CHAPTER 33

Configuring Flexible NetFlow

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

The following items apply to the Catalyst 4500 series switch:

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. Supervisor Engine 7-E supports a 100,000 entry hardware flow table, which is shared across all the ports and VLANs on the switch. To limit the number of table entries on a given interface or VLAN, enter the **cache entries** *number* command.

   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 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 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
!
```

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).

Note The switch does not support tunnels and SVI statistics.

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 34-10.)
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 not allowed. However, packet COS and Layer 3 packet field selection is allowed.
10. Only permanent and normal flow cache types are supported.
11. Supervisor Engine 7-E does not support predefined records like traditional routers (**record neflow ipv4 original-input**).
12. Supervisor Engine 7-E does not support 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 33-1 lists the packet TTL and reported values.

Table 33-1 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 33-2 lists the packet length and reported values.

  1

Table 33-2 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 33-3 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 vid	802.1Q VLAN ID	Only output VLAN as collect option is supported.
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 non key. DF flag is not supported
is-multicast	Indicator of an IPv4 multicast packet (0 - if it's not, 1 - if it is)	Supported as non-key
Precedence	IPv4 precedence	
Protocol	IPv4 protocol	
source address	IPv4 source address	
total length	IPv4 datagram	Values are reported based on Table 33-2.
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 33-1.
ttl minimum	FNF supports this field only in mon-key mode	
ttl maximum	FNF supports this field only in mon-key mode	
IPv6 Fields		
destination address	IPv6 destination address	

Table 33-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
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 33-1.
hop-limit minimum	IPv6 minimum hop limit value seen in the flow. It can be used as a non-key field only.	
hop-limit maximum	IPv6 maximum hop limit value seen in the flow. It can be used as a non-key field only.	
next-header	IPv5 next header type	Only first next header is reported
total length	IPv6 total packet length	Values are based on Table 33-2.
Total length minimum	Minimum packet size seen	
Total length maximim	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	

Table 33-3 Options Available through FNF and the Supported Fields

Field	Description	Comments
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	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.