



Configuring Flexible NetFlow

**Note**

Flexible NetFlow is supported only on Supervisor Engine 8-E, Supervisor Engine 7-E, Supervisor Engine 7L-E, and Catalyst 4500X.

Flow is defined as a unique set of key fields attributes, which might include fields of packet, packet routing attributes, and input and output interface information. A NetFlow feature defines a flow as a sequence of packets that have the same values for the feature key fields. Flexible NetFlow (FNF) allows you to collect and optionally export a flow record that specifies various flow attributes. NetFlow collection supports IP, IPv6 and Layer 2 traffic.

**Note**

This chapter provides Catalyst 4500 switch specific information. For more information, refer to the URL:

http://www.cisco.com/en/US/products/ps6965/products_ios_protocol_option_home.html

**Note**

When IP routing is disabled, on the interface configured with NetFlow Lite, packets are not received on NetFlow collector. Enable IP routing for the NetFlow collector to work.

This chapter addresses both VSS and non-VSS environments:

- [VSS Environment, page 67-1](#)
- [Non-VSS Environment, page 67-8](#)

VSS Environment

The following items apply to a Catalyst 4500 series switch that belongs to a Virtual Switch System (VSS):

1. The Catalyst 4500 series switch supports ingress flow statistics collection for switched and routed packets; it does not support Flexible Netflow on egress traffic.
2. Each switch in an VSS has an independent NFE (Netflow Engine). This means that when there is ingress traffic on both the VSS Active and Standby switches, each is capable of creating flows for its ingress traffic

3. Configuration is performed on the VSS Active switch, which is synchronized to the VSS Standby switch.
4. Netflow **show** commands including Top Talkers, aggregate cache, and **clear** commands must be executed independently on VSS Active and Standby switch. The VSS Standby console will be available via remote console access from the VSS Active switch.
5. Supervisor Engine 8-E, Supervisor Engine 7-E, Supervisor Engine 7L-E, and Catalyst 4500X support a 100,000 entry hardware flow table. Both VSS Active and Standby switch have independent hardware flow tables of 100,000 entries. The hardware flow table is shared by all the flow monitors on a switch. To prevent one monitor from using all the flow table entries, the number of entries that it uses on a switch can be limited by the **cache entries number** command. This limit is per flow monitor, irrespective of the number of targets it is attached to.

The following example illustrates how to configure the flow monitor *m1* cache to hold 1000 entries. With this configuration, interface gig 1/3/1 (on the VSS Active) can create a maximum of 1000 flows and interface gig 2/3/2 (on the VSS Standby) can create a maximum of 1000 flows:

```

flow exporter e1
    ! exporter specifies where the flow records are send to
    destination 20.1.20.4
!
flow record r1
    ! record specifies packet fields to collect
    match ipv4 source address
    match ipv4 destination address
    collect counter bytes long
    collect counter packets long
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
!
flow monitor m1
    ! monitor refers record configuration and optionally exporter
    ! configuration. It specifies the cache size i.e. how many unique flow
    ! records to collect
    record r1
    exporter e1
    cache timeout active 60
    cache timeout inactive 30
    cache entries 1000

!interface GigabitEthernet 1/3/1
    ! layer2-switched allows collection of flow records even when the packet is
    ! bridged
    ip flow monitor m1 layer2-switched input
!
interface GigabitEthernet 2/3/2
    ip flow monitor m1 input
!

```

6. Catalyst 4500-E Series Switches with Sup-8E have two ASICs, and the ASIC that the software programs for a given flow monitor depends on what the flow monitor is attached to. When a datalink flow monitor is attached to an SSID (WLAN), the software programs the ASIC on the Daughter Card that creates flows only for pure Layer 2 packets (no IP header). By contrast, when a datalink flow monitor is attached to a port, or a port VLAN, or a VLAN for example, the software programs the ASIC (Netflow Engine) that creates flows for all packets.
7. Flow collection is supported on multiple targets (Port, VLAN, per-port per-VLAN (FNF can be enabled on a specific VLAN on a given port)) and on a port-channel (FNF is configured on the port-channel interface, rather than individual member ports). These targets can be on the VSS Active or on the VSS Standby. For example, if the target is a VLAN, it can consist of ports belonging to

both switches. If there is ingress traffic in that VLAN on both switches, flows will be created in their independent flow caches. However, no Netflow configuration can be applied on the Virtual Switch Link (VSL) ports.

8. 64 unique flow record configurations are supported.
9. Flow QoS/UBRL and FNF cannot be configured on the same target. (For information on Flow-based QoS, see the section [Flow-based QoS, page 42-10.](#))
10. 14,000 unique IPv6 addresses can be monitored.
11. On a given target, one monitor per traffic type is allowed. However, you can configure multiple monitors on the same target for different traffic types.

For example, the following configuration is allowed:

```
! vlan config 10
  ip flow monitor <name> input
  ipv6 flow monitor <name> input
!
```

The following configuration is not allowed:

```
!
interface GigabitEthernet 3/1
  ip flow monitor m1 input
  ip flow monitor m2 input
```

12. On a given target monitoring Layer 2 and Layer 3, simultaneous traffic is not supported:

```
interface channel-group 1
  datalink flow monitor m1 input
  ip flow monitor m2 input
!
```

13. Selection of Layer 2 and Layer 3 packet fields in a single flow record definition is not allowed. However, ingress 802.1Q VLAN Id of packet and Layer 3 packet field selection is allowed.
14. To attach a monitor to port or port-vlan targets, a flow record matching on ingress 802.1Q VLANId key field, must match on input interface also as key field.



Note The **match datalink dot1q vlan input** option is unavailable prior to IOS Release XE 3.3.0; you would only see the **input** option starting with the IOS Release XE 3.3.0.

15. Flow monitor matching on ingress 802.1Q VLANId as key field cannot be attached on a VNET trunk port target.
16. Only permanent and normal flow cache types are supported.
17. Supervisor Engine 8-E, Supervisor Engine 7-E, Supervisor Engine 7L-E, and Catalyst 4500X do not support:
 - predefined records like traditional routers (**record netflow ipv4 original-input**)
 - flow-based sampler.
18. On VLAN interfaces, when you use the **interface** option with the **Cos,Tos**, **TTL** or **Packet length** options, the system displays inaccurate results for the interface input field.
19. The VSS active and standby switches independently export flows, to the same or different Netflow collectors depending on flow exporter configuration. An IP route to the Netflow collector must exist and it should be reachable from the VSS for flow export.

20. At the collector, the flow sequence numbers are local to a switch and will be monotonically increasing for each member of VSS. Additionally, the SourceId field of the v9 export packet uniquely identifies the VSS switch number that it was exported from.
21. The configuration of the flow exporter does not support the option **output features**.
22. Maximum number of VRFs that can be used for the flow exporter destination address configuration in VSS is 5. This limit includes the Global Routing Table and is common across all flow exporters in the VSS.

For example, when the user tries to configure an exporter destination address using a sixth VRF limit is exceeded, the following warning is displayed:

```
flow exporter e10
      destination 20.1.20.4 vrf blue
%%Warning - Netflow exporter on Cat4k VSS switch cannot exceed a total max of 5 vrfs
used for destination address
configuration. Flow exporter e10 cannot export in vrf blue.
```

23. Flow aging in flow cache is controlled through active and in-active timer configuration. The minimum for active and in-active aging timers is 5 seconds. The timers must be in units of 5 seconds.



Note Flows in the hardware table are deleted after 5 seconds of in-activity irrespective of the active or in-active timer configuration values. This allows you to create new hardware flows quickly.

24. First and Last-seen flow timestamp accuracy is within 3 seconds.
25. 2048 Flow monitors and records are supported.

When TTL is configured as a flow field, the following values are reported for a given packet TTL value. [Table 67-1](#) lists the packet TTL and reported values.

26. Cisco TrustSec (CTS) fields are supported. These fields use Netflow collector to monitor and troubleshoot the CTS network, and to segregate traffic based on source group tag (SGT) values.
 - When configuring the source group tag (**collect flow cts source group-tag**), note the following:

The system copies the packets to software before it retrieves the CTS field. A large number of flows mean that a large number of packets are copied to the software, possibly affecting CPU performance.

The maximum number of (unique) hosts allowed in the switch (IP addresses) is 12,000.

In case of burst packets, the software may not be able to retrieve the CTS field because the software queue is throttled.
 - When configuring the destination group tag (**collect flow cts destination group-tag**), note that this CTS field value is collected only if you have already configured an IP-to-SGT mapping.
 - When configuring switch-derived source group tags (**collect flow cts switch derived-sgt**), note that the switch derives this value locally.
 - When configuring CTS fields on Supervisor Engine 8-E, note that CTS fields are not supported on wireless interfaces (WLAN) and SSID.

Table 67-1 TTL Map: TTL Configured

Packet TT Value	Reported Value
0	0
1	1
2-10	10
11-25	25
26-50	50
51-100	100
100-150	150
150-255	255

- When packet length is configured as a flow field, the following values are reported for a given packet length value. [Table 67-2](#) lists the packet length and reported values.

Table 67-2 Packet Length Map: Packet Length Configured

Packet Length	Reported Value
0-64	64
65-128	128
129-256	256
257-512	512
513-756	756
757-1500	1500
1500-4000	4000
4000+	8192

The following table lists the options available through FNF and the supported fields.

Table 67-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
Data Link Fields (Layer 2 Flow Label + A94)		
dot1q priority	802 1Q user	
dot1q vlan	802.1Q VLAN ID	Ingress VLAN is supported as key field.
mac destination-address	Upstream destination MAC address	
mac source-address	Down stream source MAC address	
IPv4 Fields		
destination address	IPv4 destination address	Yes

Table 67-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
DSCP	IPv4 DSCP (part of TOS)	
fragmentation flags	IPv4 fragmentation flags	Supported as a non key field. DF flag is not supported
is-multicast	Indicator of an IPv4 multicast packet (0 - if it's not, 1 - if it is)	Supported as a non-key field.
Precedence	IPv4 precedence	
Protocol	IPv4 protocol	
source address	IPv4 source address	
total length	IPv4 datagram	Values are reported based on Table 67-2 .
Total length minimum	Minimum packet size seen	
Total length maximum	Maximum packet size seen	
Tos	IPv4 Type of Service (TOS)	
ttl	Pv4 Time to Live (TTL)	Values are reported based on Table 67-1 .
ttl minimum		Supported as a non-key field.
ttl maximum		Supported as a non-key field.
CTS Fields		
flow cts destination group-tag		Supported as a non-key field; configuring the IPv4 destination address is a prerequisite to using this field.
flow cts source group-tag		Supported as a non-key field; configuring the IPv4 source address is a prerequisite to using this field.
flow cts switch derived-sgt	Switch-derived source group-tag	Supported as a non-key field; configuring the IPv4 source address is a prerequisite to using this field.
IPv6 Fields		
destination address	IPv6 destination address	
dscp	IPv6 DSCP (part of IPv6 traffic class)	
flow-label	IPv6 flow label	

Table 67-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
is-multicast	Indicator of an IPv6 multicast packet (0 - if it's not, 1 - if it is)	Supported as a non-key field
hop-limit	IPv6 hop limit (replaces IPv4 ttl)	Values are reported based on Table 67-1 .
hop-limit minimum	IPv6 minimum hop limit value seen in the flow.	Supported as a non-key field.
hop-limit maximum	IPv6 maximum hop limit value seen in the flow.	Supported as a non-key field.
next-header	IPv6 next header type	Only first next header is reported
total length	IPv6 total packet length	Values are based on Table 67-2 .
Total length minimum	Minimum packet size seen	
Total length maximum	Maximum packet size seen	
protocol	IPv6 next header type in the last IPv6 extension header	
source address	IPv6 source address	
traffic-class	IPv6 traffic class	Yes
Routing Attributes		
forwarding-status	Forwarding status for the packet (forwarded, terminated in the router, dropped by ACL, RPF, CAR)	Supported as a non-key field
Layer 4 Header Fields		
Field	Description	Comments
TCP Header Fields		
destination-port TCP destination number	TCP destination port	
flags [ack] [fin] [psh] [rst] [syn] [urg]	TCP flags.	Supported as non-key fields.
source-port	TCP source port	
UDP Header Fields		
destination-port	UDP destination port	
source-port	UDP source port	
ICMP Header Fields		
code	ICMP code	

Table 67-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
type	ICMP type	
IGMP Header Fields		
type	IGMP	
Interface Fields		
input	Input interface index	
output	Input interface index	Output interface can be supported only as non-key.
Flexible NetFlow feature related fields		
direction: input		
Counter Fields		
bytes	32 bit counters	
bytes long	64 bit counter	
packets	32 bit counters	
packets long	64 bit counter of the packets in the flow	
Timestamp		
first seen	Time-stamp of the first packet that is accounted in the flow (in milliseconds, starting from the router boot-up)	3 sec accuracy
last seen	Time-stamp of the last packet that is accounted in the flow (in milliseconds, starting from the router boot-up)	3 sec accuracy

Configuring Flow Monitor Cache Values

Setting active cache timeout to a small value may cause the flows to be exported more frequently to the remote collector. This also causes software to delete flows from the local cache after exporting. So, cache statistics reported by switch may not display the actual flows being monitored.

Non-VSS Environment

The following items apply to the Catalyst 4500 series switch:

The Catalyst 4500 series switch supports ingress flow statistics collection for switched and routed packets; it does not support Flexible Netflow on egress traffic.

1. Supervisor Engine 8-E, Supervisor Engine 7-E, Supervisor Engine 7L-E, and Catalyst 4500X support a 100,000 entry hardware flow table. The hardware flow table is shared by all the flow monitors on a switch. To prevent one monitor from using all the flow table entries, the number of entries that it uses on a switch can be limited by the **cache entries number** command. This limit is per flow monitor, irrespective of the number of targets it is attached to.

The following example illustrates how to configure the flow monitor *m1* cache to hold 1000 entries. With this configuration, interface gig 3/1 can create a maximum of 1000 flows and interface gig 3/2 can create a maximum of 1000 flows:

```

flow exporter e1
    ! exporter specifies where the flow records are sent to
    destination 20.1.20.4
!
flow record r1
    ! record specifies packet fields to collect
    match ipv4 source address
    match ipv4 destination address
    collect counter bytes long
    collect counter packets long
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
!
flow monitor m1
    ! monitor refers record configuration and optionally exporter
    ! configuration. It specifies the cache size i.e. how many unique flow
    ! records to collect
    record r1
    exporter e1
    cache timeout active 60
    cache timeout inactive 30
    cache entries 1000

!interface GigabitEthernet 3/1
    ! layer2-switched allows collection of flow records even when the packet is
    ! bridged
    ip flow monitor m1 layer2-switched input
!
interface GigabitEthernet 3/2
    ip flow monitor m1 input
!

```

2. Catalyst 4500-E Series Switches with Sup-8E have two ASICs, and the ASIC that the software programs for a given flow monitor depends on what the flow monitor is attached to. When a datalink flow monitor is attached to an SSID (WLAN), the software programs the ASIC on the Daughter Card that creates flows only for pure Layer 2 packets (no IP header). By contrast, when a datalink flow monitor is attached to a port, or a port VLAN, or a VLAN for example, the software programs the ASIC (Netflow Engine) that creates flows for all packets.
3. Flow collection is supported on multiple targets (Port, VLAN, per-port per-VLAN (FNF can be enabled on a specific VLAN on a given port)) and on a port-channel (FNF is configured on the port-channel interface, rather than individual member ports).
4. 64 unique flow record configurations are supported.
5. Flow QoS/UBRL and FNF cannot be configured on the same target. (For information on Flow-based QoS, see the section [Flow-based QoS, page 42-10.](#))
6. 14,000 unique IPv6 addresses can be monitored.

7. On a given target, one monitor per traffic type is allowed. However, you can configure multiple monitors on the same target for different traffic types.

For example, the following configuration is allowed:

```
! vlan config 10
  ip flow monitor <name> input
  ipv6 flow monitor <name> input
!
```

The following configuration is not allowed:

```
!
interface GigabitEthernet 3/1
  ip flow monitor m1 input
  ip flow monitor m2 input
```

8. On a given target monitoring Layer 2 and Layer 3, simultaneous traffic is not supported:

```
interface channel-group 1
  datalink flow monitor m1 input
  ip flow monitor m2 input
!
```

9. Selection of Layer 2 and Layer 3 packet fields in a single flow record definition is disallowed. However, ingress 802.1Q VLAN Id of packet and Layer 3 packet field selection is allowed.
10. To attach a monitor to port or port-vlan targets, a flow record matching on ingress 802.1Q VLAN Id as the key field, must also match on the input interface as the key field.



Note Flow monitor matching on ingress 802.1Q VLAN Id as the key field cannot be attached on a VNET trunk port target.

11. Only permanent and normal flow cache types are supported.
12. Supervisor Engine 8-E, Supervisor Engine 7-E, Supervisor Engine 7L-E, and Catalyst 4500X do not support:
- predefined records like traditional routers (**record netflow ipv4 original-input**)
 - flow based sampler.
13. On VLAN interfaces, when you use the **interface** option with the **Cos,Tos**, **TTL** or **Packet length** options, the system displays inaccurate results for the interface input field.
14. The configuration of the flow exporter does not support the option **output features**.
15. Flow aging in flow cache is controlled through active and in-active timer configuration. The minimum for active and in-active aging timers is 5 seconds. The timers must be in units of 5 seconds.



Note Flows in the hardware table are deleted after 5 seconds of in-activity irrespective of the active or in-active timer configuration values. This allows you to create new hardware flows quickly.

16. First and Last-seen flow timestamp accuracy is within 3 seconds.
17. 2048 Flow monitors and records are supported.
- When TTL is configured as a flow field, the following values are reported for a given packet TTL value. [Table 67-4](#) lists the packet TTL and reported values.
18. Cisco TrustSec (CTS) fields are supported. These fields use Netflow collector to monitor and troubleshoot the CTS network, and to segregate traffic based on source group tag (SGT) values.

- When configuring the source group tag (**collect flow cts source group-tag**), note the following:
The system copies the packets to software before it retrieves the CTS field. A large number of flows mean that a large number of packets are copied to the software, possibly affecting CPU performance.
The maximum number of (unique) hosts allowed in the switch (IP addresses) is 12,000.
In case of burst packets, the software may not be able to retrieve the CTS field because the software queue is throttled.
- When configuring the destination group tag (**collect flow cts destination group-tag**), note that this CTS field value is collected only if you have already configured an IP-to-SGT mapping.
- When configuring switch-derived source group tags (**collect flow cts switch derived-sgt**), note that the switch derives this value locally.
- When configuring CTS fields on Supervisor Engine 8-E, note that CTS fields are not supported on wireless interfaces (WLAN) and SSID.

Table 67-4 TTL Map: TTL Configured

Packet TT Value	Reported Value
0	0
1	1
2-10	10
11-25	25
26-50	50
51-100	100
100-150	150
150-255	255

- When packet length is configured as a flow field, the following values are reported for a given packet length value. [Table 67-5](#) lists the packet length and reported values.

Table 67-5 Packet Length Map: Packet Length Configured

Packet Length	Reported Value
0-64	64
65-128	128
129-256	256
257-512	512
513-756	756
757-1500	1500
1500-4000	4000
4000+	8192

The following table lists the options available through FNF and the supported fields.

Table 67-6 Options Available through FNF and the Supported Fields

Field	Description	Comments
Data Link Fields (Layer 2 Flow Label + A94)		
dot1q priority	802.1Q user	
dot1q vlan	802.1Q VLAN ID	Ingress VLAN is supported as key field.
mac destination-address	Upstream destination MAC address	
mac source-address	Down stream source MAC address	
IPv4 Fields		
destination address	IPv4 destination address	Yes
DSCP	IPv4 DSCP (part of TOS)	
fragmentation flags	IPv4 fragmentation flags	Supported as a non-key field. DF flag is not supported
is-multicast	Indicator of an IPv4 multicast packet (0 - if it's not, 1 - if it is)	Supported as a non-key field.
Precedence	IPv4 precedence	
Protocol	IPv4 protocol	
source address	IPv4 source address	
total length	IPv4 datagram	Values are reported based on Table 67-5 .
Total length minimum	Minimum packet size seen	
Total length maximum	Maximum packet size seen	
Tos	IPv4 Type of Service (TOS)	
ttl	Pv4 Time to Live (TTL)	Values are reported based on Table 67-4 .
ttl minimum		Supported as a non-key field.
ttl maximum		Supported as a non-key field.
CTS Fields		

Table 67-6 Options Available through FNF and the Supported Fields

Field	Description	Comments
flow cts destination group-tag		Supported as a non-key field; configuring the IPv4 destination address is a prerequisite to using this field.
flow cts source group-tag		Supported as a non-key field; configuring the IPv4 source address is a prerequisite to using this field.
flow cts switch derived-sgt	Switch-derived source group-tag	Supported as a non-key field; configuring the IPv4 source address is a prerequisite to using this field.
IPv6 Fields		
destination address	IPv6 destination address	
dscp	IPv6 DSCP (part of IPv6 traffic class)	
flow-label	IPv6 flow label	
is-multicast	Indicator of an IPv6 multicast packet (0 - if it's not, 1 - if it is)	Supported as a non-key field
hop-limit	IPv6 hop limit (replaces IPv4 ttl)	Values are reported based on Table 67-4 .
hop-limit minimum	IPv6 minimum hop limit value seen in the flow.	Supported as a non-key field.
hop-limit maximum	IPv6 maximum hop limit value seen in the flow.	Supported as a non-key field.
next-header	IPv6 next header type	Only first next header is reported
total length	IPv6 total packet length	Values are based on Table 67-5 .
Total length minimum	Minimum packet size seen	
Total length maximum	Maximum packet size seen	
protocol	IPv6 next header type in the last IPv6 extension header	

Table 67-6 Options Available through FNF and the Supported Fields

Field	Description	Comments
source address	IPv6 source address	
traffic-class	IPv6 traffic class	Yes
Routing Attributes		
forwarding-status	Forwarding status for the packet (forwarded, terminated in the router, dropped by ACL, RPF, CAR)	Supported as a non-key field
Layer 4 Header Fields		
Field	Description	Comments
TCP Header Fields		
destination-port TCP destination number	TCP destination port	
flags [ack] [fin] [psh] [rst] [syn] [urg]	TCP flags.	Supported as non-key fields.
source-port	TCP source port	
UDP Header Fields		
destination-port	UDP destination port	
source-port	UDP source port	
ICMP Header Fields		
code	ICMP code	
type	ICMP type	
IGMP Header Fields		
type	IGMP	
Interface Fields		
input	Input interface index	
output	Output interface index	Output interface can be supported only as non-key.
Flexible NetFlow feature related fields		
direction: input		
Counter Fields		
bytes	32 bit counters	
bytes long	64 bit counter	
packets	32 bit counters	
packets long	64 bit counter of the packets in the flow	
Timestamp		

Table 67-6 Options Available through FNF and the Supported Fields

Field	Description	Comments
first seen	Time-stamp of the first packet that is accounted in the flow (in milliseconds, starting from the router boot-up)	3 sec accuracy
last seen	Time-stamp of the last packet that is accounted in the flow (in milliseconds, starting from the router boot-up)	3 sec accuracy

Configuring Flow Monitor Cache Values

Setting active cache timeout to a small value may cause the flows to be exported more frequently to the remote collector. This also causes software to delete flows from the local cache after exporting. So, cache statistics reported by switch may not display the actual flows being monitored.

