displays the NetFlow Aggregation cache information.
Router # show mls netflow aggregation flowmask	Displays the NetFlow Aggregation flow mask information.



Note

The PFC3B does not support NetFlow ToS-based router Aggregation.

This example shows how to display the NetFlow Aggregation cache information:

```
Router# show ip cache flow aggregation destination-prefix module 1
IPFLOW_DST_PREFIX_AGGREGATION records and statistics for module :1
IP Flow Switching Cache, 278544 bytes
2 active, 4094 inactive, 6 added
236 ager polls, 0 flow alloc failures
Active flows timeout in 30 minutes
Inactive flows timeout in 15 seconds
Dst If Dst Prefix Msk AS Flows Pkts B/Pk Active
Gi7/9 9.1.0.0 /16 0 3003 12M 64 1699.8
Gi7/10 11.1.0.0 /16 0 3000 9873K 64 1699.8
Router#
```

This example shows how to display the NetFlow Aggregation flow mask information:

```
Router# show mls netflow aggregation flowmask
Current flowmask set for netflow aggregation : Vlan Full Flow
Netflow aggregations configured/enabled :
AS Aggregation
PROTOCOL-PORT Aggregation
SOURCE-PREFIX Aggregation
DESTINATION-PREFIX Aggregation
Router#
```

Enabling NetFlow for Ingress-Bridged IP Traffic

NetFlow for ingress-bridged IP traffic on the PFC3B is enabled when you configure NetFlow for ingress-bridged IP traffic on the PISA. See the [“Enabling NetFlow for Ingress-Bridged IP Traffic”](#) section on page 47-11.

Enabling NetFlow for Multicast IP Traffic

NetFlow for multicast IP traffic on the PFC3B is enabled when you configure NetFlow for multicast IP traffic on the PISA.

For additional information, see the “[Enabling NetFlow for Multicast IP Traffic](#)” section on page 47-12.

Displaying PFC3B NetFlow Information

To display information about NetFlow on the PFC3B, use the following command:

Command	Purpose
Router(config)# show mls netflow { aggregation aging creation flowmask ip ipv6 mpls table-contention usage }	Displays information about NetFlow on the PFC3B.

Configuring NetFlow on the PISA

These sections describe how to configure NetFlow on the PISA:

- [Summary of NetFlow Commands on the PISA](#), page 47-10
- [Enabling NetFlow on the PISA](#), page 47-11
- [Configuring NetFlow Aggregation on the PISA](#), page 47-11
- [Enabling NetFlow for Ingress-Bridged IP Traffic](#), page 47-11
- [Enabling NetFlow for Multicast IP Traffic](#), page 47-12

Summary of NetFlow Commands on the PISA

Table 47-5 shows the NetFlow commands available on the PISA.

Table 47-5 Summary of PISA NetFlow Commands

Command	Purpose
interface x ip flow ingress	Enables NetFlow on the PISA and the PFC3B for the specified interface.
ip flow-aggregation cache	Configure NetFlow aggregation. Note that configuring aggregation on the PISA also enables aggregation for the PFC3B.
export version {8 9}	Specifies aggregation data export format 8 or 9.
mask source minimum x	Specifies the aggregation minimum mask.
ip flow ingress layer2-switched vlan x	Enables NetFlow for Layer 2 switched traffic.
interface x ip multicast netflow {ingress egress}	Enables NetFlow multicast traffic on the specified interface (for PISA and PFC3B).
show ip cache flow aggregation	Shows the configuration for NetFlow aggregation.
show ip cache verbose flow	Shows the configuration for multicast NetFlow.

Enabling NetFlow on the PISA

To enable NetFlow on the PISA, perform this task for each Layer 3 interface from which you want NetFlow:

	Command	Purpose
Step 1	Router(config)# interface {vlan vlan_ID} {type slot/port} {port-channel port_channel_number}	Selects a Layer 3 interface to configure.
Step 2	Router(config-if)# ip flow ingress	Enables NetFlow on the selected interface, for flows routed in hardware or software. You must also enable NetFlow on the PFC3B to enable NetFlow for flows routed in hardware.

Configuring NetFlow Aggregation on the PISA

To configure NetFlow aggregation on the PISA, use the procedures at this document:

Cisco IOS NetFlow Configuration Guide.

To configure NetFlow ToS-based router aggregation on the PISA, use the procedures at this URL:document:

Cisco IOS NetFlow Configuration Guide.



Note

- When you configure NetFlow aggregation on the PISA, it is configured automatically on the PFC3B (see the “[Configuring NetFlow Aggregation on the PFC3B](#)” section on page 47-9).
- The PFC3B does not support NetFlow ToS-based router aggregation.

Enabling NetFlow for Ingress-Bridged IP Traffic

NetFlow supports ingress-bridged IP traffic.



Note

- When you enable NetFlow for ingress-bridged IP traffic, the statistics are available to the Sampled NetFlow feature (see the “[NetFlow Sampling](#)” section on page 46-7).
- To enable NetFlow for bridged IP traffic on a VLAN, you must create a corresponding VLAN interface, assign it an IP address, and enter the **no shutdown** command to bring up the interface.

To enable NetFlow for ingress-bridged IP traffic in VLANs, perform this task:

Command	Purpose
Router(config)# ip flow ingress layer2-switched vlan <i>vlan_ID</i> [- <i>vlan_ID</i>] [, <i>vlan_ID</i> [- <i>vlan_ID</i>]]	Enables NetFlow for ingress-bridged IP traffic in the specified VLANs. Note NetFlow for ingress-bridged IP traffic in a VLAN requires that NetFlow on the PFC3B be enabled with the mls netflow command.
Router(config)# no ip flow ingress layer2-switched vlan <i>vlan_ID</i> [- <i>vlan_ID</i>] [, <i>vlan_ID</i> [- <i>vlan_ID</i>]]	Disables NetFlow for ingress-bridged IP traffic in the specified VLANs.

This example shows how to enable NetFlow for ingress-bridged IP traffic in VLAN 200:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip flow ingress layer2-switched vlan 200
```

Enabling NetFlow for Multicast IP Traffic

To enable NetFlow for multicast IP, perform this task:

Command	Purpose
Step 1 Router(config)# interface { <i>vlan vlan_ID</i> } { <i>type slot/port</i> } { port-channel <i>port_channel_number</i> }	Selects a Layer 3 interface to configure.
Step 2 Router(config-if)# ip flow ingress	Enables NetFlow on the interface.
Step 3 Router(config-if)# ip multicast netflow { ingress egress }	Enables NetFlow multicast traffic on the specified interface (for PISA and PFC3B). <ul style="list-style-type: none"> Specify ingress to enable NetFlow multicast ingress accounting Specify egress to enable NetFlow multicast egress accounting

For additional information about NetFlow for multicast IP, refer to the NetFlow Multicast Support documentation, available at the following document:

Cisco IOS NetFlow Configuration Guide.

The NetFlow Multicast Support document contains a prerequisite specifying that you need to configure multicast fast switching or multicast distributed fast switching (MDFS). However, this prerequisite does not apply when configuring NetFlow multicast support on the Supervisor Engine 32 PISA.