

---

## In This Chapter

This chapter provides information to configure sFlow.

Topics in this chapter include:

- [sFlow Overview on page 530](#)
- [sFlow Features on page 531](#)
  - [sFlow Counter Polling Architecture on page 531](#)
  - [sFlow Support on Logical Ethernet Ports on page 532](#)
  - [sFlow SAP Counter Map on page 532](#)
- [sFlow Record Formats on page 533](#)

## sFlow Overview

Some Layer 2 network deployments collect statistics on physical Ethernet ports and on Layer 2 interfaces at a high-frequency using a push model to, among others, monitor traffic, diagnose network issues, and/or provide billing. SR OS supports cflowd and XML accounting; however, those mechanisms are either Layer-3 specific, or focus on providing statistics at extremely large scale (thus use a pull model and cannot support high-frequency counter updates). To meet the statistics collection requirements of such Layer 2 deployments, SROS supports sFlow statistics export using sFlow version 5.

The following list main caveats for sFlow support:

- sFlow is supported on 7950 XRS product family and on 7750 SR12E only
- sFlow data sources require multi-core line cards, enabling sFlow on a card that is not a multi-core is not blocked and can be detected by SNMP trap/log generated by sFlow
- To meet high-frequency export of counters, sFlow implementation is targeted for low per-port VLL/VPLS SAP scale only. The configuration is blocked if the per-port VLL/VPLS SAP limit exceeds sFlow limit. Contact your Alcatel-Lucent representative for per-platform scaling limits applicable.

## sFlow Features

This section describes sFlow functionality supported in SR OS.

---

### sFlow Counter Polling Architecture

When sFlow is enabled on an SROS router, the system takes upon a role of an sFlow network device as described in sFlow protocol version 5. A single sFlow agent can be configured for counter polling (flow sampling is not supported). There is no support for sub-agents.

The sFlow agent sends sFlow data to an operator-configured sFlow receiver. A single receiver is supported with configurable primary and backup IPv4 or IPv6 UDP destination sockets for redundancy (each sFlow packet exported is duplicated to both sockets when both are configured). The receiver's UDP sockets can be reachable either in-band or out-of-band (default) and must both be IPv4 or IPv6. An operator can also set the maximum size of the sFlow datagrams. Operators are expected to set this value to avoid IP fragmentation (Datagrams exceeding the specified size are fragmented before handed to IP layer).

The sFlow agent manages all sFlow data sources in the system. SROS supports sFlow data that are physical ports. When a port is configured as an sFlow data source, counters for that port and all VPLS and ePipe SAPs on that port are collected and exported using sFlow (see later on section for record format). Flow data sources can only be configured when an sFlow receiver is configured. To remove the sFlow receiver, all sFlow data sources must first be de-configured at the port level.

Each data source is processed at a 15-second, non-configurable interval. If multiple data sources exist on a line card, the line card distributes the processing of each data source within a 15 second interval to avoid sFlow storms. When a timer expires to trigger a data source processing, data is collected for the physical port and for all VLL and VPLS SAPs on that port and exported using sFlow version 5 records as described in later subsections of this document. Each port and all SAP records for a given data source for a given interval are collected and sent with the counter sequence number and the timestamp value (the time value corresponds to the time counters were actually collected by a line card). The timestamp value uses line card's sysUptime value, which is synchronized with CPM time automatically by the system. A line card sends the counters to a CPM card, where sFlow UDP datagrams are created, sequenced with the CPM sequence number and sent to the receiver. If no UDP sockets are configured, no errors are generated because data is not sent. If no UDP sockets are reachable, the created UDP sFlow datagrams are dropped.

Note that line cards will reset the counter record sequence numbers if, as a result of configuration or operational change, the return statistics no longer provide continuity with the previous interval. The following lists examples of when this takes place:

- The card hard or soft resets
- The MDA resets

- The sFlow agent counter map changes

Note that CPM will reset the sFlow datagram sequence numbers if, as a result of configuration or operational change, the sFlow datagram to be sent no longer provides continuity with the previous datagram. The following lists examples of when this takes place:

- HA switch
- CTL reboot
- Creation of an sFlow receiver

---

## sFlow Support on Logical Ethernet Ports

sFlow data sources operate in a context of physical Ethernet port. To enable sFlow on Ethernet logical ports and their SAPs, an operator must explicitly enable sFlow on every physical Ethernet port that is a member of the given logical port. Currently only LAG logical ports are supported (including MC-LAG). Note that sFlow configuration does not change automatically when a port is added or removed to or from a LAG.

For SAPs on a LAG, egress statistics will increment based on ports used by each SAP on LAG egress while ingress statistics will increment based on ports used by each SAP on LAG ingress unless LAG features like, for example, per-fp-ingress-queuing or per-fp-sap-optimization result in SAP statistics collection against a single LAG port.

If logical-level view is required, for example, per LAG statistics, a receiver is expected to perform data correlation based on per-physical port interface and SAP records exported for the given logical port's physical ports and their SAPs. sFlow data records contain information that allows physical ports/SAP records correlation to a logical port. See [sFlow Record Formats](#).

**Note:** Correlation of records must allow for small difference in timestamp values returned for member ports or SAP on a LAG because all ports run independent timestamps.

---

## sFlow SAP Counter Map

To allow per SAP sFlow statistics export, operators must configure ingress and egress sFlow counter maps. The counter maps are required, because SROS systems support more granular per policer/queue counters and not IF-MIB counters per VLL/VPLS SAPs. In an absence of a map configured, 0's will be returned in corresponding statistics records.

A single ingress and a single egress counter map are supported. The maps specify which ingress and which egress SAP QoS policy queue/policer statistics map to sFlow unicast, multicast, and broadcast counters returned in an sFlow SAP record. Multiple queues and/or policers can map to each of unicast, multicast, broadcast counters. A single queue/policer can only map to one type of

traffic. Queues, policers configured in a SAP QoS policy but not configured in an sFlow map or vice-versa are ignored when sFlow statistics are collected.

## sFlow Record Formats

Table 54 describes sFlow record used and exported:

**Table 54: sFlow Record Fields**

Record	Field	Value
sFlow Datagram Header (SAP and port)	Datagram version	5
	Agent Address	Active CPM IPv4 address (from BoF)
	Sub-agent ID	0
	Sequence number	CPM inserted sFlow datagram sequence number
	SysUptime	sysUptime when the counters for records included in the datagram were collected by the line card
	NumSamples	Number of counter records in the datagram
Counter header (SAP and Port)	Enterprise	0 (standard sFlow)
	sFlow Sample Type	4 (Expanded counter sample)
	Sample Length	sFlow packet size excluding header
	Sequence number	Line card-inserted sequence number
	Source ID Type	0
	Source ID Index	tmnxPortId of the physical port (sFlow data source)
	Counter records	Count of counter records in the datagram

**Table 54: sFlow Record Fields (Continued)**

Record	Field	Value
Ethernet Interface Counters (EIC) – port (Ethernet Layer)	Enterprise	Statistics returned are based on dot3StatsEntry in EtherLike-MIB.mib. Statistics support may depend on hardware type.
	Format	
	Flow data length	
	Alignment Errors	
	FCS Errors	
	Single Collision Frames	
	Multiple Collision Frames	
	SQE Test Errors	
	Deferred Transmissions	
	Late Collisions	
	Excessive Collisions	
	Internal Mac Transmit Errors	
	Carrier Sense Errors	
	Frame Too Longs	
Internal Mac Receive Errors		
Symbol Errors		

**Table 54: sFlow Record Fields (Continued)**

Record	Field	Value
Generic Interface Counters (GIC) – port/SAP	Enterprise	0 (standard sFlow)
	Format	1 (GIC)
	Flow data length	88
	ifIndex	Port: ifIndex (tmnxPortId) of phys port SAP: SapEncapValue - part of SAP SNMP key
	ifType	Port: 6 (EthernetCsmacd) SAP: 1 (Other)
	ifSpeed	Port: Port speed value SAP: <ul style="list-style-type: none"> <li>top 32 bits: svcId for SAP (TIMETRA-SAP.mib)</li> <li>lower 32 bits: sapPortId (TIMETRA-SAP.mib)</li> </ul> <b>Notes:</b> The values plus ifIndex in the record are SAP SNMP key. <b>Notes:</b> SapPortId is LAG's tmnxPortId for SAPs on a LAG and port's tmnxPortId for SAPs on physical port
	ifDirection	Derived from MAU MIB (0 = unknown, 1 = full duplex, 2 = half duplex, 3 = in, 4 = out)
	ifAdminStatus	0 (down) 1 (up)
	ifOperStatus	0 (down) 1 (up)
	Input Octets	Statistics return for port are based on ifEntry or ifXEntry in IF-MIB.mib as applicable.  Statistics returned for SAPs are sum of counters based on the sFlow ingress/egress counter map configured.
	Input Packets	
	Input Multicast packets	
	Input Broadcast packets	
Input Discarded packets		

**Table 54: sFlow Record Fields (Continued)**

Record	Field	Value
Generic Interface Counters (GIC) – port/SAP (Continued)	Input Errors	Statistics return for port are based on ifEntry or ifXEntry in IF-MIB.mib as applicable. Statistics returned for SAPs are sum of counters based on the sFlow ingress/egress counter map configured.
	Input Unknown Protocol Packets	
	Output Octets	
	Output Packets	
	Output Multicast packets	
	Output Broadcast packets	
	Output Discarded packets	
	Output Errors	
	Promiscuous Mode	0 (FALSE)

**Notes:**

- 0 is returned for statistics that are not supported by a given hardware type.
- If required, CPM executes rollover logic to convert internal 64-bit counters to a 32-bit sFlowd counter returned.