

---

## In This Chapter

This chapter provides information to configure Cflowd.

Topics in this chapter include:

- [Cflowd Overview on page 560](#)
  - [Operation on page 561](#)
  - [Cflowd Filter Matching on page 565](#)
- [Cflowd Configuration Process Overview on page 566](#)
- [Configuration Notes on page 567](#)

## Cflowd Overview

Cflowd is a tool used to sample IPv4, IPv6 and MPLS traffic data flows through a router. Cflowd enables traffic sampling and analysis by ISPs and network engineers to support capacity planning, trends analysis, and characterization of workloads in a network service provider environment.

Cflowd is also useful for Web host tracking, accounting, network planning and analysis, network monitoring, developing user profiles, data warehousing and mining, as well as security-related investigations. Collected information can be viewed several ways such as in port, AS, or network matrices, and pure flow structures. The amount of data stored depends on the cflowd configurations.

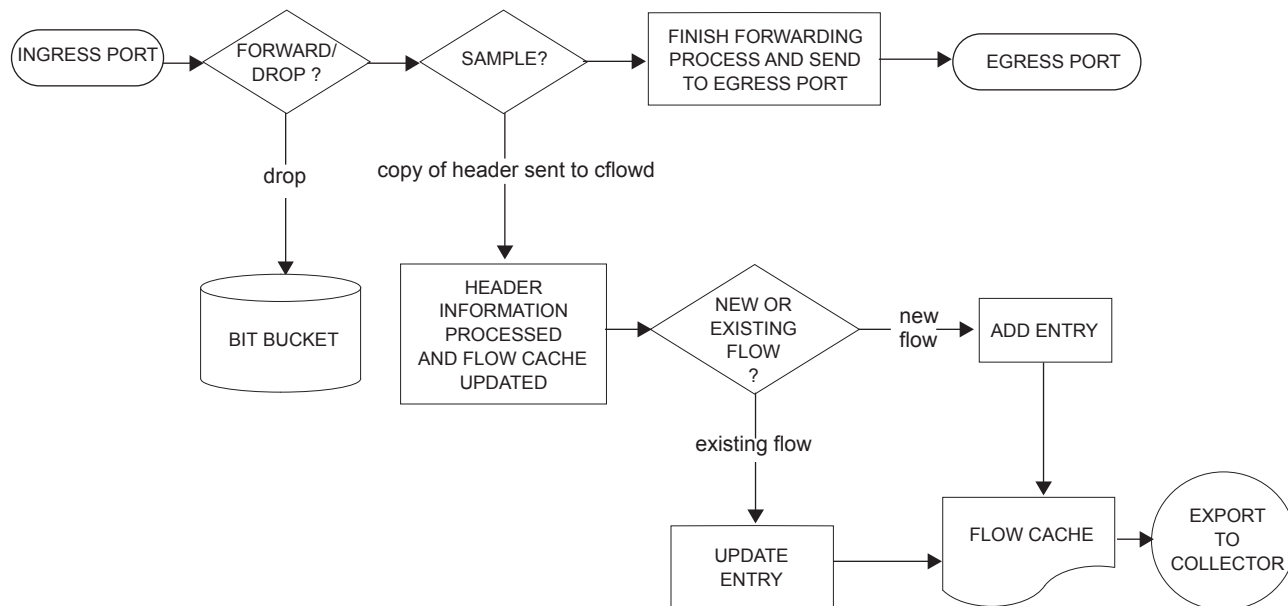
Cflowd maintains a list of data flows through a router. A flow is a uni-directional traffic stream defined by several characteristics such as source and destination IP addresses, source and destination ports, inbound interface, IP protocol and TOS bits.

When a router receives a packet for which it currently does not have a flow entry, a flow structure is initialized to maintain state information regarding that flow, such as the number of bytes exchanged, IP addresses, port numbers, AS numbers, etc. Each subsequent packet matching the same parameters of the flow contribute to the byte and packet count of the flow until the flow is terminated and exported to a collector for storage.

Cflowd is not supported on the 7750 SR-1 chassis.

## Operation

Figure 27 depicts the basic operation of the cflowd feature. This sample flow is only used to describe the basic steps that are performed. It is not intended to specify implementation.



**Figure 27: Basic Cflowd Steps**

1. As a packet ingresses a port, a decision is made to forward or drop the packet.
2. If the packet is forwarded, it is then decided if the packet should be sampled for cflowd.
3. If a new flow is found, a new entry is added to the cache. If the flow already exists in the cache, the flow statistics are updated.
4. If a new flow is detected and the maximum number of entries are already in the flow cache, the earliest expiry entry is removed. The earliest expiry entry/flow is the next flow that will expire due to the active or inactive timer expiration.
5. If a flow has been inactive for a period of time equal to or greater than the inactive timer (default 15 seconds), then the entry is removed from the flow cache.
6. If a flow has been active for a period of time equal to or greater than the active timer (default 30 minutes), then the entry is removed from the flow cache.

When a flow is exported from the cache, the collected data is sent to an external collector which maintains an accumulation of historical data flows that network operators can use to analyze traffic patterns.

Data is exported in one of the following formats:

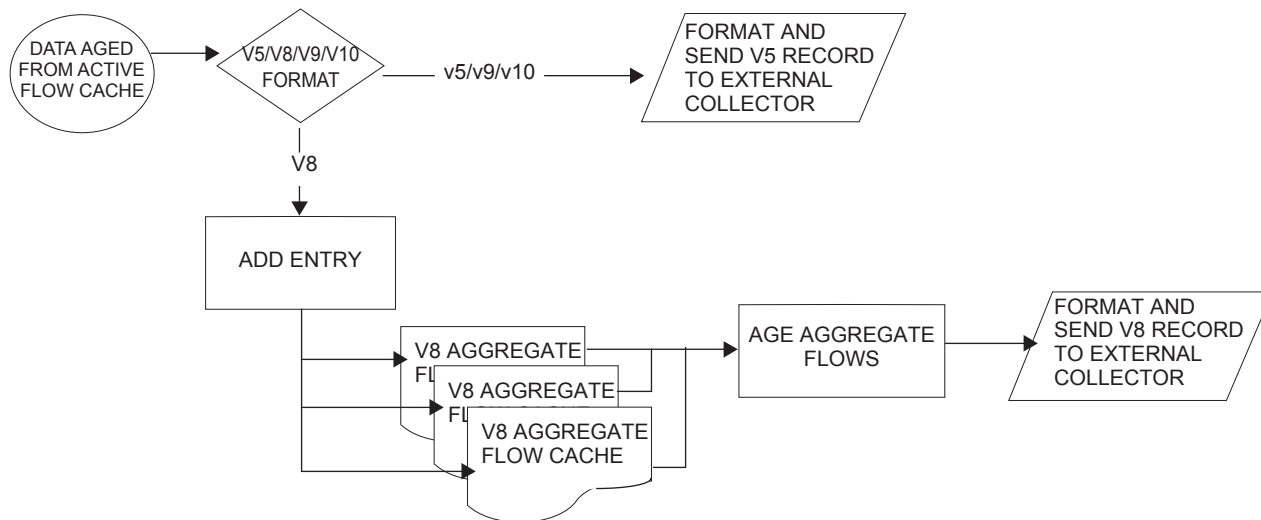
- Version 5 — Generates a fixed export record for each individual flow captured.
- Version 8 — Aggregates multiple individual flows into a fixed aggregate record.
- Version 9 — Generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.
- Version 10 (IPFIX) — Generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.

There are several different aggregate flow types including:

- AS matrix
- Destination prefix matrix
- Source prefix matrix
- Prefix matrix
- Protocol/port matrix.

V8 is an aggregated export format. As individual flows are aged out of the raw flow cache, the data is added to the aggregate flow cache for each configured aggregate type. Each of these aggregate flows are also aged in a manner similar to the method the active flow cache entries are aged. When an aggregate flow is aged out, it is sent to the external collector in the V8 record format.

[Figure 28](#) depicts Version 5, Version 8, Version 9, and Version 10 flow processing.



**Figure 28: V5, V8, V9, V10, and Flow Processing**

1. As flows are expired from the active flow cache, the export format must be determined, either Version 5, Version 8, Version 9, and Version 10.
2. If the export format is Version 5 or Version 9 and Version 10, no further processing is performed and the flow data is accumulated to be sent to the external collector.
3. If the export format is Version 8, then the flow entry is added to one or more of the configured aggregation matrices.

As the entries within the aggregate matrices are aged out, they are accumulated to be sent to the external flow collector in Version 8 format.

The sample rate and cache size are configurable values. The cache size default is 64K flow entries.

A flow terminates when one of the following conditions is met:

- When the inactive timeout period expires (default: 15 seconds). A flow is considered terminated when no packets are seen for the flow for N seconds.
- When an active timeout expires (default: 30 seconds). Default active timeout is 30 minutes. A flow terminates according to the time duration regardless of whether or not there are packets coming in for the flow.
- When the user executes a **clear cflowd** command.
- When other measures are met that apply to aggressively age flows as the cache becomes too full (such as *overflow percent*).

## Version 9

The Version 9 format is a more flexible format and allows for different templates or sets of cflowd data to be sent based on the type of traffic being sampled and the template set configured.

Version 9 is interoperable with RFC 3954, *Cisco Systems NetFlow Services Export Version 9*.

---

## Version 10

Version 10 is a new format and protocol that inter-operates with the specifications from the IETF as the IP Flow Information Export (IPFIX) standard. Like Version 9, the version 10 format uses templates to allow for different data elements regarding a flow that is to be exported and to handle different type of data flows such as IPv4, IPv6, and MPLS.

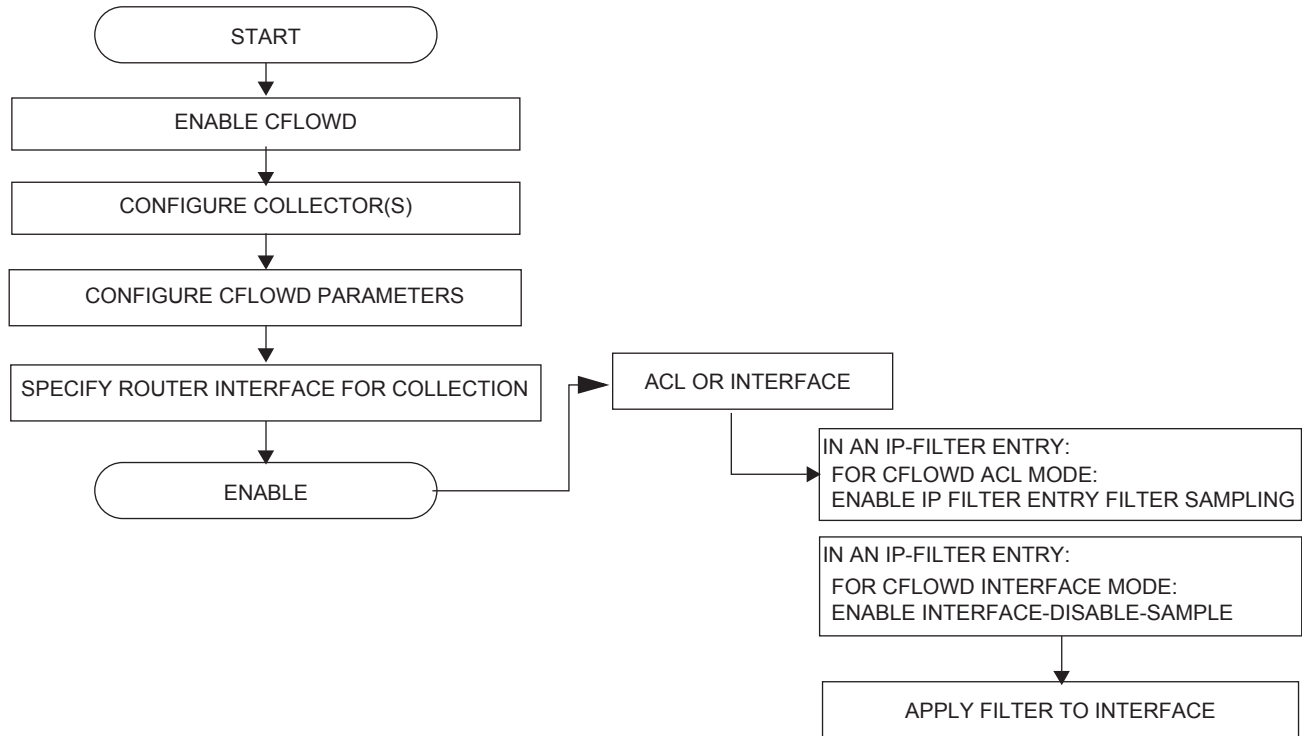
Version 10 is interoperable with RFC 5150 and 5102.

## Cflowd Filter Matching

In the filter-matching process, normally, every packet is matched against filter (access list) criteria to determine acceptability. With cflowd, only the first packet of a flow is checked. If the first packet is forwarded, an entry is added to the cflowd cache. Subsequent packets in the same flow are then forwarded without needing to be matched against the complete set of filters. Specific performance varies depending on the number and complexity of the filters.

# Cflowd Configuration Process Overview

Figure 29 displays the process to configure Cflowd parameters.



**Figure 29: Cflowd Configuration and Implementation Flow**

There are three modes in which cflowd can be enabled to sample traffic on a given interface:

- Cflowd interface, where all traffic entering a given port will be subjected to sampling as the configured sampling rate
- Cflowd interface plus the definition of IP filters which specify an action of interface-disable-sample, in which traffic that matches these filter entries will not be subject to cflowd sampling.
- Cflowd ACL, where IP filters must be created with entries containing the action filter-sampled. In this mode only traffic matching these filter entries will be subject to the cflowd sampling process.



## Configuration Notes

The following cflowd components must be configured for cflowd to be operational:

- Cflowd is enabled globally.
- At least one collector must be configured and enabled.
- A cflowd option must be specified and enabled on a router interface.
- Sampling must be enabled on either:
  - An IP filter which is applied to a port or service.
  - An interface on a port or service.

