

Chapter 14

Traffic Sampling and Forwarding Configuration

To configure forwarding options and traffic sampling, include statements at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
accounting group-name {
  output {
    cflowd [ host-names ] {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-names ] {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
family (inet | inet6 | mpls) {
  filter {
    input filter-name;
  }
}
hash-key {
  family inet {
    layer-3;
    layer-4;
  }
  family mpls {
    label-1;
    label-2;
  }
}
```

```

helpers {
  bootp {
    description description-of-service;
    interface interface-group {
      description description-of-interface;
      maximum-hop-count number;
      minimum-wait-time seconds;
      no-listen;
      server [ addresses ]
      maximum-hop-count number;
      minimum-wait-time seconds;
      server address < [ routing-instance routing-instance-name ] >;
    }
  }
  domain {
    description description-of-service;
    server address;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address < [ routing-instance routing-instance-name ] >;
    }
  }
}
tftp {
  description description-of-service;
  server address < [ routing-instance routing-instance-name ] >;
  interface interface-name {
    description description-of-interface;
    no-listen;
    server address;
  }
}
}
traceoptions {
  file filename {
    files number;
    size bytes;
  }
  flag flag;
  level level;
}
}
monitoring group-name {
  family inet {
    output {
      cflowd [ host-names ] {
        aggregation {
          autonomous-system;
          destination-prefix;
          protocol-port;
          source-destination-prefix {
            caida-compliant;
          }
          source-prefix;
        }
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
  }
}

```

```

export-format cflowd-version-5;
flow-active-timeout seconds;
flow-inactive-timeout seconds;
interface [ interface-names ] {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
}
}
}
next-hop-group [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
    }
}
port-mirroring {
    input {
        family inet {
            rate num;
            run-length num;
        }
    }
    output {
        interface interface-name {
            next-hop address;
        }
        no-filter-check;
    }
}
sampling {
    disable;
    input {
        family inet {
            max-packets-per-second number;
            rate number;
            run-length number;
        }
    }
    output {
        cflowd [ host-names ] {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version format;
        }
    }
}

```

```

file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}
flow-active-timeout seconds;
flow-inactive-timeout seconds;
interface [ interface-names ] {
  engine-id number;
  engine-type number;
  source-address address;
}
}
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
}

```

This chapter describes the following tasks for configuring traffic sampling and forwarding options:

- Minimum Traffic Sampling or Forwarding Configuration on page 219
- Configure a Forwarding Table Filter on page 220
- Configure Traffic Sampling on page 220
- Configure Discard Accounting on page 222
- Configure Flow Monitoring on page 223
- Configure a Next-Hop Group on page 224
- Configure Per-Flow Load-Balancing Information on page 224
- Configure the Router or Interface to Act as a DHCP/BOOTP Relay Agent on page 225
- Configure DNS and TFTP Packet Forwarding on page 226
- Disable Traffic Sampling on page 227
- Examples: Configure Traffic Sampling on page 228
- Configure Traffic Sampling Output on page 231
- Trace Traffic Sampling Operations on page 233
- Configure Flow Aggregation (cflowd) on page 233
- Configure Port Mirroring on page 236

Minimum Traffic Sampling or Forwarding Configuration

To configure traffic sampling, you must perform at least the following tasks:

Create a firewall filter to apply to the logical interfaces being sampled by including the filter statement at the [edit firewall family family-name] hierarchy level. In the filter then statement, you must specify the action modifier `sample` and the action `accept`.

```
[edit firewall family family-name]
filter filter-name {
  term term-name {
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the interfaces on which you want to sample traffic:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    family family-name {
      filter {
        input filter-name;
      }
      address address {
        destination destination-address;
      }
    }
  }
}
```

Enable sampling and specify a nonzero sampling rate:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate number;
    }
  }
}
```

Configure a Forwarding Table Filter

A forwarding table filter allows you to filter data packets based on their components and to perform an action on packets that match the filter.

To apply a forwarding table filter to a forwarding table, include the filter input statement at the [edit forwarding-options family *family*] hierarchy level:

```
[edit]
forwarding-options {
  family family {
    filter {
      input filter-name;
    }
  }
}
```

Configure Traffic Sampling

On routers containing the Monitoring Services PIC or the Generic Services PIC, you can configure traffic sampling for traffic passing through the router.

To configure traffic sampling on a logical interface, include the input statement at the [edit forwarding-options sampling] hierarchy level:

```
[edit forwarding-options sampling]
input {
  family inet {
    max-packets-per-second number;
    rate number;
    run-length number;
  }
}
```


Specify the threshold traffic value by using the max-packets-per-second statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.

Specify the sampling rate by setting the values for rate and run-length (see Figure 13).

Figure 13: Configure Sampling Rate

Rate and Run-length

Sampling rate = (run-length + 1) / rate



1746

The rate statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where $x = \text{run-length} + 1$. By default, the rate is 0, which means that no traffic is sampled.

The run-length statement specifies the number of matching packets to sample following the initial one-packet trigger event. By default, the run-length is 0, which means that no more traffic is sampled after the trigger event. The range is 0 through 20. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.

If you do not include the input statement, sampling is disabled.

To collect the sampled packets in a file, include the file statement at the [edit forwarding-options sampling output] hierarchy level. For more information about the output file formats, see “Configure Traffic Sampling Output” on page 231.

JUNOS does not sample packets originating from the router. If you configure a sampling filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from routing engine (RE) to packet forwarding engine (PFE) are not sampled.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Service PIC. To allow responses from the Monitoring Service PIC to pass through for sampling purposes, configure a term in the firewall filter to include the Monitoring Services PIC IP address. For more detailed information about configuring firewall filters, see “Firewall Filter Configuration” on page 143.

Configure Discard Accounting

On routers containing the Monitoring Services PIC or the Generic Services PIC, you can configure accounting for traffic passing through the router.

To configure discard accounting, include the accounting statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
accounting group-name {
  output {
    cflowd [ host-names ] {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      autonomous-system-type (origin | peer);
      port port-number;
      version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface [ interface-name ] {
      engine-id number;
      engine-type number;
      source-address address;
    }
  }
}
```

To configure an accounting group, include the accounting statement and specify a *group-name*. To configure the output flow aggregation, include the *cflowd* statement. For more information about flow aggregation, see [Configure Flow Aggregation \(cflowd\)](#) on page 233. To configure the interval before exporting an active flow, include the *flow-active-timeout* statement. The default value for *flow-active-timeout* is 1800 seconds. To configure the interval before a flow is considered inactive, include the *flow-inactive-timeout* statement. The default value for *flow-inactive-timeout* is 60 seconds. To configure the interface that sends out monitored information, include the interface statement. Discard accounting is supported for Monitoring Services PIC interfaces only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Service PIC. To allow responses from the Monitoring Service PIC to pass through for accounting purposes, configure a term in the firewall filter to include the Monitoring Services PIC IP address. For more detailed information about configuring firewall filters, see [“Firewall Filter Configuration”](#) on page 143.

You can use discard accounting for passive and active flow monitoring. For more detailed information about configuring passive and active flow monitoring, see the *JUNOS Internet Software Feature Guide* and the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

Configure Flow Monitoring

On routers containing the Monitoring Services PIC or the Monitoring Services II PIC, you can configure flow monitoring for traffic passing through the router. This type of monitoring method is passive monitoring.

To configure flow monitoring, include the monitoring statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
monitoring group-name {
  family inet {
    output {
      cflowd [ host-names ] {
        aggregation {
          autonomous-system;
          destination-prefix;
          protocol-port;
          source-destination-prefix {
            caida-compliant;
          }
          source-prefix;
        }
        autonomous-system-type (origin | peer);
        port port-number;
        version format;
      }
      export-format cflowd-version-5;
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      interface [ interface-names ] {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
      }
    }
  }
}
```

To configure a passive monitoring group, include the monitoring statement and specify a group name. To configure monitoring on a specified address family, include the family statement and specify an address family. To specify an interface to monitor incoming traffic, include the input statement. To configure the monitoring information that is sent out, include the output statement. To configure the output flow aggregation, include the cflowd statement. For more information about flow aggregation, see [Configure Flow Aggregation \(cflowd\)](#) on page 233. To specify the format of the monitoring information sent out, include the export-format statement and specify a version number. To configure the interval before exporting an active flow, include the flow-active-timeout statement. The default value for flow-active-timeout is 1800 seconds. To configure the interval before a flow is considered inactive, include the flow-inactive-timeout statement. The default value for flow-inactive-timeout is 60 seconds. To configure the interface that sends out the monitored information, include the interface statement. Flow monitoring is supported for Monitoring Services PIC interfaces only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Service PIC. To allow responses from the Monitoring Service PIC to pass through for monitoring purposes, configure a term in the firewall filter to include the Monitoring Services PIC IP address. For more detailed information about configuring firewall filters, see “Firewall Filter Configuration” on page 143.

For more detailed information about configuring the passive and active flow monitoring, see the *JUNOS Internet Software Feature Guide* and the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

Configure a Next-Hop Group

Next-hop groups allow you to include multiple interfaces used to forward duplicate packets used in port mirroring.

To configure a next-hop group, include the `next-hop-group` statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
next-hop-group [ group-names ] {
  interface interface-name {
    next-hop [ addresses ];
  }
}
```

You can specify one or more group names. To configure the interface that sends out sampled information, include the interface statement and specify an interface. To specify a next-hop address to send sampled information, include the `next-hop` statement and specify an IP address.

Next-hop groups have the following restrictions:

- Next-hop groups are supported for IPv4 addresses only.

- Next-hop groups are supported for M-series routers only.

- Next-hop groups support up to 16 next-hop addresses.

- You can configure up to 30 next-hop groups.

- Each next-hop group must have at least two next-hop addresses.

Next-hop groups can be used for port mirroring. For more information about configuring port mirroring, see “Configure Port Mirroring” on page 236 and the *JUNOS Internet Software Feature Guide*.


Configure Per-Flow Load-Balancing Information

By default, when there are multiple equal-cost paths to the same destination, the JUNOS software chooses one of the next-hop addresses at random. On routers with the Internet Processor II ASIC, you have two additional options. You can specify what information the router uses for per-flow load balancing based on port data (instead of on source and destination IP addresses only). For aggregated Ethernet and aggregated SONET interfaces, you can load-balance based on the Multiprotocol Label Switching (MPLS) label information. For more information, see “Configure Load-Balance Per-Packet Action” on page 127.

Configure the Router or Interface to Act as a DHCP/BOOTP Relay Agent

You can configure the router or an interface to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that a locally attached host can issue a DHCP or BOOTP request as a broadcast message. If the router or an interface sees this broadcast message, it relays the message to a specified DHCP or BOOTP server.

You should configure the router or an interface to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.



To configure the router to act as a DHCP/BOOTP relay agent, include the bootp statement at the [edit forwarding-options helpers] hierarchy level:

```
[edit forwarding-options helpers]
bootp {
  description description-of-service;
  interface interface-group {
    description;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server [ addresses ];
  }
  maximum-hop-count number;
  minimum-wait-time seconds;
  server address < [ routing-instance routing-instance-names ] >;
```

To set the description of the BOOTP service, DHCP service, or interface, include the description statement.

To set a logical interface or a group of logical interfaces with a specific DHCP-relay or BOOTP configuration, include the interface statement.

To set the routing instance of the server to forward, include the routing-instance statement. You can include as many routing instances as necessary in the same statement.

To stop packets from being forwarded on a logical interface, a group of logical interfaces, or router, include the no-listen statement.

To set the maximum allowed number in the hops field of the BOOTP header, include the maximum-hop-count statement. Headers that have a larger number in the hops field are not forwarded. If you omit the maximum-hop-count statement, the default value is 4 hops.

To set the minimum allowed number of seconds in the secs field of the BOOTP header, include the minimum-wait-time statement. Headers that have a smaller number in the secs field are not forwarded. If you omit the minimum-wait-time statement, the default value is 3 seconds.

To set the IP address or addresses that specify the DHCP or BOOTP server for the router or interface, include the server statement. You can include as many addresses as necessary in the same statement.

You can also configure an individual logical interface to be a DHCP/BOOTP relay if you have locally attached hosts and a remote DHCP or BOOTP server at the [edit interfaces] hierarchy level. For more information, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

Configure DNS and TFTP Packet Forwarding

You can configure the router to support DNS and TFTP packet forwarding for IPv4 traffic, which allows clients to send DNS or TFTP requests through the router. The responding DNS or TFTP server recognizes the client address and sends a response directly to that address. By default, the router ignores DNS and TFTP request packets.

To enable DNS or TFTP packet forwarding, include the helpers statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
helpers {
  domain {
    description description-of-service;
    server address < [ routing-instance routing-instance-names ] >;
    interface interface-name {
      description description-of-interface;
      no-listen;
      server address < [ routing-instance routing-instance-names ] >;
    }
  }
}
tftp {
  description description-of-service;
  server address;
  interface interface-name {
    description description-of-interface;
    no-listen;
    server address < [ routing-instance routing-instance-names ] >;
  }
}
}
```

To set domain packet forwarding, include the domain statement.

To set the description of the DNS or TFTP service, include the description statement.

To set TFTP packet forwarding, include the tftp statement.

To set a DNS or TFTP server (with an IPv4 address), include the server statement. Use one address for either a global configuration or for each interface.

To set the routing instance of the server to forward, include the routing-instance statement. You can include as many routing instances as necessary in the same statement.

To disable recognition of DNS or TFTP requests on one or more interfaces, include the no-listen statement. If you do not specify at least one interface with this statement, the forwarding service is global to all interfaces on the router.

Trace BOOTP, DNS, and TFTP Forwarding Operations

Tracing operations track all traffic forwarding operations and record them in a log file in the /var/log directory. By default, this file is named /var/log/fud. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace DNS and TFTP forwarding operations, include the traceoptions statement at the [edit forwarding-options helpers] hierarchy level:

```
[edit forwarding-options helpers]
traceoptions {
  file filename {
    files number;
    size bytes;
  }
  flag flag;
  level level;
}
```

Example: Configure DNS Packet Forwarding

Enable DNS packet request forwarding to all interfaces on the router except t1-1/1/2 and t1-1/1/3:

```
[edit forwarding-options helpers]
dns {
  server 10.10.10.30;
  interface {
    t1-1/1/2 {
      no-listen;
      server 10.10.10.9;
    }
    t1-1/1/3 {
      no-listen;
      server 10.10.10.4;
    }
  }
}
```

Disable Traffic Sampling

To explicitly disable traffic sampling on the router, include the disable statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
sampling {
  disable;
}
```

Examples: Configure Traffic Sampling

The following sections provide examples of configuring traffic sampling:

Sample a Single SONET Interface on page 228

Sample All Traffic from a Single IP Address on page 229

Sample All FTP Traffic on page 230

Sample a Single SONET Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET interface and collects it in a file named sonet-samples.txt.

Create the filter:

```
[edit firewall family inet]
filter {
  input sample-sonet {
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the SONET interface:

```
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input sample-sonet;
      }
      address 216.127.68.254/32 {
        destination 216.127.74.7;
      }
    }
  }
}
```

Finally, configure traffic sampling:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 100;
      run-length 2;
    }
  }
}
```

```

output {
  file {
    filename sonet-samples.txt;
    files 40;
    size 5m;
  }
}

```

Sample All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 168.45.92.31, and collects it in a file named samples-168-45-92-31.txt.

Create the filter:

```

[edit firewall family inet]
filter one-ip {
  term get-ip {
    from {
      source-address 168.45.92.31;
    }
    then {
      sample;
      accept;
    }
  }
}

```

Apply the filter to the Gigabit Ethernet interface:

```

[edit interfaces]
ge-4/1/1 {
  unit 0 {
    family inet {
      filter {
        input one-ip;
      }
      address 215.45.92.254;
    }
  }
}

```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```

[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 1;
    }
  }
}

```

```

output {
  file {
    filename samples-215-45-92-31.txt;
    files 100;
    size 100k;
  }
}

```

Sample All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using the FTP data transfer protocol in the output path of a specific T3 interface, and collects the information in a file named t3-ftp-traffic.txt.

Create a filter:

```

[edit firewall family inet]
filter ftp-stats {
  term ftp-usage {
    from {
      destination-port [ftp ftp-data];
    }
    then {
      sample;
      accept;
    }
  }
}

```

Apply the filter to the T3 interface:

```

[edit interfaces]
t3-7/0/2 {
  unit 0 {
    family inet {
      filter {
        input ftp-stats;
      }
      address 141.35.78.254/32 {
        destination 141.35.78.4;
      }
    }
  }
}

```

Finally, gather statistics on 10 percent of the candidate samples:

```

[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 10;
    }
  }
}

```

```

output {
  file {
    filename t3-ftp-traffic.txt;
    files 50;
    size 1m;
  }
}

```

Configure Traffic Sampling Output

You configure traffic sampling results to a file in the `/var/tmp` directory. To collect the sampled packets in a file, include the file statement at the [edit forwarding-options sampling output] hierarchy level:

```

[edit forwarding-options sampling output]
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}

```

To configure the period of time before an active flow is exported, include the `flow-active-timeout` statement at the [edit forwarding-options sampling output] hierarchy level:

```

[edit forwarding-options sampling output]
flow-active-timeout seconds;

```

To configure the period of time before a flow is considered inactive, include the `flow-inactive-timeout` statement at the [edit forwarding-options sampling output] hierarchy level:

```

[edit forwarding-options sampling output]
flow-inactive-timeout seconds;

```

To configure the interface, include the interface statement at the [edit forwarding-options sampling output] hierarchy level:

```

[edit forwarding-options sampling output]
interface [ interface-names ] {
  engine-id num;
  engine-type num;
  source-address address;
}

```

To configure the interval before exporting an active flow, include the `flow-active-timeout` statement. To configure the interval before a flow is considered inactive, include the `flow-inactive-timeout` statement. To configure the interface that sends out monitored information, include the interface statement.

Traffic Sampling Output Files

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the `/var/tmp` directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
# Apr 7 15:48:50
Time                Dest          Src Dest  Src Proto TOS  Pkt  Intf  IP  TCP
                   addr          addr port  port          len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195 0    0    1  0x0  84  8   0x0  0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195 0    0    1  0x0  84  8   0x0  0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195 0    0    1  0x0  84  8   0x0  0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195 0    0    1  0x0  84  8   0x0  0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195 0    0    1  0x0  84  8   0x0  0x0
```

The output contains the following fields:

Time—Time at which the packet was received (displayed only if you include the stamp statement in the configuration)

Dest addr—Destination IP address in the packet

Src addr—Source IP address in the packet

Dest port—TCP or UDP port for the destination address

Src port—TCP or UDP port for the source address

Proto—Packet's protocol type

TOS—Contents of the type-of-service (ToS) field in the IP header

Pkt len—Length of the sampled packet, in bytes

Intf num—Unique number that identifies the sampled logical interface

IP frag—IP fragment number, if applicable

TCP flags—Any TCP flags found in the IP header

To set the timestamp option for the file `my-sample`, enter the following:

```
[edit forwarding-options sampling output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the stamp option, the Time field is displayed.

```
# Apr 7 15:48:50
# Time                Dest          Src Dest  Src Proto TOS  Pkt  Intf  IP  TCP
#                   addr          addr port  port          len num frag flags
# Feb 1 20:31:21
#                   Dest          Src Dest  Src Proto TOS  Pkt  Intf  IP  TCP
#                   addr          addr port  port          len num frag flags
```

Trace Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the `/var/log` directory. By default, this file is named `/var/log/sampled`. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the `traceoptions` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
traceoptions {
  file filename {
    files number;
    size bytes;
    (world-readable | no-world-readable);
  }
}
```

Configure Flow Aggregation (cflowd)

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs the `cflowd` application available from CAIDA (<http://www.caida.org>). By using `cflowd`, you can obtain various types of byte and packet counts of flows through a router.

The `cflowd` application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process. To do this, include the `route-record` statement at the `[edit routing-options]` hierarchy level (for routing instances, include the statement at the `[edit routing-instances routing-instance-name routing-options]` hierarchy level):

```
[edit]
routing-options {
  route-record;
}
```

By default, flow aggregation is disabled. To enable the collection of flow aggregates, include the `cflowd` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```
[edit forwarding-options sampling output]
cflowd host-name {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
}
```

```

autonomous-system-type (origin | peer);
(local-dump | no-local-dump);
port port-number;
source-address address;
version format;
}

```

In the cflowd statement, specify the name, identifier, and source-address of the host that collects the flow aggregates. You must also include the UDP port number on the host and the version, which gives the format of the exported cflowd aggregates. To specify an IPv4 source address, include the source-address statement. To collect cflowd records in a log file before exporting, include the local-dump statement. To specify the cflowd version number, include the version statement. The cflowd version is either 5 or 8.



Note

You cannot specify both host (cflowd) sampling and port mirroring in the same configuration.

To specify aggregation of specific types of traffic, include the aggregation statement. This conserves memory and bandwidth enabling cflowd to export targeted flows rather than all the aggregated traffic.



Note

Aggregation is valid only if cflowd version 8 is specified.

To specify a flow type, include the aggregation statement at the [edit forwarding-options sampling output cflowd *host-name*] hierarchy level:

```

[edit forwarding-options sampling output cflowd host-name]
aggregation {
  source-destination-prefix;
}

```

You specify the aggregation type using one of the following options:

autonomous-system—Aggregate by autonomous system (AS) number; may require setting the separate cflowd autonomous-system-type statement to include either origin or peer AS numbers. The origin option specifies to use the origin AS of the packet source address in the Source Autonomous System cflowd field. The peer option specifies to use the peer AS through which the packet passed in the Source Autonomous System cflowd field. By default, cflowd exports the origin AS number.

destination-prefix—Aggregate by destination prefix (only).

protocol-port—Aggregate by protocol and port number; requires setting the separate cflowd port statement.

source-destination-prefix—Aggregate by source and destination prefix. Version 2.1b1 of CAIDA's cflowd application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the caida-compliant statement, the JUNOS software complies with Version 2.1b1 of cflowd. If you do not include the caida-compliant statement in the configuration, the JUNOS software records source and destination mask length values in compliance with the *cflowd Configuration Guide*.

source-prefix—Aggregate by source prefix (only).

Collection of sampled packets in a local ASCII file is not affected by the cflowd statement.

Debug cflowd Flow Aggregation

To collect the cflowd flows in a log file before they are exported, include the local-dump option at the [edit forwarding-options sampling output cflowd *hostname*] hierarchy level:

```
[edit forwarding-options sampling output cflowd host-name]
local-dump;
```

By default, the flows are collected in /var/log/sampled; to change the filename, include the filename statement at the [edit forwarding-options sampling traceoptions] hierarchy level. For more information about changing the filename, see "Configure Traffic Sampling Output" on page 231.



Because the local-dump option adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43   Src addr: 192.53.127.1
Jun 27 18:35:43   Dst addr: 192.6.255.15
Jun 27 18:35:43   Nhop addr: 192.6.255.240
Jun 27 18:35:43   Input interface: 5
Jun 27 18:35:43   Output interface: 3
Jun 27 18:35:43   Pkts in flow: 15
Jun 27 18:35:43   Bytes in flow: 600
Jun 27 18:35:43   Start time of flow: 7230
Jun 27 18:35:43   End time of flow: 7271
Jun 27 18:35:43   Src port: 26629
Jun 27 18:35:43   Dst port: 179
Jun 27 18:35:43   TCP flags: 0x10
Jun 27 18:35:43   IP proto num: 6
Jun 27 18:35:43   TOS: 0xc0
Jun 27 18:35:43   Src AS: 7018
Jun 27 18:35:43   Dst AS: 11111
Jun 27 18:35:43   Src netmask len: 16
Jun 27 18:35:43   Dst netmask len: 0
```

[... 41 more v5 flow entries; then the following header:]

```

Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43   Num-records: 42
Jun 27 18:35:43   Version: 5
Jun 27 18:35:43   Flow seq num: 118
Jun 27 18:35:43   Engine id: 0
Jun 27 18:35:43   Engine type: 3

```

Configure Port Mirroring

On routers containing a Internet Processor II ASIC, you can send a copy of an IPv4 packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

One application for port mirroring sends a duplicate packet to a virtual tunnel. Then, a next-hop group can be configured to forward copies of this duplicate packet to several interfaces. For more information about next-hop groups, see “Configure a Next-Hop Group” on page 224.

To configure port mirroring, include the port-mirroring statement at the [edit forwarding-options] hierarchy level:

```

[edit forwarding-options]
port-mirroring {
  input {
    family inet {
      rate num;
      run-length num;
    }
  }
  output {
    interface interface-name {
      next-hop address;
    }
    no-filter-check;
  }
}

```

To configure port mirroring, include the port-mirroring statement. To configure the address family, rate of sampling, and length of sampling for port mirroring, include the input statement. To specify which interface to send duplicate packets and the next-hop address to send packets, include the output statement. To see if there are any filters on the specified interface, include the no-filter-check statement.

In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, rather than another router. If you must send this traffic over a network, you should use tunnels. For more information about tunnel interfaces, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

The following restrictions apply:

You cannot configure both cflowd sampling and port mirroring in the same configuration.

You cannot configure firewall filters on the port-mirroring interface.

The interface you configure for port mirroring should not participate in any kind of routing activity.

The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of 190.68.9.10 and the port-mirrored traffic is sent to 190.68.20.15 for analysis, the device associated with the latter address should not know a route to 190.68.9.10. Also, it should not send the sampled packets back to the source address.

Only IPv4 traffic is supported.

Only transit data is supported.

You can configure only one port-mirroring interface per router. If you include more than one interface in the port-mirroring statement, the previous one is overwritten.

You must include a firewall filter with both the accept action and the sample action modifier on the inbound interface. Do not include the discard action, or port mirroring will not work.

