



SRv6 Traffic Accounting

Effective management of Segment Routing over IPv6 (SRv6) networks necessitates granular visibility into traffic flow and resource consumption. This chapter addresses the critical requirement for traffic accounting within an SRv6 domain, leveraging the inherent programmability and explicit path control offered by the architecture. It details the methodologies, mechanisms, and tools employed to systematically monitor and quantify data traversing specific segments, policies, and Segment Identifiers (SIDs). The objective is to provide actionable data essential for optimizing network resource utilization, validating Service Level Agreements (SLAs), facilitating informed capacity planning, and enabling precise troubleshooting within complex SRv6 deployments.

- [SRv6 traffic accounting, on page 1](#)
- [Usage guidelines for SRv6 traffic accounting, on page 3](#)
- [How SRv6 traffic accounting works, on page 4](#)
- [Configure SRv6 traffic accounting, on page 7](#)

SRv6 traffic accounting

SRv6 traffic accounting is a network management solution that:

- enables routers to record the number of packets and bytes transmitted on specific egress interfaces for IPv6 traffic using per-locator SRv6 counters
- monitors aggregated traffic flows as they enter, traverse, and leave the SRv6-enabled network, enabling traffic analysis from ingress to egress nodes, and
- provides essential data for capacity planning, traffic forecasting, SLA compliance, and anticipating network failures or expansion needs.

Per-locator SRv6 counters are measurement tools within SRv6 traffic accounting that track packets and bytes for each SRv6 locator, providing granular insight into traffic volumes for detailed monitoring and capacity planning.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
SRv6 Traffic Accounting	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100]) This feature is now supported on: • 8712-MOD-M • 8011-4G24Y4H-I
SRv6 Traffic Accounting	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM
SRv6 Traffic Accounting	Release 7.10.1	You can now enable the router to record the number of packets and bytes transmitted on a specific egress interface for IPv6 traffic using the SRv6 locator counter. You can use this data to create deterministic data tools to anticipate and plan for future capacity planning solutions. This feature introduces or modifies the following changes: CLI: • accounting prefixes ipv6 mode per-prefix per-nexthop srv6-locators YANG Data Models: • <code>Cisco-IOS-XR-accounting-cfg</code> • <code>Cisco-IOS-XR-fib-common-oper.yang</code> (see GitHub, YANG Data Models Navigator)

Importance of SRv6 traffic accounting

Modern networks handle increasing traffic volumes and require precise, real-time insights into how traffic moves through the SRv6 domain. Without accurate traffic accounting, operators risk congestion, inefficient resource allocation, difficulty meeting SLAs, and challenges in planning for growth or failures.

Benefits of SRv6 traffic accounting

Monitoring the traffic provides numerous benefits, and here are a few:

- To optimize network utilization and achieve a balance between underutilized and overutilized paths.
- To plan and optimize network capacity and avoid congestion.
- To plan the service provisioning and choose the right path and create an optimized backup path (for using SRLG's affinity, and so on).

Usage guidelines for SRv6 traffic accounting

Supported traffic types

- IPv6 packets.
- SRv6 VPNv4
- SRv6 VPNv6
- SRv6 INETv4
- SRv6 INETv6

SRv6 packet accounting behavior

- SRv6 packets with the local SID as the top SID.
 - If the top SID is a local uN, traffic is counted against the remote locator prefix of the next SID.
 - Traffic is not counted if the top SID is a local uA.

Measurement and counting details

- Ethernet header is considered for bytes accounting.
- SRv6 traffic accounting does not count locally generated control plane packets such as ping to the remote locator.
- Packets aren't counted if the local uA is the top SID.

Configuration requirements

- Supports a minimum telemetry pull interval of 30 seconds.

Configuration requirements

- SRv6 traffic accounting is not supported with SRv6 TE policy.
- No additional MIBs are supported to retrieve SRv6 traffic statistics. We recommend to use telemetry through the newly added sensor-path in `Cisco-IOS-XR-fib-common-oper` to retrieve these statistics.

How SRv6 traffic accounting works

Summary


The key components that are involved in the SRv6 locator counters and traffic accounting process are:

- CE1 (Customer Edge 1): An endpoint that originates external traffic in the network topology.
- CE2 (Customer Edge 2): An endpoint that receives both external and internal traffic.
- PE1 (Provider Edge 1): A network device that learns reachability, imposes traffic with SRv6 locators, and performs external traffic accounting.
- PE4 (Provider Edge 4): A network device that originates internal traffic destined for other parts of the network, typically passing through provider edge devices.
- P21 and P22: Specific interfaces or paths on PE1 that are used for forwarding traffic, contributing to the ECMP (Equal-Cost Multi-Path) setup.
- SRv6 Locator Counters (LOC.INT.E): The mechanism used for accounting traffic on a per-prefix and per-egress interface basis within the SRv6 domain.
- Demand Matrix (DM): The Demand Matrix (DM) is a representation of the amount of data transmitted between every pair of routers. Each cell in the DM represents a traffic volume from one router to another. DM gives a complete view of the traffic in your network.

The SRv6 locator counters and traffic accounting process involves tracking traffic flows and quantifying data transmission. The key components work together to track external and internal traffic flows and calculate the net external traffic, providing a comprehensive view of network traffic distribution. The traffic sent and received from CE1 is considered as the external traffic. The traffic from PE4 destined to PE2 is considered as the internal traffic.

Workflow

Figure 1: Sample Topology for SRv6 Traffic Accounting



These stages describe how SRv6 locator counters and traffic accounting work.

1. PE1 learns CE2 reachability through PE2. Consider PE1 has ECMP paths via P21 and P22 to reach PE2. When traffic reaches PE1, PE1 imposes traffic with the PE2 locator `fcbb:bb00:2::`.

2. When traffic exits the PE1 interface (fa21) through P21, PE1 keeps the count of this traffic that is sent. Also, when traffic exits the PE1 interface (fa22) through P22, PE1 keeps the count of this traffic that is sent. The traffic is accounted irrespective of the path PE1 takes to send traffic.

Here is the SRv6 label of the outgoing traffic for PE2:


```
<fcbb:bb00:1::1, fcbb:bb00:2:f2d1::> [CUSTTraffic]
```

3. When the next set of packets are received and passed through PE1, the counters are incremented on fa21 or fa22 interface based on the path the traffic sent through PE2.
4. When traffic is sent from PE4 to PE2 through PE1, PE4 imposes the traffic with the PE2 locator ID `fcbb:bb00:2::`. The traffic count is recorded at PE4 for this locator ID. The traffic from PE4 to PE1 is considered as internal traffic.
5. When traffic reaches PE1, it looks for the PE2 locator ID and keeps the traffic count at PE1 when the traffic exit the fa21 interface.
6. SRv6 traffic is calculated using the demand matrix.

Let's understand the calculation with an example. In the topology, the amount of external traffic destined for PE2 is a combination of external and internal traffic.



Figure 2: SampleTopology for Demand Matrix



The formula for external traffic to PE2 is applied:

```

External traffic to PE2
= (sum of all Loc.int.E counters on PE1) - ( sum of the Loc.int.E counters of all
neighbors of PE1)
= 39 gigabits per second - (21 + 0 + 0) gigabits per second
= 18 gigabits per second external traffic

```

Table 2: Demand Matrix showing traffic transmitted from PE1 and PE4 to PE2

From/To	PE1	PE2
PE1	NA	$39 - (21 + 0 + 0) = 18$ gigabits per second PE2 receives 18 gigabits per second external traffic from PE1
PE4	$21 - (18 + 0 + 0) = 3$ gigabits per second	$39 - (18 + 0 + 0) = 21$ gigabits per second

Configure SRv6 traffic accounting

Before you begin

Ensure that you enable SRv6 and its services.

Follow these steps to configure SRv6 traffic accounting on the router.

Procedure

Step 1 Enable SRv6 traffic accounting

Example:

```
Router#configure
Router(config)#accounting
Router(config-acct)#prefixes ipv6 mode per-prefix per-nextthop srv6-locators
Router(config-acct)#commit
```

Step 2 Verify the Stats ID allocated for remote locator. The following example shows the SRv6 locator ID and the stats ID allocated for the prefixes with the locator ID.

Example:

```
Router#show route ipv6 fccc:cc00:1:: detail

Routing entry for fccc:cc00:1::/48
  Known via "isis 100", distance 115, metric 101, SRv6-locator, type level-1 <===== locator
  flag
    Installed Jun  1 11:59:10.941 for 00:00:04
  Routing Descriptor Blocks
    fe80::1, from 1::1, via Bundle-Ether1201, Protected, ECMP-Backup (Local-LFA)
      Route metric is 101
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:2          Path ref count:1
      NHID: 0x2001b (Ref: 79)
      Stats-NHID: 0x2001c (Ref: 6)
      Backup path id:1
    fe80::1, from 1::1, via TenGigE0/1/0/5/2, Protected, ECMP-Backup (Local-LFA)
      Route metric is 101
      Label: None
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0
      Path id:1          Path ref count:1
      NHID: 0x2001a (Ref: 79)
      Stats-NHID: 0x2001d (Ref: 6) <===== Stats-NHID is allocated for prefixes with locator
  flag
    Backup path id:2
    Route version is 0x68 (104)
    No local label
    IP Precedence: Not Set
    QoS Group ID: Not Set
    Flow-tag: Not Set
    Fwd-class: Not Set
```

```

Route Priority: RIB_PRIORITY_NON_RECURSIVE_LOW (8) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 2, Download Version 39779
No advertising protos.

```

Step 3 Configure the sensory path to retrieve the accounting data using telemetry:

Example:

```

Router#configure
Router(config)#grpc
Router(config-grpc)#port 57400
Router(config-grpc)#no-tls
Router(config-grpc)#commit
Router(config-grpc)#exit
Router(config)#telemetry model-driven
Router(config-model-driven)#sensor-group s1
Router(config-model-driven-snsr-grp)#sensor-path
Cisco-IOS-XR-fib-common-oper:cef-accounting/vrfs/vrf[vrf-name='default']/af$
Router(config-model-driven-snsr-grp)#exit
Router(config-model-driven)#subscription sub1
Router(config-model-driven-subs)#sensor-group-id s1 sample-interval 30000
Router(config-model-driven-subs)#commit
Router(config-model-driven-subs)#root
Router(config)#exit
Router#

```

Step 4 Verify the counters using the telemetry data. The following example shows the accounting data with the number of packets and the bytes transmitted through the interface.

Example:

```

{
  "Cisco-IOS-XR-fib-common-oper:cef-accounting": {
    "vrfs": {
      "vrf": [
        {
          "vrf-name": "default",
          "afis": {
            "afi": [
              {
                "afi-type": "ipv6",
                "pfx": {
                  "srv6locs": {
                    "srv6loc": [
                      {
                        "ipv6-address": " fccc:cc00:1::",
                        "prefix-length": 48,
                        "ipv6-prefix": " fccc:cc00:1::",
                        "ipv6-prefix-length": 48,
                        "accounting-information": [
                          {
                            "number-of-tx-packets": "1500000",
                            "number-of-tx-bytes": "378000000",
                            "path-index": 0,
                            "outgoing-interface": "Bundle-Ether1201",
                            "nexthop-addr": "fe80::2/128"
                          },
                          {
                            "number-of-tx-packets": "1000000",
                            "number-of-tx-bytes": "252000000",
                            "path-index": 1,
                            "outgoing-interface": "TenGigE0/0/0/22",
                            "nexthop-addr": "fe80::2/128"
                          }
                        ]
                      }
                    ]
                  }
                }
              }
            ]
          }
        }
      ]
    }
  }
}

```

<===== Accounting data

<===== Accounting data

<===== Accounting data

<===== Accounting data

Step 5 Run **showcef ipv6 accounting** command to display the packets per bytes:

```
Router#sh cef ipv6 accounting
fccc:cc00:33::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
  via fe80::2/128, Bundle-Ether1201
    path-idx 0
    next hop fe80::2/128
    Accounting: 0/0 packets/bytes output
fccc:cc05:2::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
  via fe80::2/128, Bundle-Ether1201
    path-idx 0
    next hop fe80::2/128
    Accounting: 0/0 packets/bytes output
fccc:cc3e:2::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
  via fe80::2/128, Bundle-Ether1201
    path-idx 0
    next hop fe80::2/128
    Accounting: 0/0 packets/bytes output
fccc:cc3e:3::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
  via fe80::2/128, Bundle-Ether1201
    path-idx 0
    next hop fe80::2/128
    Accounting: 200000/58400000 packets/bytes output <<< for prefix fccc:cc3e:3:: we can see 2lac
    packets count
```

