Chapter 4

Configuring J-Flow Statistics

This chapter describes how to configure J-Flow statistics on your ERX router; it contains the following sections:

Overview

The JUNOSe J-Flow feature provides a method by which you can collect IP traffic flow statistics on your routing devices. J-Flow does not require any special protocol for connection setup. It also does not require any external changes to networked traffic, packets, or any other devices in the network. In other words, J-Flow is transparent to the existing network, including end stations and application software and network devices such as LAN switches.

The JUNOSe implementation of J-Flow allows you to export data to the UDP port of a remote workstation for data collection and further processing. In addition, the ability to enable J-Flow on an individual virtual router, interface, or subinterface allows you to collect network statistics for specific locations within your network.

Interface Sampling

For any given IP interface, enabling J-Flow causes packets from the input stream to be sampled at a globally configured rate. For each packet sampled, the main flow cache is examined to see if there is an existing entry. If no entry exists, J-Flow creates a new entry and records attributes of the flow. If the packet matches an existing entry, J-Flow updates the existing flow.

In general, the system samples packets that it can forward. In other words, the system does not sample packets that it discards. As sampling occurs, the system records flow characteristics as they would appear for a packet that the virtual router transmits. This means, for example, that if a packet uses the address of an output interface or next-hop value altered by a policy setting, the system records the altered value in the flow record.

Aggregation Caches

Data from flow cache entries is summarized to build aggregated views or aggregation caches. Aggregation caches are created and maintained along with the main cache. Aggregation caches have their own history area where the aging aggregation cache records are collected. Aggregation caches have a set of configuration parameters: number of entries, active and inactive time out, and export destination.

Types of aggregation caches include:

- AS—Aggregates flow data based on source and destination AS, and ingress and egress interface values.
- Destination Prefix—Aggregates flow data based on the destination address, mask, destination AS, and egress interface.
- Prefix—Aggregates flow data based on source prefix, destination prefix, source mask, destination mask, source AS, destination AS, ingress interface, and egress interface.
- Protocol Port—Aggregates flow data based on protocol, source port, and destination port.
- Source Prefix—Aggregates flow data based on source address, source mask, source AS, and ingress interface.

Aggregation caches contain a subset of the fields collected in the raw flow data. For example, TCP flags, Next Hop Address, and ToS values are not maintained in any of the aggregation caches. Unlike the main cache, aggregation caches are not enabled by default.

Flow Collection

The JUNOSe J-Flow functionality allows statistics collection at the VR/VRF level. This means that each virtual router (VR)/VPN routing and forwarding (VRF) table has its own main cache for statistics gathering.

Although you can export flow statistics only at the VR level, VRF data is rolled up for each VR. The reason for supporting export flow at the VR level is that existing export formats cannot discriminate between VRs and VRFs. However, even though export formats do not allow for segregation, the JUNOSe CLI commands do. Segregating each collection by VR removes any ambiguity and aliasing that may occur with overlapping address spaces (as may occur in virtual private network [VPN] configurations).

Main Flow Cache Contents

The following 7-tuple distinguishes an entry in the flow cache for a VR:

- Source IP address (SA)
- Destination IP address (DA)
- Source port number (SP)

- Destination port number (DP)
- Layer 3 protocol type
- Type of service (ToS byte) or Differentiated Services code point (DSCP)
- Input interface

Cache Flow Export

Using UDP as the transport method, the ERX router can export the content of the flow cache as the system removes the entries. You can specify one export destination for each VR.

Each export packet contains a header and flow records. The version 5 header contains the following fields:

- Version—Format version
- Count—Number of records in this packet
- SysUpTime—System up time value when this packet was built
- Unix Timestamp—Number of seconds and nanoseconds since 0000 UTC 1970 (Coordinated Universal Time)
- Sequence Number—Number of total records sent on this export stream
- Engine type—Type of switching engine (line module or route processor)

NOTE: The J-Flow setting for Engine type is always RP = 0.

- Engine ID—SRP slot number

If, for any reason, the virtual router is unable to export records to the collector, the unsent records are discarded. However, the virtual router continues to increase the sequence number by one as if it sent the records. Discrepancies between the sequence number and sent records can assist in recognizing discontinuities at the collector end.

Aging Flows

After the virtual router creates a flow in the cache, the flow is removed at the expiration of either the active or the inactive timer.

In sampled environments, methods for detecting the end of a flow can be unreliable. The active timer places a hard limit on how long a flow may last before the virtual router closes it and gathers the necessary statistics. If the flow is still active when the active timer expires, the virtual router creates a new flow entry to replace the closed flow.

The inactive timer removes flows if they do not contain any data traffic for a specified period of time.

Operation with NAT

When functioning with Network Address Translation (NAT), J-Flow sampling occurs before NAT applies any translation.

Operation with High Availability

When high availability is enabled, the following occurs in the event of a switchover:

- Any flows that are collected but not exported off of the router are lost.
- Flow history is lost.
- Counters are reset to zero.

After the standby SRP becomes active, and all other applications indicate that they have recovered, sampling and flow-collecting resume.

Platform Considerations

For information about modules that support J-Flow statistics on ERX-14xx models, ERX-7xx models, and the ERX-310 router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support NAT.

For information about modules that support J-Flow on the E120 router and the E320 router:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support J-Flow.

Before You Configure J-Flow Statistics

Before you configure J-Flow statistics, be sure you have created IP interfaces from which J-Flow can extract traffic flow information. For information about configuring IP interfaces, see *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 1, Configuring IP*.

Configuring Flow-Based Statistics Collection

To configure J-Flow on a virtual router:

1. Enable J-Flow statistics.
2. Enable J-Flow statistics on the desired interfaces.
3. (Optional) Define the sampling interval at which you want to collect statistics.
4. (Optional) Customize the size of the main flow cache.
5. (Optional) Define flow cache aging timers.
6. (Optional) Specify to where you want to export J-Flow statistics.

Enabling Flow-Based Statistics

Use the **ip flow statistics** command to explicitly enable J-Flow.

NOTE: Issuing any configuration-level commands implicitly enables J-Flow.

ip flow statistics

- Use to enable J-Flow.
- Example

  host1(config)#**ip flow statistics**

- Use the **no** version to disable J-Flow on the virtual router.

Enabling Flow-Based Statistics on an Interface

Use the **ip route-cache flow sampled** command to enable J-Flow statistics on an interface. You can also use this command to configure an IP profile that is applied to dynamically created IP interfaces. This feature provides J-flow capability on all dynamically created IP interfaces, including those used for MPLS-to-IP forwarding scenarios.

NOTE: Issuing an interface-level flow command does not enable J-Flow on the virtual router. To enable J-Flow, issue the **ip flow statistics** command.

ip route-cache flow sampled

- Use to enable J-Flow on an interface. or in an IP profile for dynamically created IP interfaces.
- Examples

  host1(config-if)#**ip route-cache flow sampled**

  or

  host1(config-profile)#**ip route-cache flow sampled**

- Use the **no** version to disable J-Flow statistics on the interface.

Defining a Sampling Interval

Use the **ip flow-sampling-mode packet-interval** command to define the packet-sampling interval for the virtual router. The sampling interval specifies the rate at which the virtual router samples J-Flow information. This rate is used for all interfaces that have J-Flow enabled. After you enable J-Flow on an interface, the virtual router samples one packet at the specified packet interval. You can specify an interval in the range 1–4,000,000,000 packets.

When you use the **ip flow-sampling-mode packet-interval** command to define the packet-sampling interval for Gigabit Ethernet interfaces configured on the ES2 10G LM (line module) with either the ES2-S1 GE-8 IOA or the ES2-S2 10GE PR IOA on E120 routers and E320 routers, the J-Flow application makes the following internal adjustments to achieve better performance on the ES2 10G LM, regardless of the packet-sampling interval that you configure:

- J-Flow adjusts the maximum sampling interval to 8,388,608, which is the decimal equivalent of 0x800000.
- J-Flow changes the packet-sampling value to the closest integer that is a power of two and that is less than or equal to the configured value.

For performance reasons, J-Flow applies these adjustments to the sampling interval only for the interfaces configured on the ES2 10G LM on the virtual router. The configured sampling interval does not change for interfaces not configured on the ES2 10G LM on the virtual router.

When the data rate increases on a given interface, J-Flow packet sampling might not be able to maintain the configured sampling rate and might drop the intended sampled packets. If this occurs, you can address the issue by reducing the sampling rate.

NOTE: For all modules except the ES2 10G LM on the E120 router and the E320 router, packet sampling occurs individually for each processor. Because the router distributes packets over multiple processors, sampling occurs when each processor reaches the specified packet interval.

NOTE: Even though each flow is sampled, the flow sample is not necessarily cached because of system constraints.

ip flow-sampling-mode packet-interval

- Use to define the J-Flow packet-sampling interval.
- Specify a packet-sampling interval in the range 1–4000000000 packets; the default value is 4000000000.
- Specifying an interval less than 10 sets a very high sampling rate that can severely degrade performance. The lower the packet-sampling interval you configure, the faster the sampling rate.
- For information about the effects of using the **ip flow-sampling-mode packet-interval** command for the ES2 10G LM with either the ES2-S1 GE-8 IOA or the ES2-S2 10GE PR IOA on E120 routers and E320 routers, see *Defining a Sampling Interval* on page 116.
- Example—Samples 1 out of 50 packets from the line module on which the interface resides

  host1(config)#**ip flow-sampling-mode packet-interval 50**

- Use the **no** version to return the sampling interval to its default value, 4 billion.

Setting Cache Size

Use the **ip flow-cache entries** command to limit the number of main flow cache entries for the virtual router (as collected across all line modules that are running J-Flow). After the cache size exceeds the flow-cache entry limit, the least recently used flow is removed.

The possible flow-cache range is 1,024 – 524,288 entries. The default value is 65,536 entries.

ip flow-cache entries

- Use to limit J-Flow main flow cache entries.
- Example

  host1(config)#**ip flow-cache entries 80000**

- Use the **no** version to return the cache size to its default value, 65535.

Defining Aging Timers

After the virtual router creates a flow in the cache, the virtual router can remove the flow at the expiration of either the active or the inactive timer.

Specifying the Activity Timer

Use the **ip flow-cache timeout active** command to specify a value for the activity timer. The activity timer measures the amount of time that the virtual router has been recording a datagram for a given flow. When this timer expires, the virtual router exports the flow cache entry from the cache and removes the entry. This process prevents active flows from remaining in the flow cache, and allows collected data to appear in a timely manner. The possible range for the activity timer value is 1 – 60 minutes. The default value is 30 minutes.

ip flow-cache timeout active

- Use to define the activity timer, in minutes.
- Example

  host1(config)#**ip flow-cache timeout active 50**

- Use the **no** version to return the activity timer to its default value (30 minutes).

Specifying the Inactivity Timer

Use the **ip flow-cache timeout inactive** command to specify a value for the inactivity timer. The inactivity timer measures the length of time expired since the virtual router recorded the last datagram for a given flow. When this timer expires, the virtual router exports the flow cache entry from the cache and removes it. When, at a later time, another datagram begins that uses the same flow characteristics, the virtual router allocates a new flow cache entry, and the inactivity timer begins again. The possible range for the inactivity timer value is 10 – 600 seconds. The default value is 15 seconds.

ip flow-cache timeout inactive

- Use to define the inactivity timer, in seconds.
- Example

  host1(config)#**ip flow-cache timeout inactive 90**

- Use the **no** version to return the inactivity timer to its default value (15 seconds).

Specifying Flow Export

Use the **ip flow-export** command to specify the location to which you want to export the J-Flow datagrams.

ip flow-export

- Use to specify the location to which you want to export J-Flow datagrams or specify an alternate source address for outbound export J-Flow datagrams.
- Example 1—Specifies the destination address for J-Flow datagrams

  host1(config)#**ip flow-export 192.168.2.73 2055 version 5 peer-as**

- Example 2—Specifies the source address for outbound export J-Flow datagrams

  host1(config)#**ip flow-export source fastEthernet 5/0**

- Use the **no** version to remove the export setting.

Configuring Aggregation Flow Caches

Aggregation caches are disabled by default. Exporting flow records from the router does not occur while it is in the disabled state. When the configuration for an aggregation cache is changed from enabled to disabled state, all flow records from that cache are removed and flow collection stops.

For Prefix, Destination Prefix, and Source Prefix aggregation caches, you can specify a minimum source and destination mask size to affect the granularity of the IP address space captured in the aggregation cache. The commands to configure the minimum mask size for the source and destination address are issued in Flow Cache Configuration mode and are specific to each aggregation cache:

host1(config-flow-cache)#**mask source minimum *value***

host1(config-flow-cache)#**mask destination minimum *value***

The value (a number in the range 1–32) specifies the size of the minimum mask. The **no** version restores the default minimum mask size, which is 0. A mask of size *N* has the *N* most significant bits set in the corresponding bit mask.

You cannot configure a minimum mask size for aggregation caches that do not retain an IP address in their aggregation scheme (like the AS aggregation cache). You can configure the Prefix aggregation cache for both source and destination minimum mask size. You can configure only the source minimum mask size for the Source Prefix aggregation cache. You can configure only the destination minimum mask size for the Destination Prefix aggregation cache.

The peer/origin information configured with the export command for the man V5 cache is used to display the AS number of the AS aggregation cache for both the source and destination AS. If no (default) configuration is present, zero appears in the AS numbers for both V5 export and V8 export and in the **show** commands for the V8 AS aggregation cache.

Establish an aggregation cache:

1. Enter Flow Cache Configuration mode for the AS aggregation cache.

   host1(config)#**ip flow-aggregation cache as**

2. Configure the number of entries (1024—524288) in the aggregation cache; the **no** version sets the number of entries back to its default value of 4096 (flow-data may be lost if the previous setting is larger than the default).

   host1(config-flow-cache)#**cache entries *entryNumber***

3. Set the active (1—60) and inactive (10—600) aging timers.

   host1(config-flow-cache)#**cache timeout active *active-tmo***
   host1(config-flow-cache)#**cache timeout inactive *inactive-tmo***

4. Configure an export destination for the aggregation cache; the **no** version removes the destination.

   host1(config-flow-cache)#**export destination {*hostname* | *ip address*} *udp-port-number***

5. Set the source IP address for datagrams containing information from this cache: the **no** version removes the explicit setting of the source address.

   host1(config-flow-cache)#**export source *interfacetype* interface**

6. Enable the aggregation cache.

   host1(config-flow-cache)#**enabled**

   The aggregation cache starts accumulating information from the flow cache; the **no** version stops the accumulation of information from the flow cache, but does not suspend the operation of the flow cache.

cache entries

- Use to set the number of entries in the aggregation cache.
- Example

  host1(config-flow-cache)#**cache entries 524288**

- Use the **no** version to reset the number of entries to the default value 4096.

cache timeout

- Use to set the active and inactive timers.
- Example

  host1(config-flow-cache)#**cache timeout active 50**

- Use the **no** version to reset the default value.

enabled

- Use to enable the aggregation cache to accumulate information form the flow cache.
- Example

  host1(config-flow-cache)#**enabled**

- Use the **no** version to stop the information flow from the flow cache.

export destination

- Use to configure an export destination for the aggregation cache.
- Example

  host1(config-flow-cache)#**export destination myhost udp-port**

- Use the **no** version to remove the destination.

export source

- Use to configure an export source for the aggregation cache.
- Example

  host1(config-flow-cache)#**export source interface inf1**

- Use the **no** version to remove the destination.

ip flow-aggregation cache

- Use to create an aggregation cache.
- Example

  host1(config)#**ip flow-aggregation cache**

- Use the **no** version to remove the aggregation cache and its configuration.

mask destination

- Use to set the minimum mask size for the destination address for the prefix and destination prefix aggregation caches.
- Example

  host1(config-flow-cache)#**mask destination 128**

- Use the **no** version to restore the default mask size, which is 0.

mask source

- Use to set the minimum mask size for the source address for the prefix and source prefix aggregation caches.
- Example

  host1(config-flow-cache)#**mask source 60**

- Use the **no** version to restore the default mask size, which is 0.

Monitoring J-Flow Statistics

This section shows how to clear J-Flow statistics and use the **show** commands to view J-Flow settings and statistical results.

Clearing J-Flow Statistics

Use the **clear ip flow stats** command to clear all entries from all flow caches on the virtual router.

clear ip flow stats

- Use to clear entries from all flow caches on the VR/VRF.
- Example

  host1(config)#**clear ip flow stats**

- There is no **no** version.

J-Flow show Commands

You can monitor the following aspects of J-Flow statistics by using the following commands:

To Display	Command
Main cache flow operational statistics	**show ip cache flow**
J-Flow sampling state	**show ip flow sampling**
J-Flow export state and export statistics	**show ip flow export**

You can use the output filtering feature of the **show** command to include or exclude lines of output based on a text string that you specify. See *JUNOSe System Basics Configuration Guide, Chapter 2, Command-Line Interface*, for details.

show ip cache flow

- Use to display IP flow cache operational statistics.
- Field descriptions
  - Main Cache
    - Max Entries—Maximum number of entries allowed in the main cache
    - Activity Timeout—Activity timer value
    - Inactivity Timeout—Inactivity timer value
    - Size—Distribution of IP packets by size
    - Percent—Percent distribution of different-sized IP packets
  - Protocol - Port—Protocol of the sample and port destination for that sample
  - Total Flows—Total number of flows

- Flows/Sec—Number of flows per second
- Packets/Flow—Number of packets per flow
- Bytes/Packet—Number of bytes per packet
- Packets/Sec—Number of packets per second
- Src. Addr—Source address of sampled packets
- Src. Intf—Source interface of sampled packets
- Dst. Addr—Destination address of sampled packets
- Dst. Intf—Destination interface of sampled packets
- Summary
  - Total Flows Processed—Total number of flows processed
  - Total Packets—Total number of packets sampled
  - Total Bytes—Total number of bytes received

- Example 1—Brief output

```
host1#show ip cache flow active brief
29140 packets sampled.
Distribution of IP packets by size.

   Size         Percent
----------      -------
    1 - 32        0.000
       64         0.000
       96         0.000
      128         0.000
      160         0.000
      192         0.000
      224         0.000
      256         0.000
      288         0.000
      320         0.000
      352         0.000
      384         0.000
      416         0.000
      448         0.000
      480         0.000
      512         0.000
      544         0.000
      576         0.000
     1024        96.791
     1536         3.209
     2048         0.000
     2560         0.000
     3072         0.000
     3584         0.000
     4096         0.000
     4608         0.000

                     Total         Flows       Packets      Bytes      Packets
   Protocol-Port     Flows          /Sec         /Flow       /Pkt         /Sec
-------------------  ----------  ---------  ------------  ---------  -----------
TCP-telnet                1      0.000       118.000   1014.000        0.000
UDP-whois++               1      0.008       935.000   1026.000        7.664
```

```
------------------ Summary ---------------------
Total Flows Processed: 2
Total Packets 1053
Total Bytes   1078962
------------------------------------------------
```

■ Example 2—Detailed output

NOTE: The output format for this command was modified slightly to fit within the confines of this document.

```
host1#show ip cache flow active detail
Main Cache
Max Entries: 65536
Activity Timeout: 60 mins.
Inactivity Timeout: 600 secs.
Cache Enabled

32012 packets sampled.
Distribution of IP packets by size.

   Size         Percent
----------      -------
    1 - 32        0.000
       64         0.000
       96         0.000
      128         0.000
      160         0.000
      192         0.000
      224         0.000
      256         0.000
      288         0.000
      320         0.000
      352         0.000
      384         0.000
      416         0.000
      448         0.000
      480         0.000
      512         0.000
      544         0.000
      576         0.000
     1024        96.789
     1536         3.211
     2048         0.000
     2560         0.000
     3072         0.000
     3584         0.000
     4096         0.000
     4608         0.000

                                                                    Packets   Bytes       Packets
   Src.Addr     Src.Intf     Dst.Addr     Dst.Intf     Protocol Port      /Flow     /Packet     /Sec
-------------  ---------    ---------    ----------   ---------------   ---------  --------- ---------
10.20.30.41 258  GigE4/0    12.0.0.2     GigE2/0      TCP-telnet         58.000    1014.000   0.000
10.20.30.41 63   GE4/0      50.60.70.88               UDP-whois++        1028.000  1026.000   7.672

------------------ Summary ---------------------
Total Flows Processed: 2
Total Packets 1086
Total Bytes   1113540
------------------------------------------------
```

- Example 3—History output

```
host1#show ip cache flow history
35604 packets sampled.
Distribution of IP packets by size.

   Size         Percent
----------      -------
    1 - 32        0.000
       64         0.000
       96         0.000
      128         0.000
      160         0.000
      192         0.000
      224         0.000
      256         0.000
      288         0.000
      320         0.000
      352         0.000
      384         0.000
      416         0.000
      448         0.000
      480         0.000
      512         0.000
      544         0.000
      576         0.000
     1024        96.784
     1536         3.216
     2048         0.000
     2560         0.000
     3072         0.000
     3584         0.000
     4096         0.000
     4608         0.000

                      Total         Flows     Packets      Bytes      Packets
   Protocol Port      Flows          /Sec       /Flow       /Pkt        /Sec
------------------  ----------  ---------  ------------  ---------  -----------
TCP-telnet                 216      1.450       159.264   1014.000      230.879

------------------ Summary ---------------------
Total Flows Processed: 216
Total Packets 34401
Total Bytes   34882614
------------------------------------------------
```

show ip cache flow aggregation

- Use to display IP flow cache operational statistics for an aggregation cache.
- Field descriptions
  - Aggregation Cache
    - AS—AS aggregation cache
    - Destination-prefix—Destination-prefix aggregation cache
    - Prefix—Prefix aggregation cache
    - Protocol-port—Protocol-port aggregation cache
    - Source-prefix—Source-prefix aggregation cache
  - Total Flows—Total number of flows
  - Flows/Sec—Number of flows per second
  - Packets/Flow—Number of packets per flow
  - Bytes/Packet—Number of bytes per packet
  - Packets/Sec—Number of packets per second
  - Src. Addr—Source address of sampled packets
  - Src. Intf—Source interface of sampled packets
  - Dst. Addr—Destination address of sampled packets
  - Dst. Intf—Destination interface of sampled packets
  - Summary
    - Total Flows Processed—Total number of flows processed
    - Total Packets—Total number of packets sampled
    - Total Bytes—Total number of bytes received

Example—Aggregation cache flow output

```
host1#show ip cache flow aggregation as active brief
29140 packets sampled.

                                  Total
   Src AS              Dest AS     Flows     Packets/Flows   Bytes/Pkt
Packets/Sec
-----------------  ----------   -------------- ----------------  ---------
400                   100         0.000         118.000     1014.000
0.000
100                    400        0.008         935.000     1026.000
7.664

------------------ Summary ---------------------
Total Flows Processed: 2
Total Packets 1053
Total Bytes   1078962
-------------------------------------------------
```

show ip flow export

- Use to display configuration values for IP flow cache export.
- Example

```
host1#show ip flow export
Flow export is enabled using version 5 format.
Exporting to 10.0.0.2 port 9898 using source ip interface
GigabitEthernet5/0/0.
```

show ip flow sampling

- Use to display configuration values for IP flow cache sampling.
- Example

```
host1#show ip flow sampling
Flow sampling is enabled
'Packet Interval' sampling mode is configured.
1 out of every 1000 packets is being sampled.
```