

Chapter 25

Configure Traffic Sampling

Traffic sampling requires a router equipped with an Internet Processor II ASIC. You can sample IP traffic based on particular input interfaces and various fields in the packet header. You can use traffic sampling to monitor any combination of specific logical interfaces, specific protocols on one or more interfaces, a range of addresses on a logical interface, or individual IP addresses. Information about the sampled packets is saved to files on the router's hard disk.

To determine whether a router has an Internet Processor II ASIC, use the `show chassis hardware` command.

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

```
[edit forwarding-options]
sampling {
  disable;
  input {
    family inet {
      rate number;
      run-length number;
    }
  }
  output {
    cflowd host-name {
      engine-id id-number;
      (local-dump | no-local-dump);
      port port-number;
      version format;
    }
    file {
      filename filename;
      files number;
      size bytes;
      (stamp | no-stamp);
      (world-readable | no-world-readable);
    }
  }
  traceoptions {
    file filename {
      files number;
      size bytes;
      (world-readable | no-world-readable);
    }
  }
}
```

This chapter describes the following tasks you perform in configuring traffic sampling:

Minimum Traffic Sampling Configuration on page 300

Configure Parameters for Sampling Traffic on page 301

Disable Traffic Sampling on page 301

Examples: Configure Traffic Sampling on page 301

Configure the Files to Contain Traffic Sampling Output on page 305

Configure Flow Aggregation on page 307

Minimum Traffic Sampling 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] hierarchy level. In the filter's then statement, you must specify the actions sample and accept.

```
[edit firewall]
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 inet {
      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 Parameters for Sampling Traffic

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

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

In the rate statement, specify the fraction of traffic you want to sample. By default, the rate is 0, which means that no traffic is sampled. The sampling rate parameter is the denominator of the ratio ($1/n$) of packets sampled; for example, if you set a rate of 10, every tenth packet will be sampled.

In the run-length statement, set the number of samples following the initial trigger event. By default, the run length is 0, which means that no traffic is sampled after the trigger event. This feature allows you to sample adjacent packets to those already being sampled. For example, if you set a rate of 100 to trigger sampling on 1 out of every 100 samples and set the run length to be 3, the software samples the next two packets that are marked for sampling as well, giving an effective sample rate of 3 percent. *number* can be a value from 0 through 20.

If you do not configure the input statement parameters, sampling is disabled. If you configure the input parameters but not the output parameters, sampling is enabled and the sampled packets are collected in the file `/var/tmp/sampled.pkts`. The default number of files maintained is five, and the default size of each file is 1 MB. For more information about the output file formats, see “Configure the Files to Contain Traffic Sampling Output” on page 305.

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 302

Sample All Traffic from a Single IP Address on page 303

Sample All FTP Traffic on page 304

Sample a Single SONET Interface

The following example configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET interface.

Create the filter at the [edit firewall] hierarchy level.

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

Apply the filter to the SONET interface at the [edit interfaces] hierarchy level:

```
[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 at the [edit forwarding-options] hierarchy level:

```
[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 example configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 215.45.92.31.

Create the filter at the [edit firewall] hierarchy level:

```
[edit firewall]
filter one-ip {
  term get-ip {
    from {
      source-address 215.45.92.31;
    }
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the Gigabit Ethernet interface at the [edit interfaces] hierarchy level:

```
[edit interfaces]
ge-4/1/1 {
  unit 0 {
    family inet {
      filter {
        input one-ip;
      }
      address 215.45.92.254/32 {
        destination 215.45.92.7;
      }
    }
  }
}
```

Finally, gather statistics on all the candidate samples, in this case using a rate of 100 percent:

```
[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 example 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 port.

Create a filter at the [edit firewall] hierarchy level:

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

Apply the filter to the T3 Ethernet interface at the [edit interfaces] hierarchy level:

```
[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 a portion of the candidate samples:

```
[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 1;
    }
  }
  output {
    file {
      filename t3-ftp-traffic.txt ;
      files 50;
      size 1m;
    }
  }
}
```

Configure the Files to Contain Traffic Sampling Output

By default, the traffic sampling results are placed on the router's hard disk in the file `/var/tmp/sampled.pkts`, and up to five 1-MB files are created to contain the results. You can change the filename, but the file is always placed in `/var/tmp`.

To change information about the files that contain the traffic sampling results, include the `file` statement at the `[edit forwarding-options sampling output]` hierarchy level:

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

Traffic Sampling Output Files

Traffic sampling output is saved to an ASCII text file. By default, it is named `/var/tmp/sampled.pkts`. You can configure a different name by including the `filename` statement at the `[edit forwarding-options sampling output file]` hierarchy level.

To display logging information about traffic sampling, use the `show firewall log` command. For an example, see “Filter Action Statement” on page 270.

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
```

This output contains the following fields:

Time—Time at which the packet was received (displayed only if you configure timestamping).

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 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 Transmission Control Protocol 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
```

Configure the File to Contain Logging Information

Logging information is placed in a 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 change the configuration of the logging file, include the file statement at the [edit forwarding-options sampling traceoptions] hierarchy level:

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

Configure Flow Aggregation

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>).

The aggregate contains the following types of byte and packet counts:

- Per logical interface
- Source address to destination address
- Source port to destination port
- Per protocol
- Per type of service
- Per autonomous system (AS) number

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 {
  engine-id id-number;
  (local-dump | no-local-dump);
  port port-number;
  version format;
}
```

In the cflowd statement, specify the name or identifier of the host that collects the flow aggregates. Include the Routing Engine ID number, the UDP port number on the host, and the version, which gives the format of the exported cflowd aggregates. To collect cflowd records in a log file before exporting, include the local-dump option.

Collection of sampled packets in a local ASCII file is not affected by the cflowd statement. To configure flow aggregation when you are sampling traffic, you must include the route-record enable statement at the [edit routing-options] hierarchy level:

```
[edit routing-options]
route-record enable;
```

Debugging 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 this configuration, you modify the [edit forwarding-options sampling traceoptions] statement; see “Configure the Files to Contain Traffic Sampling Output” on page 305 for more information. Note that the local-dump option adds extra overhead, so you should use it only for debugging cflowd problems, not on production routers.

The following is an example of the flow information. The AS number exported is the origin AS number. Note that all the 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
```