CHAPTER 55

Configuring NetFlow

This chapter describes how to configure NetFlow statistics collection in Cisco IOS Software Release 12.2SX.

Note For complete syntax and usage information for the commands used in this chapter, see these publications:

- The Cisco IOS Software Releases 12.2SX Command References at this URL:
  http://www.cisco.com/en/US/docs/ios/mcl/122sx_mcl.html
- The Release 12.2 publications at this URL:
  http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/index.htm

This chapter contains the following sections:

Understanding NetFlow

The NetFlow feature collects traffic statistics about the packets that flow through the switch and stores the statistics in the NetFlow table. Several features use the statistics in the NetFlow table, and the statistics can be exported using the NetFlow Data Export (NDE) feature.

These sections provide additional information about NetFlow:

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 route processor (RP) captures statistics for flows routed in software and the NetFlow table on the PFC (and on each DFC) 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).

In PFC3A mode, NetFlow collects statistics only for routed traffic. With other PFCs, you can configure NetFlow to collect statistics for both routed and bridged traffic.

Collecting and exporting a large volume of statistics can significantly impact switch processor (SP) and route processor (RP) CPU 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.
- Per-interface NetFlow allows you to enable or disable NetFlow data collection on Layer 3 interfaces.
- 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. NetFlow aggregation increases SP 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 clear table space for new entries.

NetFlow on the PFC

The NetFlow table on the PFC captures statistics for flows routed in hardware.

These sections describe NetFlow on the PFC in more detail:

- Flow Masks, page 55-2
- Flow Mask Conflicts, page 55-3
- Default NetFlow Configuration, page 55-6

Flow Masks

A flow is a unidirectional stream of packets between a source and a destination. The flow mask specifies the fields in the incoming packet that NetFlow uses to match (or create) a NetFlow table entry.

All flow masks include the ingress interface in their definition. Therefore, NetFlow always collects statistics on a per-interface basis. You can also enable or disable NetFlow per-interface.

The PFC supports the following flow masks:

- interface-source—A less-specific flow mask. Statistics for all ingress flows on an interface from each source IP address aggregate into one entry.
- interface-destination—A less-specific flow mask. Statistics for all ingress flows on an interface to each destination IP address aggregate into one entry.
- interface-destination-source—A more-specific flow mask. Statistics for all ingress flows on an interface between the same source IP address and destination IP address aggregate into one entry.
- interface-full—The most-specific flow mask. The PFC creates and maintains a separate table entry for each IP flow on an interface. An interface-full entry includes the source IP address, destination IP address, protocol, and protocol ports.

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 interface-source, the NetFlow table contains one entry per source IP address. (Assume that NetFlow is enabled on only one interface). The statistics for all flows from each source are accumulated in the one entry. However, if the flow mask is configured as interface-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 55-6 for information about NetFlow table capacity.

Flow Mask Conflicts

Several features use the NetFlow table. Table 55-1 lists the flow mask requirements for each feature.

Table 55-1 Feature Requirements for Flow Masks

Feature	Interface Source	Interface Destination	Interface Destination Source	Full	Interface Full	Non-interface Full
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	
Sampled NetFlow					X	
NetFlow Aggregation		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 PFC can apply only one flow mask to each packet lookup.

The Feature Manager software in the RP is responsible for resolving feature conflicts. The Feature Manager's main purpose 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 RP.

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.

To avoid problems with feature conflicts, follow these guidelines:

- 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 RP.
- 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 translated correctly by NAT if NAT is configured for overload. Except in PFC3A mode, you can use the **mls ip nat netflow-frag-l4-zero** command to ensure that NAT functions correctly in this configuration.

NetFlow on the RP

The NetFlow feature on the RP captures statistics for flows routed in software.

For additional information about NetFlow on the RP, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfc.htm#wp1000962

NetFlow Features

NetFlow supports the following features:

Per-interface NDE

Cisco IOS Release 12.2(33)SXH and later releases support per-interface NDE, which enables PFC NetFlow data collection on a per-interface basis. With releases earlier than Release 12.2(33)SXH, NetFlow on the PFC could be only be enabled and disabled globally.

When you upgrade to a software release that supports the per-interface NDE feature, the system automatically enables per-interface NDE and configures the **ip flow ingress** command on every Layer 3 interface. This one-time action takes place on the first reload after the upgrade and maintains backward compatibility with the global NetFlow enable command. After the reload, you can configure the **no ip flow ingress** command on Layer 3 interfaces to selectively disable PFC and RP NetFlow data collection/export.

The per-interface NDE feature only applies to IPv4 unicast flows on Layer 3 interfaces. Flows for non-IPv4 protocols (such as IPv6 and MPLS) are not controlled by this feature.

NetFlow Aggregation

NetFlow supports aggregation for packets forwarded in hardware (PFC) or software (RP). For information about NetFlow aggregation schemes, see the following publications:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfov.htm#wp1001212

For information about configuring NetFlow aggregation, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfc.htm#wp1001058

NetFlow on the RP supports ToS-based router aggregation, described at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s15/dtnfltos.htm

NetFlow for Multicast IP

NetFlow is supported for multicast IP packets forwarded in hardware (PFC) or software (RP).

NetFlow multicast provides ingress accounting and egress accounting. With ingress accounting, NetFlow creates one flow per source and includes information about how many packet replications occur. With egress accounting, NetFlow creates one flow for each outgoing interface.

Optionally, NetFlow multicast keeps statistics for multicast packets that fail the reverse path fail (RPF) check.

Default NetFlow Configuration

Table 55-2 shows the default NetFlow configuration.

Table 55-2 Default NetFlow Configuration

Feature	Default Value
NetFlow	Disabled
NetFlow of routed IP traffic	Disabled
NetFlow of ingress bridged IP traffic	Disabled
Sampled NetFlow	Disabled
NetFlow Aggregation	Disabled
Per-interface NDE	Enabled
Exclude ACL-denied traffic	Disabled (NetFlow creates entries for ACL-denied traffic)

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.
- Except in PFC3A mode, NetFlow supports bridged IP traffic. PFC3A mode does not support NetFlow bridged IP traffic.
- NetFlow supports multicast IP traffic.

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

- 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 55-3 lists the recommended maximum utilization levels.

Table 55-3 NetFlow Table Utilization

PFC	Recommended NetFlow Table Utilization	Total NetFlow Table Capacity
PFC3CXL	235,520 (230 K) entries	262,144 (256 K) entries
PFC3C	117,760 (115 K) entries	131,072 (128 K) entries
PFC3BXL	235,520 (230 K) entries	262,144 (256 K) entries
PFC3B	117,760 (115 K) entries	131,072 (128 K) entries
PFC3A	65,536 (64 K) entries	131,072 (128 K) entries

Configuring NetFlow

These sections describe how to configure NetFlow:

Configuring NetFlow on the PFC

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

NetFlow PFC Commands Summary

Table 55-4 shows a summary of the NetFlow commands available on the PFC.

Table 55-4 Summary of PFC NetFlow Commands

Command	Purpose
mls netflow	Enables NetFlow on the PFC.
mls flow ip	Sets the minimum flow mask.
mls aging	Sets the configurable aging parameters.
mls exclude acl-deny	Disables the creation of flows for ACL-denied traffic.
show mls netflow {...}	Displays NetFlow PFC information for unicast and multicast traffic.
show mls netflow aggregation flowmask	Displays the NetFlow aggregation flow mask.

Enabling NetFlow on the PFC

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

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

This example shows how to disable NetFlow statistics collection on the PFC (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 PFC. The actual flow mask may be more specific than the level configured in the **mls flow** command, if other configured features need a more specific flow mask (see the "Flow Mask Conflicts" section on page 55-3).

To set the minimum IPv4 flow mask, perform this task:

Command	Purpose
`Router(config)# mls flow ip {interface-source | interface-destination | interface-destination-source | interface-full}`	Sets the minimum flow mask for IPv4 packets.
`Router(config)# no mls flow ip`	Reverts to the default flow mask (null).

This example shows how to set the minimum flow mask:

```
Router(config)# mls flow ip interface-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: if-dst
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 55-6), 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 Exclude ACL-deny

By default, NetFlow table entries are created for ACL-denied flows. These flows can cause the NetFlow table to overflow. With Release 12.2(33)SXH and later releases, to exclude ACL-denied flows from the NetFlow table, perform this task:

Command	Purpose
`Router# mls exclude acl-deny`	Excludes ACL-denied flows from the NetFlow table.

This example shows how to exclude ACL-denied flows from the NetFlow table:

```
Router(config)# mls exclude acl-deny
```

Displaying PFC NetFlow Information

To display information about NetFlow on the PFC, perform this task:

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

Configuring NetFlow Features

NetFlow features generally apply to packets forwarded in hardware (PFC) and software (RP). For the features to apply to PFC, you need to enable NetFlow on the PFC.

These sections describe how to configure NetFlow features:

Configuring NetFlow on Layer 3 Interfaces

The per-interface NDE feature allows you to enable or disable NetFlow collection on a per-interface basis for packets forwarded in hardware (PFC) or software (RP). This feature is automatically enabled in Release 12.2(33)SXH and later releases.

To enable or disable NetFlow for a Layer 3 interface, perform this task:

	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 for the specified interface. NetFlow will collect statistics for packets forwarded in hardware (PFC) or software (RP).
Step 3	`Router(config-if)# no ip flow ingress`	Disables NetFlow for the specified interface. NetFlow will stop collecting statistics for packets forwarded in hardware (PFC) or software (RP).

When you upgrade for the first time to a software image that supports per-interface NetFlow on the PFC, the system automatically configures each Layer 3 interface to enable NetFlow (this ensures backward compatibility with the global **mls netflow** command). This one-time action occurs during the first system restart after the upgrade. After this action, you can configure Layer 3 interfaces to disable or enable NetFlow data collection.

Enabling NetFlow for Ingress-Bridged IP Traffic

Except in PFC3A mode, NetFlow supports ingress-bridged IP traffic. PFC3A mode does not support NetFlow for bridged IP traffic.

Note

- When you enable NetFlow for ingress-bridged IP traffic, the statistics are available to the Sampled NetFlow feature (see the "Sampled NetFlow" section on page 56-7).
- To enable NetFlow for bridged IP traffic on a VLAN, you must create a corresponding VLAN interface and enter the **no shutdown** command. The **no shutdown** command can be followed, if necessary, by the **shutdown** command.
- For Layer 3 VLANs, enabling NetFlow for ingress-bridged IP traffic also enables NetFlow for Layer 3 flows on the specified VLANs.
- The exported bridged flows will have ingress and egress VLAN information and not the physical port information.

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

Command	Purpose
`Router(config)# ip flow ingress layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	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 PFC be enabled with the **mls netflow** command.
`Router(config)# no ip flow ingress layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	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
```

Configuring NetFlow Aggregation

To configure NetFlow aggregation, use the procedures at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfc.htm#wp1001058

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

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s15/dtnfltos.htm

Note

- When you configure NetFlow aggregation, it is configured automatically for packets forwarded in hardware (PFC) or software (RP).
- The PFC and DFCs do not support NetFlow ToS-based router aggregation.

To display NetFlow Aggregation information for the PFC or DFCs, 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 PFC and DFCs do 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
```

Configuring NetFlow for Multicast IP Traffic

To configure NetFlow for multicast IP traffic, perform this task:

	Command	Purpose
Step 1	`Router(config)# ip multicast netflow output-counters`	(Optional) Enables the calculation of output bytes/packets for an ingress flow.
Step 2	`Router(config)# ip multicast netflow rpf-failure`	(Optional) Enables NetFlow for multicast data that fails the RPF check.
Step 3	`Router(config)# interface {vlan` *vlan_ID*`} | {`*type slot/port*`} | {port-channel` *port_channel_number*`}`	Selects a Layer 3 interface to configure.
Step 4	`Router(config-if)# ip flow {ingress | egress}`	Enables NetFlow multicast traffic on the specified interface (for RP and PFC). • Specify **ingress** to enable NetFlow multicast ingress accounting. • Specify **egress** to enable NetFlow multicast egress accounting.

For additional information about configuring NetFlow for multicast traffic, see the Configuring NetFlow Multicast Accounting document, available at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/honf_c/chap10/onf_dmlc.htm

This document specifies a prerequisite 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 with 12.2SX releases.

Note The Configuring NetFlow Multicast Accounting document describes new configuration commands for Cisco IOS release 12.2(4) and newer releases. The 12.2SX releases support the new commands.