



Nokia Service Router Linux
7215 Interconnect System
7220 Interconnect Router
7250 Interconnect Router
7730 Service Interconnect Router
Release 26.3

Product Overview

3HE 22243 AAAA TQZZA
Edition: 01
March 2026

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

© 2026 Nokia.

Table of contents

1	About this guide.....	7
1.1	Precautionary and information messages.....	7
1.2	Conventions.....	7
2	What's new.....	9
3	About SR Linux.....	10
3.1	What is SR Linux?.....	10
3.2	Features overview.....	11
3.2.1	Modular network applications.....	11
3.2.2	Model-driven architecture.....	11
3.2.2.1	IDB publish/subscribe model for messaging.....	11
3.2.3	Data models support.....	11
3.2.4	Protocol buffers and gRPC for inter-process communication.....	12
3.2.5	Third-party application support.....	13
3.2.6	CLI plug-ins.....	13
3.2.7	Hardware extensibility.....	13
3.2.8	Software extensibility.....	14
3.3	SR Linux NDK.....	14
3.4	SR Linux documentation.....	14
4	Hardware overview.....	18
4.1	7215 IXS platforms.....	20
4.1.1	Architecture.....	20
4.1.2	Chassis components.....	20
4.1.3	Power and cooling.....	20
4.2	7220 IXR-Dx platforms.....	21
4.2.1	Architecture.....	21
4.2.2	Chassis components.....	21
4.2.3	Power and cooling.....	21
4.3	7220 IXR-Hx platforms.....	22
4.3.1	Architecture.....	22
4.3.2	Chassis components.....	22
4.3.3	Power and cooling.....	22

4.4	7250 IXR platforms.....	23
4.4.1	Architecture.....	23
4.4.2	Chassis components.....	24
4.4.3	Power and cooling.....	24
4.5	7250 IXR-X platforms.....	24
4.5.1	Architecture.....	25
4.5.2	Chassis components.....	25
4.5.3	Power and cooling.....	25
4.6	7730 SXR-1 platforms.....	25
4.6.1	Architecture.....	26
4.6.2	Chassis components.....	26
4.6.3	Power and cooling.....	26
5	SR Linux architecture overview.....	27
5.1	SR Linux components.....	27
5.1.1	Linux kernel.....	27
5.1.2	Operating system.....	27
5.1.3	Modular applications.....	27
5.1.3.1	Application manager.....	29
6	SR Linux management overview.....	31
6.1	SR Linux management server.....	31
6.2	SR Linux CLI.....	31
6.3	JSON-RPC server.....	32
6.4	gRPC server.....	32
6.5	gNOI.....	32
6.6	NETCONF.....	32
6.7	SNMP.....	33
6.8	Zero Touch Provisioning.....	33
6.9	Secure Boot.....	34
6.10	Measured Boot.....	34
6.11	BootZ.....	34
6.12	SR Linux configuration.....	34
6.13	Securing access.....	35
6.14	SR Linux logging.....	35
6.15	P4Runtime support.....	36

6.16	Event Handler.....	36
6.17	gRIBI.....	36
6.18	Network Synchronization.....	37
7	SR Linux interfaces.....	38
7.1	SR Linux interface types.....	38
7.2	LAG interfaces.....	39
7.3	Subinterfaces.....	40
7.4	DHCP relay.....	40
7.5	LLDP.....	40
7.6	MACsec.....	40
8	SR Linux routing functions.....	41
8.1	Network instances.....	41
8.2	Static routes.....	41
8.3	Aggregate routes.....	42
8.4	BGP feature support.....	43
8.5	IS-IS feature support.....	44
8.6	OSPF feature support.....	45
8.7	Routing policies.....	46
8.8	ECMP load balancing.....	46
8.9	Access Control Lists.....	47
8.10	Quality of Service.....	47
8.11	MPLS feature support.....	47
8.12	Segment routing with IPv6 data plane (SRv6).....	49
8.13	Multicast.....	49
9	SR Linux services.....	51
9.1	Layer 2 services.....	51
9.2	EVPN for Layer 2.....	52
9.3	EVPN for Layer 3.....	52
9.4	IP-VPN services.....	52
10	SR Linux OAM and diagnostics tools.....	54
10.1	BFD support.....	54
10.2	sFlow support.....	54
10.3	Interactive traffic monitoring tool.....	55

10.4	Packet-trace tool.....	55
10.5	Traffic mirroring to remote and local destinations.....	55
10.6	TWAMP.....	56
10.7	STAMP.....	56
10.8	Ethernet connectivity fault management.....	56
10.9	Link measurement.....	57
10.10	Delay and loss measurement.....	57
10.11	Uncolored SR-MPLS TE-Policy.....	57
10.12	Colored TE policy.....	58
10.13	Service Activation Testhead (SAT).....	58
10.14	IPFIX.....	59
11	Standards and protocol support.....	60
11.1	Bidirectional Forwarding Detection (BFD).....	60
11.2	Border Gateway Protocol (BGP).....	60
11.3	Bridging and management.....	61
11.4	Ethernet VPN (EVPN).....	61
11.5	gRPC Remote Procedure Calls (gRPC).....	61
11.6	Intermediate System to Intermediate System (IS-IS).....	62
11.7	Internet Protocol (IP) general.....	63
11.8	Internet Protocol (IP) version 4.....	63
11.9	Internet Protocol (IP) version 6.....	63
11.10	Label Distribution Protocol (LDP).....	64
11.11	Multiprotocol Label Switching (MPLS).....	64
11.12	Open Shortest Path First (OSPF).....	64
11.13	Path Computation Element Protocol (PCEP).....	65
11.14	Pseudowire (PW).....	65
11.15	Quality of Service (QoS).....	65
11.16	Segment Routing (SR).....	66
11.17	Simple Network Management Protocol (SNMP).....	66
11.18	Timing.....	67
11.19	Two-Way Active Measurement Protocol (TWAMP).....	67
11.20	Virtual Private LAN Service (VPLS).....	67
11.21	Yet Another Next Generation (YANG).....	68
11.22	Yet Another Next Generation (YANG) OpenConfig Modules.....	68

1 About this guide

This document provides a basic overview of the Nokia Service Router Linux (SR Linux). This document is intended for marketing personnel, network technicians, administrators, operators, service providers, and others who need a basic understanding of SR Linux.

**Note:**

This manual covers the current release and may also contain some content that will be released in later maintenance loads. See the *SR Linux Software Release Notes* for information about features supported in each load.

Configuration and command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

1.1 Precautionary and information messages

The following are information symbols used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.2 Conventions

The Nokia SR Linux documentation uses the following command conventions:

- **Bold** type indicates a command that the user must enter.
- Input and output examples are displayed in `Courier` text.
- A vertical bar (|) indicates a mutually exclusive argument.
- Square brackets ([]) indicate optional elements.
- Braces ({ }) indicate a required choice. When braces are contained within square brackets, they indicate a required choice within an optional element.

- *Italic* type indicates a variable.

The following table outlines platform grouping conventions used in the SR Linux documentation suite.



Note: Some platforms in the 7250 IXR support mixed systems. For more information about mixed system support, see "Chassis types" in the *Configuration Basics Guide*.

Table 1: Platform grouping legend

Platform group	Description
7215 IXS	7215 IXS-A1
7220 IXR	All 7220 IXR platforms
7220 IXR-Dx	7220 IXR-D1, 7220 IXR-D2, 7220 IXR-D2L, 7220 IXR-D3, 7220 IXR-D3L, 7220 IXR-D4, 7220 IXR-D5
7220 IXR-Hx	7220 IXR-H2, 7220 IXR-H3, 7220 IXR-H4, 7220 IXR-H4-32D, 7220 IXR-H5-32D, 7220 IXR-H5-64D, 7220 IXR-H5-64O
7250 IXR ¹	7250 IXR platforms
7250 IXR Gen 2	7250 IXR-6, 7250 IXR-10
7250 IXR Gen 2c+	7250 IXR-6e with IMM2, 7250 IXR-10e with IMM2, 7250 IXR-X1b, 7250 IXR-X3b
7250 IXR Gen 3	7250 IXR-6e with IMM3, 7250 IXR-10e with IMM3, 7250 IXR-18e, 7250 IXR-X4
7250 IXR-6e/10e (mixed system)	7250 IXR-6e (mixed system) ² , 7250 IXR-10e (mixed system) ²
7730 SXR	7730 SXR-1-32D, 7730 SXR-1d-32D, 7730 SXR-1x-44S

¹ References to the 7250 IXR platform group may be appended with (including mixed systems) or (excluding mixed systems) to indicate mixed system support.

² References to this platform as part of 7250 IXR (mixed system) indicate mixed system support of 7250 IXR Gen 2c+ (IMM2) and 7250 IXR Gen 3 (IMM3). That is, the 7250 IXR-6e and 7250 IXR-10e can hold and support both IMM2 and IMM3 at the same time.

2 What's new

Table 2: What's new in 26.3.1

Topic	Location
Adds 7250 IXR-X4 and 7730 SXR-1-32D	7250 IXR-X platforms 7730 SXR-1 platforms
New topic - SRv6	Segment routing with IPv6 data plane (SRv6)
Updated the <i>Multicast</i> topic to indicate support for Multicast Source Discovery Protocol (MSDP)	Multicast
New topic - IPFIX	IPFIX
Updates standards and protocols	Standards and protocol support

3 About SR Linux

SR Linux delivers scalability, flexibility, and ease of operations for data centers and wide area networks. SR Linux also delivers:

- An open and extensible system that is fully programmable and scalable
- Model-driven management for simplified operations, integrations, and visibility
- Plug-and-play hardware integration
- Superior support for integrating community and customer-driven applications
- Customizable Command Line Interface (CLI) and on-demand CLI commands


3.1 What is SR Linux?

SR Linux is a Network Operating System (NOS) that leverages its routing protocol stack from Nokia SR OS, while also using Linux as its underlying operating system, allowing operators to use debugging and configuration tools that they are familiar with, even if they have no experience with SR OS.

Routing functions on SR Linux run as modular, lightweight applications, configurable via external APIs. These applications use gRPC and APIs to communicate with each other and external systems over TCP. The Nokia-supplied applications can be augmented by third-party developed applications, which plug into the SR Linux framework. Application-based functions allow for modular upgrades and easy fault isolation.

[Figure 1: SR Linux management framework](#) shows the SR Linux management framework.

Figure 1: SR Linux management framework



sw4389

3.2 Features overview

SR Linux supports a robust set of features. The sections that follow highlight major functionality.

3.2.1 Modular network applications

As a Linux-based NOS, SR Linux uses modular applications that are isolated in their own failure domains. A central application manager is responsible for the lifecycle of each application and provides full control of the protocols running on the system.

On SR Linux, each protocol (BGP, IS-IS, and so on) runs as its own application. These applications may be configured using external APIs, including CLI, gNMI, and JSON-RPC. The applications run like any others do in Linux. As well, users can integrate their own applications into SR Linux.

SR Linux uses an unmodified Linux kernel as its foundation to build a suite of network applications. This provides benefits such as reliability, portability, and ease of application development. Using an unmodified kernel also speeds the availability of non-Nokia applications (for example, OpenSSH) and security patches for operating system components.

3.2.2 Model-driven architecture

SR Linux makes extensive use of structured data models. Each application has a YANG model that defines its configuration and state. SR Linux exposes the YANG models to the supported management APIs. For example, the command `t tree` in the CLI is derived from the SR Linux YANG models loaded into the system, and a gNMI client can use `Set` RPCs to configure an application based on its YANG model. When a configuration is committed, the SR Linux management server validates the YANG models and translates them into protocol buffers for the impart database (IDB).

See the [SR Linux architecture overview](#) chapter for more information about the relationship between IDB and SR Linux components).

3.2.2.1 IDB publish/subscribe model for messaging

IDB is a lightweight database that controls messaging between SR Linux applications, using a publish/subscribe (pub/sub) model. To do this, the IDB database is split up into topics. Each application owns a set of topics, to which it publishes information, and can subscribe to topics published by other applications. Applications subscribe to topics when they open a session to the IDB, and publish messages to their own topics for other applications to consume.

3.2.3 Data models support

The SR Linux model-driven management interfaces are based on a common infrastructure that uses YANG models as the core definition for network element configuration, state, and operational actions. The model-driven interfaces take the underlying YANG modules and render them for the particular management interface.

SR Linux supports the following YANG data models:

- Nokia vendor-specific data models
- OpenConfig vendor-neutral data models

SR Linux YANG data models

Each application that SR Linux supports has a Nokia vendor-specific YANG model that defines the application's configuration and state. SR Linux exposes the YANG models to supported management APIs; for example, the CLI command trees derived from the SR Linux YANG models loaded into the system. A gNMI client can use Set RPCs to configure an application based on the YANG model. When you commit a configuration, the SR Linux management server validates the YANG models and translates them into protocol buffers for the impart database (IDB).

OpenConfig data models

OpenConfig is an informal working group that provides structured, vendor-neutral YANG data models to address the use requirements of networking applications and technologies by the community. The OpenConfig data models use standards-based, YANG data-modeling language, support Remote Procedure Calls (RPCs), and allow network operators to use a single set of data models to configure and manage commonly-used network protocols, services, and devices that support the OpenConfig initiative.

Data models interaction

The SR Linux vendor-specific and OpenConfig vendor-neutral data models can be used together to configure and manage network elements. SR Linux data models include vendor-specific features and functions that OpenConfig data models do not describe, and therefore offer a more complete representation of the capabilities of the SR Linux network elements. The SR Linux configuration and operational statements map to path statements in the supported OpenConfig modules. The mappings are exposed in JSON files delivered with the software package.


3.2.4 Protocol buffers and gRPC for inter-process communication

IDB stores data as protocol buffers (protobufs). Protobufs are a language-neutral, platform-neutral mechanism for serializing structured data. Each protocol buffer message is a small logical record of information, containing a series of name-value pairs.

SR Linux uses gRPC for inter-process communication. gRPC is a client application that can directly call methods on a server application on a different machine as if it was a local object. The supported external APIs (CLI, gNMI, and JSON-RPC) communicate with the SR Linux device and retrieve state information using gRPC.

SR Linux applications share state details with each other using the pub/sub model (see [Figure 2: SR Linux infrastructure](#)).

Figure 2: SR Linux infrastructure



3.2.5 Third-party application support

Third-party applications can be fully integrated into SR Linux with the same functionality as Nokia applications. This includes configuration using YANG, telemetry support, and life-cycle management. Because third-party applications are not managed independently, it allows a reduction in operational overhead.

3.2.6 CLI plug-ins

The SR Linux CLI is itself an application that can load dynamic plugins from other applications. You can develop custom **show** commands and run them from the SR Linux CLI. The CLI plugins allow for integration with remote systems, supporting retrieval of state information.

3.2.7 Hardware extensibility

SR Linux supports a variety of network chipsets through the Nokia eXtensible Data Path (XDP). XDP serves as a hardware abstraction layer that facilitates adoption of new or non-Nokia network chipsets. It provides a common set of software instructions that northbound applications use so that they are not directly dependent on ASIC vendor SDKs. XDP borrows from the development experience for high-performance VNFs and makes use of user space acceleration for traffic destined for the control plane and any non-ASIC interfaces.

3.2.8 Software extensibility


Every SR Linux application, including third-party applications, supports its own YANG model, which can be loaded into the system. Operators can see and define the syntax and semantics of their application in a simple and standardized form. With this design, the YANG data model is defined first, then the CLI, APIs, and **show** output formats are derived from it.

SR Linux handles management and operations using the gRPC Network Management Interface (gNMI). Because SR Linux is natively model-driven, it can stream telemetry without requiring any translation layers. Telemetry is supported using POLL, ON_CHANGE, and ONCE streaming.

Third-party applications have access to the full streaming telemetry framework. This allows these applications to operate and be monitored, configured, and debugged the same as any other application on the system (see [Figure 3: SR Linux software extensibility](#)).

In addition to the gNMI interface, SR Linux includes a CLI and a JSON-RPC API for management. The CLI provides a framework for accessing the underlying data models of the system. The JSON-RPC API interface supports requests against the data models, as well as allowing a programmable interface access to the extensible plugin framework in the CLI.

Figure 3: SR Linux software extensibility



3.3 SR Linux NDK

SR Linux provides the NetOps Development Kit (NDK), a software development kit with a suite of libraries to assist operators with developing alongside SR Linux applications. The NDK is provided in the form of header files written in C++. This allows the operator to add SR Linux functionality to their own applications, by using these header files and the methods they provide to interact directly with the SR Linux IDB server.

See the *SR Linux NDK API Reference Guide* for reference information about the gRPC APIs used with the NDK.

3.4 SR Linux documentation

The SR Linux documentation set consists of the documents listed in [Table 3: SR Linux documentation set](#). These documents are available in PDF and HTML formats.

Table 3: SR Linux documentation set

Document	Description
<i>SR Linux Product Overview Guide</i>	High-level description of SR Linux functionality, including key components, and where it fits into the network.
<i>7215 IXS-A1 Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-D Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-D2L and 7220 IXR-D3L Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-D4 Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-D5 Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-H2 and 7220 IXR-H3 Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-H4 Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-H4-32D Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-H5-32D Chassis Installation Guide (SR Linux)</i> <i>7220 IXR-H5-64D and 7220 IXR-H5-64O Chassis Installation Guide (SR Linux)</i> <i>7250 IXR-6 and 7250 IXR-10 Chassis Installation Guide (SR Linux)</i> <i>7250 IXR-6e and 7250 IXR-10e Chassis Installation Guide (SR Linux)</i> <i>7250 IXR-X Chassis Installation Guide (SR Linux)</i> <i>7730 SXR-1d-32D Chassis Installation Guide (SR Linux)</i> <i>7730 SXR-1x-44S Chassis Installation Guide (SR Linux)</i>	Site preparation, chassis and component installation procedures, and hardware component configuration procedures for hardware platforms that support SR Linux.
<i>SR Linux Software Installation Guide</i>	Basic concepts behind Linux kernel operation on SR Linux, and provides procedures for upgrading the software and provisioning SR Linux using Zero Touch Provisioning (ZTP).


Document	Description
<i>SR Linux Configuration Basics Guide</i>	Basic configuration concepts for SR Linux, including accessing and using the CLI, and how to manage the system. Descriptions and examples of how to configure key features are provided.
<i>SR Linux Routing Protocols Guide</i>	Supported protocols and examples of how to configure and implement them.
<i>SR Linux Interfaces Guide</i>	Supported interfaces and naming conventions. Provides examples to configure and implement supported interfaces.
<i>SR Linux ACL and Traffic Steering Guide</i>	Supported access control list types, actions for each platform type, match conditions, statistics collection, and TCAM allocation. Also describes how to configure traffic steering using policies and filters.
<i>SR Linux VPN Services Guide</i>	Basic configuration concepts for EVPN Layer 2 (L2), Layer 3 (L3), and IP-VPN services functionality. Provides examples to configure and implement various protocols and services.
<i>7220 IXR and 7250 IXR SR Linux Quality of Service Guide</i>	Basic configuration concepts for Quality of Service (QoS) functionality on IXR devices. Provides examples to configure QoS features and classifier policies.
<i>7730 SXR SR Linux Quality of Service Guide</i>	Basic configuration concepts for Quality of Service (QoS) functionality on SXR devices. Provides examples to configure QoS features and classifier policies.
<i>SR Linux Segment Routing Guide</i>	Basic configuration concepts for segment routing functionality. Provides examples to configure segment routing and use of tools that provide operational information about segment routing.
<i>SR Linux MPLS Guide</i>	Basic configuration concepts for the Multiprotocol Label Switching (MPLS) protocol. Provides examples to configure MPLS and LDP.
<i>SR Linux gRIBI Guide</i>	Basic configuration concepts for the gRPC Routing Information Base Interface (gRIBI) protocol. Provides configuration examples.

Document	Description
<i>SR Linux P4RT Guide</i>	Basic configuration concepts for the P4Runtime client. Provides configuration examples.
<i>SR Linux Event Handler Guide</i>	Basic configuration concepts for event handler Provides configuration examples.
<i>SR Linux Network Synchronization Guide</i>	Basic configuration for the implementation of Synchronous Ethernet and the Precision Time Protocol on SR Linux.
<i>SR Linux Data Model Reference Guide</i>	Descriptions of the configuration and state data models available for SR Linux.
<i>SR Linux System Management Guide</i>	Descriptions of the interfaces used with SR Linux, which include the CLI, gNMI, and JSON. This document also provides an overview to CLI plug-ins and describes Nokia-defined general, operation, and show commands.
<i>SR Linux Troubleshooting Toolkit Guide</i>	How to use and configure diagnostic tools for SR Linux, including interactive traffic monitoring and packet tracing.
<i>SR Linux CLI Plug-In Guide</i>	How to create custom show routines for SR Linux, as well as how to install, modify, and remove them.
<i>SR Linux Log Events Guide</i>	Contents of the log messages generated by SR Linux.
<i>SR Linux Advanced Configuration Guide</i>	Scenarios for configuring complex network-level configurations where additional guidance and more detailed procedures may be required.
<i>SR Linux OAM and Diagnostics Guide</i>	Description of features such as mirroring and sFlow, Operations, Administration, and Maintenance (OAM) and diagnostics tools. Provides configuration examples for mirroring, sFlow, and OAM tools.
<i>SR Linux Multicast guide</i>	Basic configuration concepts for multicast routing. Provides configuration examples for Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), and Protocol Independent Multicast (PIM).
<i>SR Linux Software Release Notes</i>	The most up-to-date information about supported features, supported hardware, known limitations, and resolved issues.

4 Hardware overview

SR Linux is supported on a variety of Nokia hardware platforms that can be deployed in the network as WAN/data center interconnect (DCI), spine, and leaf/top of rack (TOR) devices. The following diagram shows typical deployment for SR Linux hardware platforms.

Figure 4: SR Linux hardware platforms deployed as WAN/DCI, spine, and leaf/TOR devices



sw4765

The linked sections in the following table describe the specifications of each hardware platform group. Each router series has a dedicated installation guide containing complete specifications, recommendations for preparing the installation site, and procedures to install and ground the routers. See the respective chassis installation guides listed in [SR Linux documentation](#) for more information.

Table 4: SR Linux usage by hardware platform group

Usage	Hardware Platform Group
WAN / DC interconnect	7250 IXR platforms
Spine	7220 IXR-Hx platforms 7250 IXR platforms
Leaf/ToR	7215 IXS platforms 7220 IXR-Dx platforms

The following table lists the devices that make up each hardware platform group that supports SR Linux.

Table 5: Supported platforms

Platform group	Description
7215 IXS	<ul style="list-style-type: none"> 7215 IXS-A1
7220 IXR-Dx	<ul style="list-style-type: none"> 7220 IXR-D1 7220 IXR-D2

Platform group	Description
	<ul style="list-style-type: none"> • 7220 IXR-D2L • 7220 IXR-D3 • 7220 IXR-D3L • 7220 IXR-D4 • 7220 IXR-D5
7220 IXR-Hx	<ul style="list-style-type: none"> • 7220 IXR-H2 • 7220 IXR-H3 • 7220 IXR-H4 • 7220 IXR-H4-32D • 7220 IXR-H5-32D • 7220 IXR-H5-64D • 7220 IXR-H5-64O
7250 IXR Gen 2	<ul style="list-style-type: none"> • 7250 IXR-6 • 7250 IXR-10
7250 IXR Gen 2c+	<ul style="list-style-type: none"> • 7250 IXR-6e • 7250 IXR-10e • 7250 IXR-X1b • 7250 IXR-X3b
7250 IXR Gen 3	<ul style="list-style-type: none"> • 7250 IXR-6e • 7250 IXR-10e • 7250 IXR-18e • 7250 IXR-X4
7730 SXR	<ul style="list-style-type: none"> • 7730 SXR-1-32D • 7730 SXR-1d-32D • 7730 SXR-1x-44S

The sections that follow describe the specifications of each platform. Each router series has a dedicated installation guide containing complete specifications, recommendations for preparing the installation site, and procedures to install and ground the routers. See the respective chassis installation guides listed in [SR Linux documentation](#) for more information.

4.1 7215 IXS platforms

SR Linux is supported on the 7215 IXS-A1 hardware platform. The following table summarizes the features of the router.

Table 6: 7215 IXS-A1 specifications

Parameter	7215 IXS-A1
Height	1 RU
Width	43.84 cm
Depth	25.4 cm
System Capacity (FD)	88G
Module Connectors	4 x 10G SFP+ ports and 48 x 1G RJ45 ports

4.1.1 Architecture

The 7215 IXS-A1 is a top-of-the rack data center management switch. It is designed to provide management connectivity for leaf-spine data center fabrics in enterprise, service provider, and webscale data center environments.

4.1.2 Chassis components

The 7215 IXS-A1 chassis is a 1 RU platform with the following features:

- 4 x 10G SFP+ ports and 48 x 1G RJ45 ports
- USB 2.0 port for future software updates and configuration
- system status LEDs:
 - Stat
 - Fan
 - PS
- dual built-in redundant AC or DC power supply units (PSUs)
- front-to-back or back-to-front cooling airflow with redundant fans
- one RJ45 console (RS-232 serial interface) connector
- designed for rack-mounting with fixed brackets or slide rails (included with the chassis)

4.1.3 Power and cooling

The 7215 IXS-A1 chassis can be powered by two built-in redundant AC or DC PSUs.

The 7215 IXS-A1 chassis is equipped with built-in fans with front-to-back or back-to-front airflow for cooling. The fans are controlled by the system software, and their speed is set according to the environmental temperature surrounding the chassis.

4.2 7220 IXR-Dx platforms

SR Linux is supported by the 7220 IXR-D1, 7220 IXR-D2, 7220 IXR-D2L, 7220 IXR-D3, 7220 IXR-D3L, 7220 IXR-D4, and 7220 IXR-D5 platforms. The following table summarizes the features of these routers.

Table 7: 7220 IXR-Dx platform specifications

Parameter	7220 IXR-D1	7220 IXR-D2	7220 IXR-D2L	7220 IXR-D3	7220 IXR-D3L	7220 IXR-D4	7220 IXR-D5
Height	1 RU	1 RU	1 RU	1 RU	1 RU	1 RU	1 RU
Width	43.8 cm	43.8 cm	43.8 cm	43.8 cm	43.8 cm	43.8 cm	43.8 cm
Depth	40 cm	46 cm	53.6 cm	46 cm	51.5 cm	53.6 cm	59 cm
System Capacity (FD)	88G	2T	2T	3.2T	3.2T	6T	12.8T
Module connectors	48T 4SFP+	48SFP28 8QSFP28	48SFP28 8QSFP28 2SFP+	32QSFP28 2SFP+	32QSFP28 2SFP+	28QSFP28 8QSFP-DD	32QSFPDD 2SFP+

4.2.1 Architecture

The 7220 IXR-Dx platforms are high-performance, high-density, fixed configuration devices designed for data center leaf-spine deployments. The 7220 IXR-Dx platforms deliver capabilities including IP routing, Layer 2 switching, QoS, router security, scalable telemetry, and model-driven management.

4.2.2 Chassis components

The 7220 IXR-Dx chassis vary in system capacity. [7220 IXR-Dx platforms](#) compares the differences among 7220 IXR-D1, 7220 IXR-D2, 7220 IXR-D2L, 7220 IXR-D3, 7220 IXR-D3L, 7220 IXR-D4, and 7220 IXR-D5 platforms. They provide high-density QSFPDD, QSFP28, QSFP+, SFP28, SFP+, SFP, and RJ-45 ports. The router supports native 400GE, 100GE, 50GE, 40GE, 25GE, 10GE, and 1GE port options.

4.2.3 Power and cooling

The 7220 IXR-Dx platform chassis can be powered by either dual removable AC PSUs or DC PSUs. The chassis support two PSUs with 1+1 redundancy.

The chassis are equipped with N+1 hot-swappable fans with front-to-back or back-to-front airflow for cooling. The 7220 IXR-D1 has three fan modules, 7220 IXR-D2 has four fan modules, 7220 IXR-D3 has

five fan modules, 7220 IXR-D2L, 7220 IXR-D3L, 7220 IXR-D4, and 7220 IXR-D5 platforms have six fan modules.

4.3 7220 IXR-Hx platforms

SR Linux is supported by the 7220 IXR-H2, 7220 IXR-H3, 7220 IXR-H4, 7220 IXR-H4-32D, 7220 IXR-H5-32D, 7220 IXR-H5-64D, and 7220 IXR-H5-64O platforms. The following table summarizes the features of these routers.

Table 8: 7220 IXR-H specifications

Parameter	7220 IXR-H2	7220 IXR-H3	7220 IXR-H4	7220 IXR-H4-32D	7220 IXR-H5-32D	7220 IXR-H5-64D	7220 IXR-H5-64O
Height	4 RU	1 RU	2 RU	1 RU	1 RU	2 RU	2 RU
Depth	55 cm	55 cm	64.9 cm	55 cm	65 cm	63 cm	63 cm
System Capacity (FD)	12.8T	12.8T	25.6T	12.8T	25.6T	51.2T	51.2T
Module connectors	128 QSFP28	32QSFP-DD 2SFP+	64QSFP-DD 2SFP+	32QSFP-DD 1SFP+	32QSFP-112DD 2SFP+	64QSFP-112DD 2SFP+	64OSFP 2SFP+

4.3.1 Architecture

The 7220 IXR-H platforms are high-performance, high-density, fixed configuration devices designed for data center leaf-spine deployments. The 7220 IXR-H platforms deliver capabilities including IP routing, Layer 2 switching, QoS, router security, scalable telemetry, and model-driven management.

4.3.2 Chassis components

The 7220 IXR-H chassis vary in system height. [Table 8: 7220 IXR-H specifications](#) summarizes the differences between the platform variants. They provide high-density QSFPDD, OSFP, QSFP28, SFP+, and RJ-45 ports. The router supports native 800GE, 400GE, 200GE, 100GE, 50GE, 40GE, 25GE, and 10GE port options.

4.3.3 Power and cooling

The 7220 IXR-H platform chassis can be powered by either removable AC PSUs or DC PSUs. The maximum PSUs supported are as follows:

- 7220 IXR-H2 supports four PSUs with 2+2 redundancy
- 7220 IXR-H3 supports two PSUs with 1+1 redundancy
- 7220 IXR-H4 supports two PSUs with 1+1 redundancy

- 7220 IXR-H4-32D supports two PSUs with 1+1 redundancy
- 7220 IXR-H5-32D supports two PSUs with 1+1 redundancy
- 7220 IXR-H5-64D supports two PSUs with 1+1 redundancy
- 7220 IXR-H5-64O supports two PSUs with 1+1 redundancy

The chassis are equipped with N+1 hot-swappable fans.

- 7220 IXR-H2 has eight fan modules and supports front-to-back or back-to-front airflow for cooling
- 7220 IXR-H3 has six fan modules and supports front-to-back or back-to-front airflow for cooling
- 7220 IXR-H4 has four fan modules and supports only front-to-back airflow for cooling
- 7220 IXR-H4-32D has seven fan modules and supports front-to-back or back-to-front airflow
- 7220 IXR-H5-32D has seven fan modules and supports front-to-back airflow
- 7220 IXR-H5-64D has four fan modules and supports front-to-back airflow
- 7220 IXR-H5-64O has four fan modules and supports front-to-back airflow

4.4 7250 IXR platforms

SR Linux is supported on the 7250 IXR-6, 7250 IXR-10, 7250 IXR-6e, 7250 IXR-10e, and 7250 IXR-18e hardware platforms. The following table summarizes the specifications of each product.

Table 9: 7250 IXR platform specifications

Parameter	7250 IXR-6	7250 IXR-10	7250 IXR-6e	7250 IXR-10e	7250 IXR-18e
Height	7 RU	13 RU	10 RU	16 RU	35 RU
Depth	81.28 cm	81.28 cm	92.2 cm	92.2 cm	104.1 cm
Number of IMM slots	4	8	4	8	16
Slot capacity (full duplex)	9.6 T	9.6 T	14.4 T	14.4 T	28.8 T
Maximum system capacity per chassis (half duplex)	76.1 T or 115.2 T HD with local switching	153.6 T or 230.4 T with local switching	115.2 T	230.4 T	460.8 T

4.4.1 Architecture

The 7250 IXR platforms deliver capabilities that include IP routing, Layer 2 Ethernet, QoS, router security, scalable telemetry and model-driven programmability. Flexible traffic management includes big buffering, and per-port queuing, shaping and policing.

Each chassis uses an orthogonal direct cross-connect architecture, with Integrated Media Modules (IMMs) connecting in front and switch fabrics and fans connecting at the rear. The lack of a backplane, midplane,

or midplane connector system provides a compact chassis design, optimal cooling, and easy capacity upgrades.

4.4.2 Chassis components

The 7250 IXR chassis include fan trays, Power Supply Units (PSUs), Control Processing Modules (CPMs), Switch Fabric Modules (SFMs), and IMMs. The chassis vary in system capacity, height, and number of IMM slots. [Table 9: 7250 IXR platform specifications](#) summarizes the differences between the platform variants.

The system uses a complete Faraday Cage design to ensure EMI containment, a critical requirement for platform evolution that will support next-generation Application-Specific Integrated Circuits (ASICs). The routers are high-density, high-performance modular devices that are designed for data spine deployments. They provide hardware support for 400GE, 100GE, 40GE, 25GE, and 10GE interfaces for intra-fabric and server connectivity.

4.4.3 Power and cooling

The 7250 IXR platforms can be AC or DC powered with hot-swappable and load-sharing PSUs. The PSUs are N+M power redundant, and each PSU is cooled by a fan independent of the chassis fans. The maximum PSUs supported are as follows:

- 7250 IXR-6 - 6 PSUs
- 7250 IXR-10 - 10 PSUs
- 7250 IXR-6e - 9 PSUs
- 7250 IXR-10e - 12 PSUs
- 7250 IXR-18e - 24 PSUs

The chassis have dual fans that support front-to-back airflow. The fan design uses a stainless-steel orthogonal mesh honeycomb placed in front of the line card for air intake. The perforation rate in the honeycomb is approximately 90%, which allows a large surface area that is exposed to cool air, causing a reduction in power consumption.

4.5 7250 IXR-X platforms

SR Linux is supported on 7250 IXR-X platforms, which includes the 7250 IXR-X1b, 7250 IXR-X3b, and 7250 IXR-X4 hardware platforms. All platform variants are 1 RU in height. The following table summarizes the features of these routers.

Table 10: 7250 IXR-X series specifications

Parameter	7250 IXR-X1b	7250 IXR-X3b	7250 IXR-X4
Depth	54.7 cm	64.8 cm	64.8 cm
System capacity (full duplex)	7.2T	14.4T	25.6T

Parameter	7250 IXR-X1b	7250 IXR-X3b	7250 IXR-X4
Transceiver cages	24 QSFP28, 12 QSFP-DD	36 QSFP56-DD	32 QSFP112-DD

4.5.1 Architecture

The 7250 IXR-X platforms are high-performance, high-density devices that deliver capabilities including IP routing, Layer 2 Ethernet, QoS, router security, scalable telemetry and model-driven programmability. Flexible traffic management includes big buffering, and per-port queuing, shaping and policing.

4.5.2 Chassis components

The 7250 IXR-X platform chassis vary in system capacity. The 7250 IXR-X1b and 7250 IXR-X3b platforms support high-density QSFP56-DD ports. The 7250 IXR-X4 platform supports QSFP112-DD ports.

4.5.3 Power and cooling

The 7250 IXR-X platform chassis can be powered by either dual removable AC PSUs or LVDC PSUs. The chassis support two PSUs with 1+1 redundancy. On the rear of the chassis, there are two slots for the PSUs.

The 7250 IXR-X platform chassis are equipped with N+1 hot-swappable fans with front-to-back or back-to-front airflow for cooling. On the rear of the chassis, there are three slots for the fan trays.

4.6 7730 SXR-1 platforms

SR Linux is supported on 7730 SXR-1 platforms, which includes the 7730 SXR-1-32D, 7730 SXR-1d-32D and 7730 SXR-1x-44S platforms. The following table summarizes the features of these routers.

Table 11: 7730 SXR platform specifications

Parameter	7730 SXR-1-32D	7730 SXR-1d-32D	7730 SXR-1x-44S
Height	1 RU	2 RU	1 RU
Depth	45 cm	27.5 cm	45 cm
Slot capacity (full duplex)	4.4 T	4.4 T	4.4 T
Module connectors	28 QSFP28, 4 QSFPDD	28 QSFP28, 4 QSFPDD	28 QSFP28, 4 QSFPDD

4.6.1 Architecture

The 7730 SXR-1 platforms deliver comprehensive access, aggregation and edge features in a fixed compact platform with superior scale, capability, performance, longevity, and upgradability.

4.6.2 Chassis components

The 7730 SXR chassis vary in system capacity. They provide high-density QSFPDD, QSFP28, SFPDD, and RJ 45 ports. The routers support up to 400GE native options on QSFPDD ports and up to 100GE on QSFP28 and SFPDD ports.

4.6.3 Power and cooling

The 7730 SXR-1 platform chassis can be powered by either removable AC PSUs or DC PSUs. The chassis support two PSUs with 1+1 redundancy.

The chassis are equipped with hot-swappable fans. The 7730 SXR-1-32D and 7730 SXR-1x-44S have three fan modules. The 7730 SXR-1d-32D has two fan modules. The systems only support front-to-back airflow direction.

5 SR Linux architecture overview

The sections that follow describe components of the SR Linux architecture and how they work together.

5.1 SR Linux components

At a high level, SR Linux can be divided into three components: a mainline Linux kernel, an operating system, and a suite of modular, lightweight applications, each supporting a different protocol or function (IS-IS, BGP, AAA, and so on). These applications use gRPC and APIs to communicate with each other and external systems over TCP.

5.1.1 Linux kernel

The kernel is the core of a computer operating system, with complete control over everything in the system. The kernel is loaded and executed by the boot loader, then handles the system startup as well as input/output requests from software, translating them into data-processing instructions for the CPU. The kernel handles memory and peripherals, as well as interactions with the switch ASIC.

SR Linux uses a mainline Linux kernel with no modifications or customization. This allows kernel upgrades to be performed using normal upgrade mechanisms (yum on CentOS/RHEL-based systems, and aptitude on Debian-based systems).

5.1.2 Operating system

SR Linux uses a mainline Linux kernel without any modifications or customization. This ensures a seamless upgrade process for SR Linux users.


See the "Appendix: Migrating from CentOS to Debian OS" chapter of the *SR Linux Software Installation Guide* for information about the high-level changes in regard to migration from CentOS to Debian OS.

To follow the Linux Filesystem Hierarchy Standard (FHS), SR Linux software is distributed and laid out as a closed-source, third-party application on the OS. The FHS clearly defines where configuration and binaries should be kept; for SR Linux, all application configuration is kept in the `/opt/srlinux` directory.

5.1.3 Modular applications

SR Linux is a suite of applications running like any others would in a Linux environment. The applications communicate with the IDB to process configuration and state. [Figure 5: SR Linux applications and IDB](#) shows the relationship between the IDB and the applications that run in SR Linux.

Figure 5: SR Linux applications and IDB



sv4049

Messaging between applications is controlled by IDB, and configuration via supported APIs is controlled by the management server application.

[Table 12: SR Linux applications](#) describes the key SR Linux applications. Use the **show system application** command to display information about all applications running on the system.

Table 12: SR Linux applications

Application name	Description
IDB	Controls messaging between SR Linux applications.
mgmt_server	Controls interaction between external APIs and the IDB, handles application-specific YANG models, and translates them into protobufs for IDB. See SR Linux management server .
aaa_mgr	Performs AAA for end users connecting to the system.
fib_mgr	Responsible for route resolution and route selection.
lldp_mgr	Responsible for sending and receiving LLDP packets via XDP.

Application name	Description
static_route_mgr	Creates/updates/deletes static routes.
arp_nd_mgr	Responsible for ARP resolution for IPv4 and Neighbor Discovery for IPv6.
bgp_mgr	Runs the BGP control plane.
chassis_mgr	Monitors interface/subinterface state and synchronizes interfaces between Linux and the ASIC.
xdp_mgr	Handles data path programming. The xdp_mgr runs on each line card and is split into two components: platform independent and platform dependent. The platform-independent component handles communication between IDB and the ASIC, and the platform-dependent component is an extensible framework for plugging in data plane programming software.
linux_mgr	Creates routes and neighbors in Linux. As control packets enter the Linux host (and kernel), the Linux network stack responds to them. Synchronizing routes and neighbors to Linux stops Linux from ARPing/routing via the default gateway when routes exist at the data plane but not within Linux.
plcy_mgr	Enforces routing policies.
app_mgr	Monitors the health of the processes running SR Linux applications, and restarts them if they fail.
net_inst_mgr	Synchronizes VRFs between the switch ASIC and Linux.
log_mgr	Controls the log infrastructure, implemented by rsyslog.
acl_mgr	Controls ACLs in the system, both on the Linux host and on the CPM.
bfd_mgr	Controls BFD sessions on the system.

IDB stores data as protobufs. Protobufs use a .proto file to define how data is structured. Protobufs are not human-readable, but an IDB client is provided with the IDB server, which allows you to read from different topics, decoding the protobuf in real time.

5.1.3.1 Application manager

The application manager is responsible for monitoring the health of the processes running each SR Linux application, and restarting them if they fail.

Each application has specific YAML configuration. The application manager reads in the application-specific YAML configuration and starts each application. It allows applications to not start if no configuration exists for them.

The application manager functions as a replacement for the Linux systemd. At boot time, systemd starts the application manager, which in turn starts SR Linux applications that it manages (based on YAML configuration). The device manager (dev_mgr) is the first application started, then IDB, then other applications are started based on the YAML configuration. The application manager loads the YANG model for each application into the management server.

The application manager is also responsible for restarting the entire system if a critical application cannot be restarted successfully; this restart mechanism is controlled through configuration.

6 SR Linux management overview

This chapter describes the system management functions of SR Linux, including the role of the SR Linux management server and external management APIs (CLI, gNMI, and JSON-RPC). It describes SR Linux configuration modes, methods for securing access to the device, and logging functions.

6.1 SR Linux management server

Configuration and state for the modular SR Linux applications are driven by centralized data models (YANG) that are managed by the SR Linux management server application.

The SR Linux management server provides a central point for external clients and APIs to access the system. The supported external APIs (CLI, JSON-RPC, and gNMI) communicate with the management server via its gRPC interface.

The management server manages the YANG models, which are loaded into the management server by the application manager, based on the requirements of each application. The management server translates these models into protobufs for the IDB. This allows other applications to read their own configuration, by subscribing to the management server topic representing the configuration.

The management server owns the configuration of each application, so each application does not have write access to its own configuration. This provides a central point of configuration enforcement.

Agents built with the NDK function similar to other applications provided with SR Linux. SR Linux applications share state details with each other using a publish/subscribe (pub/sub) architecture. Agents have their own table space within the IDB and can subscribe and receive a notification to events occurring on the device, or create their own table space and publish data to it. This data can be read by other applications within SR Linux, allowing route modifications by publishing routes to the IDB for selection by the FIB manager.

6.2 SR Linux CLI

The SR Linux CLI is an interactive interface for configuring, monitoring, and maintaining the SR Linux via an SSH or console session. The SR Linux CLI operates as a client that communicates with the management server via gRPC. The command `t tree` in the CLI is derived from the SR Linux YANG models.

The SR Linux CLI supports command autocompletion, aliases, annotation, and standard Linux output modifiers such as `grep`. Command output can be displayed in JSON format.

See the “CLI interface” chapter of the *SR Linux System Management Guide* for information about CLI features.

6.3 JSON-RPC server

A JSON-RPC server can be enabled on the SR Linux device, which allows JSON-formatted requests to be issued to the device to retrieve and set configuration and state. You can use the JSON-RPC API to run CLI commands and standard Get and Set methods. The SR Linux device returns responses in JSON-format.

When the JSON-RPC server is enabled, the application passes the requests to the SR Linux management server via the gRPC interface. This JSON-RPC API uses HTTP and HTTPS for transport, and users are authenticated with the `aaa_mgr` application. HTTPS requests can be authenticated using TLS.

See the "JSON interface" chapter of the *SR Linux System Management Guide* for more information.

6.4 gRPC server

The gRPC-based gNMI protocol is used for the modification and retrieval of configuration from a target device, as well as the control and generation of telemetry streams from a target device to a data collection system.

SR Linux can enable a gRPC server that allows external gRPC clients to connect to the device and modify the configuration and collect state information.

When the gRPC server is enabled, the SR Linux `gnmi_mgr` application functions as a target for gRPC clients. The `gnmi_mgr` application validates gNMI clients and passes Get, Set, and Subscribe RPCs to the SR Linux `mgmt_server` application via the gRPC interface.

See the "gNMI interface" chapter in the *SR Linux System Management Guide* for information about the supported RPCs.

6.5 gNOI

The gRPC Network Operations Interface (gNOI) defines a set of gRPC-based services for executing operational commands on network devices. The individual RPCs and messages that perform the operations required for certificate management on the node are defined at the following location: <https://github.com/openconfig/gnoi>.

The gNOI file service on SR Linux allows a client to transfer files to and from a target node. This service can be used to extract debugging information through the transfer of system logs and core files.

See the "gNOI" chapter of the *SR Linux System Management Guide* for more information.

6.6 NETCONF

NETCONF is a standardized IETF configuration management protocol, specified in RFC 6241. It is an XML-based protocol that can be used as an alternative to the CLI or gNMI for managing SR Linux.

The SR Linux NETCONF interface supports configuration and state. NETCONF uses RPC messaging to facilitate communication between a NETCONF client and the NETCONF server running on the SR Linux device. RPC message and configuration or state data is encoded in XML documents, which are exchanged

between the NETCONF client and NETCONF server in a series of request and response message interactions.

See the "NETCONF" chapter of the *SR Linux System Management Guide* for more information.

6.7 SNMP

SNMP is an application-layer protocol that enables communication between managers (the management system) and agents (the network devices). It provides a standard framework to monitor devices in a network from a central location.

An SNMP manager can get a value from an SNMP agent. The manager uses definitions in the management information base (MIB) to perform operations on the managed device such as retrieving values from variables and processing traps.

The following actions can occur between the agent and the manager:

- The manager gets information from the agent.
- The agent sends traps to notify the manager of significant events that occur on the system.

The SNMP agent provides management information to support object identifiers (OIDs) in standard and Nokia MIBs.

SR Linux provides a customizable SNMP framework that allows you to define custom MIBs for get requests and traps. This framework also powers SR Linux's built-in MIBs and traps, providing flexibility to customize SNMP MIBs to the specific requirements for your network. The framework consists of:

- Mapping files (YAML): Define MIB tables and OIDs.
- Conversion scripts (uPython): Process data from the management server via gNMI and convert it for SNMP.

See the "SNMP" chapter of the *SR Linux System Management Guide* for more information.

6.8 Zero Touch Provisioning

Zero Touch Provisioning (ZTP) automates the process of booting the SR Linux device, obtaining an address on the network, then downloading and executing a Python script to configure the system.

The node storage device ships with the SR Linux image and the `grub.cfg` for auto-boot using the out-of-band port or the in-band. When the system boots, if `autoboot = enabled` is set in the `grub.cfg` file, it initiates the ZTP process. The system then obtains an IP address via DHCPv4 or v6, downloads a Python script to configure the device, and executes the script.

The script can contain URLs to an updated image and a new kernel. The ZTP process can download these and place them on the compact flash, where they become active at the next reboot.

See the "Zero Touch Provisioning" chapter of the *SR Linux Software Installation Guide* for information about using ZTP to initialize SR Linux.

6.9 Secure Boot

Secure Boot ensures that the software executed by the system is trusted and originated from Nokia IP Routing. The system verifies that the boot chain up to the operating system is digitally signed by Nokia, ensuring the integrity and authenticity of every software element.

See "Secure Boot" in the *SR Linux Software Installation Guide* for more information.

6.10 Measured Boot

Measured Boot complements Secure Boot by providing cryptographic evidence in local or remote attestation to verify that the system booted in a known and trusted state. The state of each component in the boot process is measured and recorded at every boot in the Trusted Platform Module (TPM) of the control card.

See "Measured Boot" in the *SR Linux Software Installation Guide* for more information.

6.11 BootZ

SR Linux supports the secure bootstrapping protocol, BootZ. BootZ allows operators to securely bootstrap their devices following the zero-touch provisioning approach, eliminating the need for staging or pre-configuration before deployment.

See the "BootZ" chapter of the *SR Linux Software Installation Guide* for more information.

6.12 SR Linux configuration

SR Linux uses transaction-based configuration, which allows the operator to make changes to the configuration, and then explicitly commit the configuration to apply it.

By default, the SR Linux configuration file is located in `/etc/opt/srlinux/config.json`. At boot time, the management server loads the configuration and publishes content to IDB for applications to consume.

Configuration modes define how the system is running when transactions are performed. Supported modes are the following:

- **Candidate** – is used to modify a configuration. Modifications are not applied to the running system until a **commit** command is issued. When committed, the changes are copied to the running configuration and become active.
- **Running** – is used to display the currently running or active configuration. Configurations cannot be edited in this mode.

When a configuration is committed, it is first validated for YANG syntax, then tested by each application, then validated by the forwarding plane. If validation fails at any stage, the configuration is not committed.

A configuration candidate can be either shared or private:

- **Shared** – is the default configuration candidate for CLI sessions. Multiple users can modify the shared candidate concurrently. When the configuration is committed, the changes from all of the users are applied.
- **Private** – is the default configuration candidate when using JSON-RPC or gNMI clients, and can optionally be used in the CLI. With a private candidate, each user modifies their own separate instance of the configuration candidate. When a user commits their changes, only the changes from that user are committed.

By default, there is a single unnamed global configuration candidate. You can optionally configure one or more named configuration candidates, which function identically to the global configuration candidate. Both shared and private configuration candidates support named versions.

You can optionally create a rescue configuration, which is loaded if the startup configuration fails to load. If the startup configuration fails to load, and no rescue configuration exists, the system is started using the factory default configuration.

When you upgrade the SR Linux software image, the configuration in the startup `config.json` file is read into the running configuration and automatically upgraded to ensure compatibility with the new software version.

See the "Configuration management" chapter in the *SR Linux Configuration Basics Guide* for more information.

6.13 Securing access

SR Linux is able to secure access to the device for users connecting via SSH or the console port, as well as for applications and FTP access. SR Linux performs authentication, authorization, and accounting (AAA) functions for each user type.

Authentication can be performed for users configured within the underlying Linux OS, and for administrative users configured within SR Linux.

Authorization is performed through role-based access control. Users can be configured with a set of one or more roles that indicate the privileges for which they are authorized in the system. You can configure SR Linux to use information from a TACACS+ or RADIUS server to assign roles to an authenticated user.

SR Linux supports command accounting, including the entire CLI string that a user enters on the command line, including any pipes or output redirects specified in the command. The accounting records can be sent to a destination such as a TACACS+ or RADIUS server group or the local system.

See the "Securing access" chapter of the *SR Linux Configuration Basics Guide* for more information.

6.14 SR Linux logging

SR Linux implements logging via the standard Linux syslog libraries. The SR Linux device uses rsyslog in the underlying Linux OS to filter logs and pass them on to remote servers or other specified destinations.

SR Linux supports configuration of Linux facilities and SR Linux subsystems as sources for log messages to filter. See the "Logging" chapter of the *SR Linux Configuration Basics Guide* for information about configuring input sources, filters, and output destinations for log messages.

See the *SR Linux Log Events Guide* for properties and descriptions of the log messages that can be generated by SR Linux subsystems.

6.15 P4Runtime support

Programming Protocol-Independent Packet Processors (P4) is an open-source language for programming the data plane on networking devices. P4Runtime is an API for controlling the data plane on devices defined in a P4 program. The P4 language and P4Runtime specification are maintained at p4.org.

The SR Linux eXtensible Data Path (XDP) is not programmed in P4, and no programming of the data path is done using P4. However, SR Linux supports using P4 programs and P4Runtime for some use cases that involve packet input/output. These include the following:

- Packet injection/reception, where a P4Runtime client injects packets on one side of a link and receives them on the other. The P4Runtime client then builds a topology based on embedded data within the payload of each packet.
- Redirecting traceroute packets with TTL=0, TTL=1, and TTL=2 to a P4Runtime client, so they can be enriched with information that is not visible to the device.

See the *SR Linux P4RT Guide* for configuration information.

6.16 Event Handler

Event Handler is a framework that enables SR Linux to react to specific system events, using programmable logic for the actions to take in response to the events.

The Event Handler framework allows you to write a custom script and have the script be invoked when specific events occur, such as when a port goes operationally down. The script can generate a list of actions to be executed on the SR Linux device. The actions can include updating the SR Linux configuration, changing the operational state of a group of ports, executing a **tools** command, or running another script.

A primary use for Event Handler is the ability to change the operational state of one group of ports based on the state of another group of ports. This type of usage is known as operational groups. See the *SR Linux Event Handler Guide* for more information.

6.17 gRIBI

The gRIBI (gRPC Routing Information Base Interface) is a gRPC-based protocol that allows external clients to inject routes into the RIB of a network device. The gRIBI clients add and remove RIB entries using an API that is defined by the `gribi.proto`. RIB entries are defined using the OpenConfig Abstract Forwarding Table (AFT) model, translated to protobuf.

The protobuf definition of gRIBI is maintained in the gRIBI GitHub repository: `gribi.proto`. The corresponding `gribi_aft.proto` containing the OpenConfig AFT model is available in the same repository.

See the *SR Linux gRIBI Guide* for more information, including a list of supported RPCs and AFTs.

6.18 Network Synchronization

SR Linux supports network synchronization capabilities, including physical layer frequency distribution via Synchronous Ethernet and packet based distribution of time using the precision time protocol (PTP) of IEEE 1588.

The synchronization network is designed so a clock always receives timing from a clock of equal or higher stratum level or quality level. This ensures that if an upstream clock has a fault condition (for example, it loses its reference and enters a holdover or free-run state) and begins to drift in frequency, the downstream clock is able to follow it. For greater reliability and robustness, most offices and nodes have at least two synchronization references that can be selected in priority order (such as primary and secondary).

Additional resiliency can be provided by the ability of the node clock to operate within prescribed network performance specifications in the absence of any reference for a specified period. A clock operating in this mode is said to hold the last known state over (or holdover) until the reference lock is once again achieved. Each level in the timing hierarchy is associated with minimum levels of network performance.

Each synchronization-capable port can be independently configured to transmit data using the node reference timing.

Specifically for synchronous Ethernet, transmission of a reference clock through a chain of Ethernet equipment requires that all equipment supports synchronous Ethernet. A single piece of equipment that is not capable of performing synchronous Ethernet breaks the chain. Ethernet frames still get through, but downstream devices should not use the recovered line timing because it is not traceable to an acceptable stratum source.

7 SR Linux interfaces

This chapter describes SR Linux interface types, subinterfaces, and support for Link Aggregation Groups (LAGs). See the "Interfaces" chapter in the *SR Linux Interfaces Guide* for configuration examples.

7.1 SR Linux interface types

On the SR Linux device, an interface is any physical or logical port through which packets can be sent to or received from other devices.

SR Linux supports the following interface types:

- Loopback

Loopback interfaces are virtual interfaces that are always up, providing a stable source or destination from which packets can always be originated or received. SR Linux supports up to 256 loopback interfaces system-wide, across all network instances. Loopback interfaces are named **loN**, where *N* is 0 to 255.

- System

The system interface is a type of loopback interface that has characteristics that do not apply to regular loopback interfaces:

- The system interface can be bound to the default network-instance only.
- The system interface does not support multiple IPv4 addresses or multiple IPv6 addresses.
- The system interface cannot be administratively disabled. When configured, it is always up.

SR Linux supports a single system interface named **system0**. When the system interface is bound to the default network-instance, and an IPv4 address is configured for it, the IPv4 address is the default local address for multi-hop BGP sessions to IPv4 neighbors established by the default network-instance, and it is the default IPv4 source address for IPv4 VXLAN tunnels established by the default network-instance. The same functionality applies with respect to IPv6 addresses / IPv6 BGP neighbors / IPv6 VXLAN tunnels.

- Network

Network interfaces carry transit traffic, as well as originate and terminate control plane traffic and in-band management traffic.

The physical ports in line cards installed in the SR Linux device are network interfaces. A typical line card has a number of front-panel cages, each accepting a pluggable transceiver. Each transceiver may support a single channel or multiple channels, supporting one Ethernet port or multiple Ethernet ports, depending on the transceiver type and its breakout options.

In the SR Linux CLI, each network interface has a name that indicates its type and its location in the chassis. The location is specified with a combination of slot number and port number, using the following format:

```
ethernet-slot[/mda-slot][/connector]/port
```

where:

- *slot* is a slot number (1 to 9). On 7220 IXR-Dx and 7220 IXR-Hx fixed systems there is only one slot, numbered 1; even though no other slot value is possible, the slot number 1 is still part of the interface name.
- *mda-slot* is a number representing the media dependent adapter (MDA) position within the parent card or chassis on devices that support MDA assemblies.
- *connector* is a number indicating the front-panel connector/cage to which a breakout cable is connected. It is omitted when there is no breakout configuration.
- *port* is a number referring to the front-panel connector/cage to which a non-breakout cable is connected, or else it refers to the channel number (1 to 4) in a breakout configuration.

For example:

- Interface ethernet-1/1 can refer to port 1 of a 7220 IXR-Dx system or to port 1 of the 7250 IXR IMM installed in slot 1.
 - Interface ethernet-2/1 refers to port 1 of the 7250 IXR IMM installed in slot 2.
 - Interface ethernet-1/1/1 can refer to breakout port 1 of connector 1 in a 7220 IXR-Dx chassis.
 - Interface ethernet-1/m1/2/4 refers to breakout port 4 of connector 2 on the first MDA of the chassis.
- Management
Management interfaces are used for out-of-band management traffic. SR Linux supports a single management interface named **mgmt0**.
The **mgmt0** interface supports the same functionality and defaults as a network interface, except for the following:
 - Packets sent and received on the **mgmt0** interface are processed completely in software.
 - The **mgmt0** interface does not support multiple output queues, so there is no output traffic differentiation based on forwarding class.
 - The **mgmt0** interface does not support pluggable optics. It is a fixed 10/100/1000-BaseT copper port.
 - Integrated Routing and Bridging (IRB)
IRB interfaces enable inter-subnet forwarding. Network instances of type **mac-vrf** are associated with a network instance of type **ip-vrf** via an IRB interface.

7.2 LAG interfaces

A LAG, based on the IEEE 802.1ax standard (formerly 802.3ad), increases the bandwidth available between two network devices, depending on the number of links installed. A LAG also provides redundancy if one or more links participating in the LAG fail. All physical links in a LAG combine to form one logical interface.

LAGs can be either statically configured, or formed dynamically with Link Aggregation Control Protocol (LACP). Load sharing is executed in hardware, which provides line rate forwarding for all port types. A LAG consists of ports of the same speed.

7.3 Subinterfaces

On the SR Linux device, each type of interface can be subdivided into one or more subinterfaces. A subinterface is a logical channel within its parent interface.

Traffic belonging to one subinterface can be distinguished from traffic belonging to other subinterfaces of the same port using encapsulation methods such as 802.1Q VLAN tags.

While each port can be considered a shared resource of the router that is usable by all network instances, a subinterface can only be associated with one network instance at a time. To move a subinterface from one network instance to another, you must disassociate it from the first network instance before associating it with the second network instance.

You can configure ACL policies to filter IPv4 and IPv6 packets entering or leaving a subinterface.

SR Linux supports policies for assigning traffic on a subinterface to forwarding classes or remarking traffic at egress before it leaves the router. DSCP classifier policies map incoming packets to the appropriate forwarding classes, and DSCP rewrite-rule policies mark outgoing packets with an appropriate DSCP value based on the forwarding class.

7.4 DHCP relay

DHCP relay refers to the router's ability to act as an intermediary between DHCP clients requesting configuration parameters, such as a network address, and DHCP servers when the DHCP clients and DHCP servers are not attached to the same broadcast domain, or do not share the same IPv6 link (in the case of DHCPv6).

SR Linux supports DHCP relay for IRB subinterfaces and Layer 3 subinterfaces. The DHCP server network can be in the same IP-VRF network-instance of the Layer 3 subinterfaces that require DHCP relay, or it can be in a different IP-VRF network-instance or the default network instance.

7.5 LLDP

The IEEE 802.1ab Link Layer Discovery Protocol (LLDP) is implemented in SR Linux using the `lldp_mgr` application, which controls sending of packets using in-band and out-of-band interfaces (a Linux-native LLDP such as `lldpad` is not used). The `lldp_mgr` crafts packets based on its configuration and forwards them using `xdp`. The SR Linux `xdp` process manages the sending and receiving of frames using in-band and out-of-band ports.

7.6 MACsec

The IEEE 802.1AE standardized Media Access Control Security (MACsec) is implemented in SR Linux to provides point-to-point security for almost all types of traffic on Ethernet links between directly connected nodes. For more information on the SR Linux implementation of MACsec, see the "MACsec" chapter in the *SR Linux Configuration Basics Guide*.

8 SR Linux routing functions

This chapter describes key elements of SR Linux routing functions: network instances, support for standard routing protocols including BGP, OSPF, and IS-IS, as well as ACLs, routing policies, and Quality of Service.

8.1 Network instances

On the SR Linux device, you can configure one or more virtual routing instances, known as network instances. Each network instance has its own interfaces, its own protocol instances, its own route table, and its own FIB.

When a packet arrives on a subinterface associated with a network instance, it is forwarded according to the FIB of that network instance. Transit packets are normally forwarded out another subinterface of the network instance.

SR Linux supports the following types of network instances:

- default
- ip-vrf
- mac-vrf
- vpws

The initial startup configuration for SR Linux has a single default network instance. By default, there are no ip-vrf or mac-vrf network instances; these must be created by explicit configuration. The ip-vrf network instances are the building blocks of Layer 3 IP VPN services, and mac-vrf network instances are the building blocks of EVPN services.

Within a network instance, you can configure BGP, OSPF, and IS-IS protocol options that apply only to that network instance. See the *SR Linux Configuration Basics Guide* for configuration information.

8.2 Static routes

Within a network instance, you can configure static routes. Each static route is associated with an IPv4 prefix or an IPv6 prefix, which represents the packet destinations matched by the static route. Each static route belongs to a specific network instance. Different network instances can have overlapping routes (static or otherwise) because each network instance installs its own routes into its own set of route tables and FIBs.

Each static route must be associated with a statically configured next-hop group, which determines how matching packets are handled: either perform a blackhole discard action or a forwarding action. The next-hop group can specify a list of one or more next-hops, each identified by an IPv4 or IPv6 address and a resolve flag. If the resolve flag is set to false, only a direct route can be used to resolve the IPv4 or IPv6 next-hop address; if the resolve flag is set to true, any route in the FIB can be used to resolve the IPv4 or IPv6 next-hop address.

Each static route has a specified metric and preference. The metric is the IGP cost to reach the destination. The preference specifies the relative degree this static route is preferred compared to other static and non-static routes available for the same IP prefix in the same network instance.

A static route is installed in the FIB for the network instance if the following conditions are met:

- The route has the lowest preference value among all routes (static and non-static) for the IP prefix.
- The route has the lowest metric value among all static routes for the IP prefix.

If BGP is running in a network instance, all static routes of that same network instance are automatically imported into the BGP local RIB, so that they can be redistributed as BGP routes, subject to BGP export policies.

See the "Network instances" chapter of the *SR Linux Configuration Basics Guide* for configuration information.

8.3 Aggregate routes

You can specify aggregate routes for a network instance. Each aggregate route is associated with an IPv4 prefix or an IPv6 prefix, which represents the packet destinations matched by the aggregate route. As with static routes, each aggregate route belongs to a specific network instance, though different network instances can have overlapping routes because each network instance installs its own routes into its own set of route tables and FIBs.

An aggregate route can become active when it has one or more contributing routes. A route contributes to an aggregate route if all of the following conditions are met:

- The prefix length of the contributing route is greater than the prefix length of the aggregate route.
- The prefix bits of the contributing route match the prefix bits of the aggregate route up to the prefix length of the aggregate route.
- There is no other aggregate route that has a longer prefix length that meets the previous two conditions.
- The contributing route is actively used for forwarding and is not an aggregate route itself.

That is, a route can only contribute to a single aggregate route, and that aggregate route cannot recursively contribute to a less-specific aggregate route.

Aggregate routes have a fixed preference value of 130. If there is no route to the aggregate route prefix with a numerically lower preference value, then the aggregate route, when activated by a contributing route, is installed into the FIB with a blackhole next-hop. It is not possible to install an aggregate route into the route table or as a BGP route without also installing it in the FIB.

The aggregate routes are commonly advertised by BGP or another routing protocol so that the individual contributing routes no longer need to be advertised. This can speed up routing convergence and reduce RIB and FIB sizes throughout the network. If BGP is running in a network instance, all active aggregate routes of that network instance are automatically imported into the BGP local RIB so they can be redistributed as BGP routes, subject to BGP export policies.

See the "Network instances" chapter of the *SR Linux Configuration Basics Guide* for configuration information.

8.4 BGP feature support

As with other functions on SR Linux, BGP operates as a modular application. The BGP manager application is responsible for running the BGP control plane. It subscribes to the IDB for configuration updates and listens for network instance and routing policy updates.

SR Linux supports the following BGP features:

Basic features

- Global AS configuration with local AS override per session. SR Linux always advertises 4-byte ASN capability.
- Local-address/source selection per neighbor
- EBGP and IBGP sessions
- Peer groups
- Configurable route table preference, with separate control for EBGP and IBGP routes
- Configurable TCP MSS per-neighbor
- EBGP split-horizon (always enabled)
- EBGP multihop
- BGP import and export policies
- IPv4/IPv6 default originate independent from default routes in the FIB
- AS path loop detection, with configurable threshold
- RFC 7606 error handling
- Tools commands to hard reset peer or soft reset with `ROUTE_REFRESH`
- Configurable trace options
- Routing policy actions to replace `AS_PATH`
- Options for handling the `AS_PATH` in received BGP routes
- BGP route reflection
- BGP AIGP metric
- Seamless MPLS with BGP-LU
- Segment routing with BGP-LU prefix SID
- BGP-LS
- BGP shortcuts
- BGP RT Constrained Route Advertisement

Failure detection features

- Session `KEEPALIVES`, with configurable hold-time and keepalive
- BGP next-hop tracking
- Fast failover (subinterface down)

- Bidirectional Forwarding Detection (BFD) for single-hop and multi-hop IPv4 and IPv6 sessions

Convergence features

- Configurable min-route-advertisement interval
- Rapid-withdrawal
- Option to wait for FIB install before advertising reachability
- Option to delay route advertisement until the address family has reached converged state (or timeout occurs)
- BGP-LU FRR

Graceful restart

- Helper/receiving router role
- Restarting router role during warm restart

Dynamic/unconfigured BGP neighbors

- Accept incoming sessions from allowed prefix/AS ranges

Address family support

- IPv4 unicast address family with IPv4 next-hops
- IPv4 unicast address family with IPv6 next-hops: requires MP_REACH_NLRI encoding and RFC 5549 capability advertisement
- IPv6 unicast address family with IPv6 next-hops
- Configurable limits on received routes per session per-AF > log when any threshold is exceeded

Multipath/ECMP

- Each BGP route supports multiple ECMP next-hops
- Two-level ECMP: Level 1 selects BGP path/next-hop, Level 2 selects a next-hop of the route that resolves the BGP next-hop
- Max Level 1 next-hops per route and max Level 2 next-hops per BGP next-hop are configurable per NLRI address family
- Weighted ECMP using Link Bandwidth community

See the *SR Linux Routing Protocols Guide* for BGP configuration examples. See the *SR Linux Advanced Configuration Guide* for a BGP underlay routing example.

8.5 IS-IS feature support

The SR Linux IS-IS manager application includes support for the following features:

- Level 1, Level 2, and Level 1/2 IS types
- Configurable Network Entity Title (NET) per IS-IS instance
- Support for IPv4/v6 routing

- ECMP support per destination
- IS-IS export policies (redistribution of other types of routes into IS-IS)
- Authentication of CSNP, PSNP, and IIH PDUs with authentication type and key, configurable per instance and per instance level. Authentication type and key for IIH PDUs are also configurable per interface and level.
- Support for authentication keychains
- Purge Originator ID TLV (RFC 6232)
- Options to ignore and suppress the attached bit
- Ability to set the overload bit immediately or to set the bit after each subsequent restart of the IS-IS manager application and leave it on for a configurable duration each time
- Control over the Link-State PDU (LSP) MTU size, with range from 490 bytes to 9490 bytes
- Configuration control over timers related to LSP lifetime, LSP refresh interval, SPF calculation triggers, and LSP generation
- Support for hello padding (strict, loose, and adaptive modes)
- Support for graceful restart, but only acting as a helper of the restarting router
- Level 1 to Level 2 route summarization
- BFD for fast failure detection
- Configurable hello timer/multiple per interface and level
- Support for wide metrics (configurable per level)
- Configurable route preference for each route type: Level 1-internal, Level 1-external, Level 2-internal and Level 2-external.
- Use of route policies to add/remove/replace one IS-IS route tag
- Multi-instance IS-IS (MI-ISIS), which allows multiple instances of IS-IS to operate on a single circuit
- Multi-Topology Intermediate System to Intermediate System (MT-ISIS) MT0 and MT2 support
- IS-IS support for Traffic Engineering TLVs and advertising interface delay
- Non-Stop Routing (NSR) IS-IS

See the *SR Linux Routing Protocols Guide* for an IS-IS configuration example.

8.6 OSPF feature support

SR Linux supports OSPFv2 and OSPFv3. The following features are supported:

- Reference bandwidth
- OSPF areas
- Stub areas
- Not So Stubby Areas (NSSAs)
- Passive interfaces
- Authentication
- Route policies

- Route redistribution to other protocols
- BFD for monitoring OSPF adjacencies
- Overload support and associated options
- Export policies (route redistribution from other protocols into OSPF, including ASBR support)
- Graceful restart

See the *SR Linux Routing Protocols Guide* for an OSPF configuration example.

8.7 Routing policies

Routing policies control the size and content of the routing tables, the routes that are advertised, and the best route to take to reach a destination. SR Linux routing policies allow for detailed control of IP routes learned and advertised by routing protocols such as BGP.

Each routing policy has a sequence of rules (called entries or statements) and a default action. Each statement has a numerical sequence identifier that determines its order relative to other statements in that policy. The statement supports both alphanumerical and numerical sequence identifiers. When a route is analyzed by a policy, it is evaluated by each statement in sequential order.

Each policy statement has zero or more match conditions and a base action (either accept or reject); the statement may also have route-modifying actions. A route matches a statement if it meets all of the specified match conditions.

The first statement that matches the route determines the actions that are applied to the route. If the route is not matched by any statements, the default action of the policy is applied. If there is no default action, then a protocol- and context-specific default action is applied.

See the "Routing policies" chapter in the *SR Linux Routing Protocols Guide* for a list of valid match conditions and policy actions, as well as configuration examples.

8.8 ECMP load balancing

Static, BGP, OSPF, and IS-IS routes to IPv4 and IPv6 destinations are programmed into the datapath by their respective applications, with multiple IP ECMP next-hops. SR Linux load-balances packets across these IP ECMP next-hops.

When an IPv4/IPv6 packet is received on a subinterface, and it matches a route with a number of ECMP hops, the next-hop used to forward the packet is determined from a hash calculation based on the packet type (IPv4/IPv6, TCP/UDP).

SR Linux attempts to keep packets in the same flow on the same network path while distributing traffic so that each of the N ECMP next-hops carries approximately $1/N$ th of the load.

On some platforms, SR Linux supports resilient hashing, which allows SR Linux to move as few flows as possible when removing or adding members to the ECMP set. Resilient hashing is particularly useful when the ECMP next-hops of an IP route correspond to network appliances or host servers that maintain state for the flows that they service, and moving flows would require state to be rebuilt.

8.9 Access Control Lists

An Access Control List (ACL) is an ordered set of rules that are evaluated on a packet-by-packet basis to determine whether access should be provided to a specific resource. ACLs can be used to drop unauthorized or suspicious packets from entering or leaving a routing device via specified interfaces.

SR Linux supports the following types of ACLs:

- Interface ACLs restrict the traffic allowed to pass through a specific set of subinterfaces. IPv4 and IPv6 filters can be applied to a subinterface to restrict IP traffic entering or exiting that subinterface, while MAC filters can match Ethernet frames carrying an IP or non-IP payload and can be applied to the input or output traffic of one or more subinterfaces.
- CPM-filter ACLs filter IPv4 or IPv6 traffic that is locally terminating on the router, regardless of the subinterface, port, or line card where it enters.
- Capture-filter ACLs filter all transit and terminating IPv4 or IPv6 traffic that is arriving on any subinterface of the router. Traffic matching a capture-filter ACL can be copied to the control plane for packet capture and analysis.
- System filter ACLs evaluate traffic early in the ingress pipeline, at a stage before tunnel termination occurs and before interface filters are run. For VXLAN traffic, system filters can match and drop unauthorized VXLAN tunnel packets before they are decapsulated, based on information in the outer header.

The rules of each ACL policy are evaluated in sequential order. Filter evaluation stops at the first matching entry where all match conditions are met, at which point the actions specified in the ACL are applied to the packet.

For information about configuring ACLs, see the *SR Linux ACL and Traffic Steering Guide*.

8.10 Quality of Service

SR Linux supports QoS policies for assigning traffic to forwarding classes or remarking traffic at egress before it leaves the router. Classifier policies map incoming packets to the appropriate forwarding classes, and rewrite-rule policies mark outgoing packets based on the forwarding class. Each received packet is classified to a forwarding class based on the packet type and subinterface configuration.

Supported QoS features and configuration are platform-dependent. See the *7730 SXR SR Linux Quality of Service Guide* and the *7220 IXR and 7250 IXR SR Linux Quality of Service Guide* for information about how transit and router-originated traffic is processed, and for configuration information.

8.11 MPLS feature support

Multiprotocol Label Switching (MPLS) provides the ability to set up connection-oriented paths over a connectionless IP network. MPLS facilitates network traffic flow and provides a mechanism to engineer network traffic patterns independently from routing tables.

Label Distribution Protocol (LDP) is a protocol used to distribute MPLS labels in non-traffic-engineered applications. LDP allows routers to establish label switched paths through a network by mapping network-layer routing information directly to data link layer-switched paths.

SR Linux supports the following MPLS and LDP functionality:

Tunnels

- SR-MPLS
 - LER and LSR functionality
 - Unnumbered subinterfaces
 - Flexible algorithm (Flex-Algo) support with delay normalization
- Configurable label range
 - Null label support, PHP use-cases
 - SR-ISIS with MT=2 support, LFA, ti-LFA
 - LER and LSR ECMP
- Entropy label based ECMP, LER, and LSR
 - Uncolored SR-MPLS TE-Policy with the following path computation options for both IPv4 and IPv6 next hops:
 - Explicit-path
 - PCE computed, PCC initialized
 - Local CSPF
 - sBFD based liveness tracking
 - B-SID support
 - Active/standby segment lists
 - FRR/LFA support
 - Tag-set (also known as, admin-tag) for steering services/shortcuts to a preferred segment list
 - Advertising link delay within IS-IS traffic engineering extensions
- Colored TE policy for SR-MPLS allows 32 programmed segment lists per candidate path, allowing ECMP of traffic across parallel segment lists.

LDP

- LDPv4 implementation compliant with RFC 5036
- LDP support in the default network-instance only
- Label distribution using DU (downstream unsolicited), ordered control
- Platform label space only
- Configurable label range (dynamic, non-shared label-block)
- Support for overload TLV when label-block has no free entries
- Configurable timers (hello-interval, hello-holdtime, keepalive-interval)
- Ingress LER, transit LSR, and egress LER roles for /32 IPv4 FECs
- Automatic FEC origination of the system0.0 /32 IPv4 address prefix
- /32 IPv4 FEC resolution using IGP routes, with longest prefix match option
- ECMP support with configurable next-hop limit (up to 64)

- Automatic installation of all LDP /32 IPv4 prefix FECs into TTM
- Per-peer configurable FEC limit
- Graceful restart helper capability
- BGP shortcuts for IPv4 traffic
- Protocol debug/trace-options
- LDP-IGP synchronization
- Advertise address mapping message for primary IPv4 address of the adjacent interface only.
- Non-configurable capability advertisement in INITIALIZATION messages only claiming support for:
 - State advertisement control (SAC) with interest in IPv4 prefix FECs only
 - Fault tolerance (Graceful Restart)
 - Nokia overload TLV
 - Unrecognized notification
- Split-horizon support: A label-mapping message is not advertised to a peer if the FEC matches an address sent by that peer in an Address Mapping message.
- Remote LFA including auto tLDP session creation to P,Q node
- IPv4 and IPv6 next hops
- BFD support
 - Configurable local-lsr-id for both I-LDP and t-LDP
 - FEC-originate
 - Per peer import and export policies
 - ECMP for LER and LSR

See the *SR Linux MPLS Guide* for configuration information.

8.12 Segment routing with IPv6 data plane (SRv6)

SRv6 is an IPv6-based segment routing approach that introduces the Segment Routing Header (SRH) as a new IPv6 extension header. SRv6 provides more than just IPv6 transport with shortest-path and source-routing capabilities; it provides a framework for IPv6 network programmability that leverages the large IPv6 address space.

For more information, see the "Segment routing with IPv6 data plane (SRv6)" chapter of the *SR Linux Segment Routing Guide*.

8.13 Multicast

IP multicast allows for one-to-many and many-to-many communication, allowing data sources to transmit simultaneously to multiple destinations.

SR Linux supports Internet Group Management Protocol (IGMP), Multicast Listener Discovery Protocol (MLDP), Protocol Independent Multicast (Sparse Mode) (PIM-SM), and Multicast Source Discovery Protocol (MSDP) protocols for managing multicast traffic.

See *SR Linux Multicast Routing Guide* for more information.

9 SR Linux services

SR Linux services facilitate EVPN-VXLAN and EVPN-MPLS deployments in data centers and service provider networks. Ethernet Virtual Private Network (EVPN) is a technology that allows Layer 2 and Layer 3 traffic to be tunneled across an IP network.

The SR Linux EVPN solution supports the following features:

- EVPN for VXLAN/MPLS tunnels (Layer 2), extending a BD in overlay multi-tenant DCs and service provider networks
- EVPN for VXLAN/MPLS tunnels (Layer 3), allowing inter-subnet-forwarding for unicast traffic within the same tenant infrastructure

These features are summarized in the following sections. See the *SR Linux VPN Services Guide* for descriptions of supported features and configuration examples.

9.1 Layer 2 services


Layer 2 services refers to the infrastructure implemented on SR Linux to support multiple virtual switches on the same system.

To do this, SR Linux uses a network instance of type **mac-vrf**, which functions as a broadcast domain. Each **mac-vrf** network instance builds a bridge table composed of MAC addresses that can be learned via the data path on network instance interfaces or via static configuration. You can configure the size of the bridge table for each **mac-vrf** network instance, as well as the aging for dynamically learned MAC addresses and other parameters related to the bridge table.

The **mac-vrf** network instance is associated with a network instance of type **default** or **ip-vrf** via an Integrated Routing and Bridging (IRB) interface. IRB interfaces enable inter-subnet forwarding.

[Figure 6: MAC-VRF, IRB interface, and IP-VRF](#) shows the relationship between an IRB interface and **mac-vrf**, and **ip-vrf** network instance types.

Figure 6: MAC-VRF, IRB interface, and IP-VRF



sw4397

SR Linux also supports **vpws** network instances for point-to-point or ELINE services.

See the *SR Linux VPN Services Guide* for a description of Layer 2 services components and configuration examples.

9.2 EVPN for Layer 2

EVPN for Layer 2 allows for the extension of a broadcast domain. To support this topology, SR Linux includes the following features:

- Bridged subinterface extensions, including a default subinterface that captures untagged and non-explicitly configured VLAN-tagged frames on tagged subinterfaces
- EVPN control and data plane extensions as described in RFC 8365 and RFC 7432
- Distributed security and protection
- EVPN L2 multi-homing, including the ES model definition for all-active and single-active multi-homing

SR Linux also supports static and TLDP-signaled pseudowires, which can be used on MAC-VRF and VPWS instances.

See the *SR Linux VPN Services Guide* for a description of supported features, basic configuration information, and EVPN L2 multihoming configuration examples.

9.3 EVPN for Layer 3

SR Linux supports EVPN for Layer 3 for inter-subnet-forwarding for unicast traffic within the same tenant infrastructure. SR Linux features that support this topology fall into the following categories:

- EVPN L3 control plane (EVPN IP prefix routes, RT5s, MAC/IP Advertisement routes, or RT2s) and data plane as described in RFC 9135 and RFC 9136
- EVPN L3 multi-homing on MAC-VRFs with IRB interfaces that use anycast GW IP and MAC addresses in all leafs attached to the same BD
- Host route mobility procedures to allow fast mobility of hosts between leaf nodes attached to the same BD

Other supported features include:

- Interface-less (IFL) model interoperability with unnumbered interface-ful (IFF) model
- ECMP over EVPN, including unequal ECMP, IP aliasing, and combined ECMP
- Support for interface-level OAM (ping) in anycast deployments
- EVPN interoperability with VLAN-aware bundle services

See the *SR Linux VPN Services Guide* for EVPN Layer 3 basic configuration information and examples.

9.4 IP-VPN services

IP-VPN services use a combination of MP-BGP and MPLS to distribute IPv4/v6 routing information and provide Layer 3 VPN services.

Each IP-VPN consists of a set of customer end points connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each IP-VPN instance. The PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the IP-VPN association and handles any potential IP address overlap.

Multi-Protocol BGP (MP-BGP) is used to exchange the routes of a particular VPN among the PE routers that are attached to that VPN.

See the *SR Linux VPN Services Guide* for IP-VPN configuration information and examples.

10 SR Linux OAM and diagnostics tools

This chapter describes the Operations, Administration, and Maintenance (OAM) and diagnostic tools available on SR Linux, including BFD, sFlow, traffic monitoring, packet-tracing, and IP performance measurement.

10.1 BFD support

Bidirectional Forwarding Detection (BFD) is a lightweight mechanism used to monitor the liveness of a remote neighbor. Because of this lightweight nature, BFD can send and receive messages at a much higher rate than other control plane hello mechanisms. This attribute allows connection failures to be detected faster than other hello mechanisms.

SR Linux supports BFD asynchronous mode, where BFD control packets are sent between two systems to activate and maintain BFD neighbor sessions between them.

BFD can be configured to monitor connectivity for the following:

- BGP peers
- Next-hops for static routes
- OSPF adjacencies
- IS-IS adjacencies

See "Bidirectional Forwarding Detection" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

Micro-BFD, where BFD sessions are established for individual members of a LAG, is also supported. If the BFD session for one of the links indicates a connection failure, the link is taken out of service from the perspective of the LAG.

See "Micro-BFD" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

SR Linux supports seamless bidirectional forwarding detection (S-BFD), which is a simplified mechanism that speeds up a BFD session by eliminating the negotiation and state establishment process. This is accomplished primarily by predetermining the session discriminator and using specific mechanisms to distribute the discriminators to a remote network entity. This allows client applications or protocols to quickly initiate and perform connectivity tests.

See "Seamless bidirectional forwarding detection (S-BFD)" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.2 sFlow support

SR Linux supports sFlow version 5 behavior and formats. sFlow is used to monitor data traffic flows traversing different points in a network. The sFlow functionality uses an sFlow agent and an sFlow collector. The agent is software that runs on a network element and samples and reports flow headers and

statistics. The collector is software that typically runs on a remote server and receives the flow headers and statistics from one or more sFlow agents.

On the SR Linux device, sFlow samples flow data and reports the samples to configured sFlow collectors. Up to eight sFlow collectors can be configured. When sFlow is enabled on an interface, the sFlow agent streams interface statistics to the configured sFlow collectors.

See the "sFlow" chapter in the *SR Linux OAM and Diagnostics Guide* for configuration information and examples.

10.3 Interactive traffic monitoring tool

SR Linux features an interactive traffic monitoring tool that samples packets entering the system on any interface matching a set of parameters, and streams the header details either to the current login session or to a specified output file.

When the traffic monitoring tool is activated, mirroring policies are dynamically populated on all ingress ports, and matching packets are sent to the CPM for display. Header information for the matching packets is displayed in either tcpdump format or hex format, depending on the options chosen.

When the traffic monitoring tool is deactivated, the mirroring policies are automatically removed from all ingress interfaces.

See the "Interactive traffic monitoring" chapter of the *SR Linux Troubleshooting Toolkit Guide* for usage information.

10.4 Packet-trace tool

SR Linux includes a packet-trace tool that reports the forwarding behavior of a probe packet. The probe packet is injected into a specified interface forwarding context, and the packet-trace tool records the forwarding destination or egress port for the probe packet, as well as any matched ACL records or reasons for discarding the packet. The probe packet can be specified in Scapy format, base64 format, or pcap file format. See the "Packet-trace tool" chapter of the *SR Linux Troubleshooting Toolkit Guide* for usage information.

10.5 Traffic mirroring to remote and local destinations

Traffic mirroring sends copies of IPv4 and IPv6 packets from a designated source to a designated destination.

The source for the mirrored traffic can be an interface (port), a subinterface (VLAN), a LAG, or an ACL filter. The mirrored traffic can be copied to a local destination, such as a locally-attached traffic analyser (local mirroring), or encapsulated into a tunnel toward a remote destination (remote mirroring).

See the "Mirroring" chapter of the *SR Linux OAM and Diagnostics Guide* for usage information and configuration examples.

10.6 TWAMP

Two-Way Active Measurement Protocol (TWAMP) is a standards-based method to measure the IP performance between two devices including packet loss, delay, and jitter. TWAMP leverages the methodology and architecture of One-Way Active Measurement Protocol (OWAMP) to define a method to measure two-way or round-trip metrics.

There are four logical entities in TWAMP: the control client, the session sender, the server, and the session reflector. The control client and session sender are typically implemented in one physical device and the server and session reflector in a second physical device. SR Linux acts as the server and the session reflector.

See "TWAMP" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.7 STAMP

The Simple Two-Way Active Measurement Protocol (STAMP) defined in RFC 8762 is a standards-based method to measure the IP performance without the use of a control channel to pre-signal session parameters.

The PDU structure allows for the collection of frame delay, frame delay range, inter-frame delay variation, and frame loss ratio. The RFC 8972 STAMP Optional Extensions specification maintains the existing structure of the STAMP PDU but redefines existing fields and adds the capability to include TLVs. SR Linux supports RFC 8762 and the structural changes with TLV processing in the options draft RFC 8972.

For each routed network instance, the STAMP session sender transmits STAMP test packets to the destination UDP port of the session reflector. The session reflector receives the packets, processes the STAMP test packet, and sends them back to the session sender. The session sender receives the reflected packets and uses the timestamps and sequence numbers to calculate delay and loss performance metrics. The session reflector supports a prefix list which filters based on IPv4 or IPv6 addressing. The reflector is stateful and uses the tuple SIP, DIP, SP, DP, and SSID to identify individual STAMP test sessions.

See "STAMP" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.8 Ethernet connectivity fault management

Ethernet Connectivity Fault Management (ETH-CFM) is a set of protocols designed to monitor, detect, and troubleshoot issues within Ethernet networks. Defined by standards IEEE 802.1ag and ITU-T Y.1731, ETH-CFM operates at the Ethernet layer to ensure network reliability and fault isolation.

It provides several key functionalities, including:

- maintenance domains (MD) that define the scope of fault detection
- maintenance associations (MA) that group network endpoints
- maintenance endpoints (MEPs) that serve as active points for initiating and terminating fault management tasks.

ETH-CFM tools like continuity check messages (ETH-CCM) are used for continuous monitoring, while Loopback (ETH-LBM) and Linktrace (ETH-LTM) messages help in verifying connectivity and tracing

the path between endpoints. Additionally, the Remote Defect Indication (ETH-RDI) alerts MEPs of any detected faults. Together, these tools provide comprehensive fault detection, isolation, and performance monitoring in Ethernet networks.

See "Ethernet OAM tools and protocols" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.9 Link measurement

A link measurement test is a method for measuring and reporting delay information for directly connected IP peers. Link measurement uses the STAMP protocol defined in RFC 8762 to measure delay for an IP interface. The feature uses reporting options and thresholds to determine the values to be reported. This allows for the network element to advertise delay metrics for the IP interface using the IGP. The unidirectional delay values can be used to influence delay sensitive paths through the network.

See "Link measurement" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.10 Delay and loss measurement

Performance monitoring encompasses a variety of tools and protocols designed to measure, report, analyze, and optimize network performance, allowing network administrators to detect and resolve issues proactively. SR Linux supports performance monitoring using STAMP for delay and packet loss measurements in IP networks.

The STAMP OAM performance monitoring architecture for delay and loss measurement consists of the following foundational components:

- session: This is the overall collection of the test parameters, measurement intervals, thresholds and storage of results. It is the overall container that defines the attributes of the session.
- standard performance monitoring packets: SR Linux supports STAMP to measure performance metrics such as delay and packet loss.
- measurement interval: These are time-based non-overlapping windows that capture results that are received in that window of time.
- data structures: These are the unique counters and measurement results that represent the specific protocol.
- bin group: These are ranges in microseconds that count the results that fit into the range.

See "Performance monitoring" in the *SR Linux OAM and Diagnostics Guide* for configuration information.

10.11 Uncolored SR-MPLS TE-Policy

Uncolored SR-MPLS TE policies define the rules for traffic engineering in an SR-enabled network. They determine how traffic should be routed across the network based on various constraints.

SR Linux supports several path computation methods, which can be configured individually for each TE policy segment list, including explicit paths, local Constrained Shortest Path First (CSPF), and Path Computation Element Protocol (PCEP) methods.

SR Linux supports tags (analogous to SR-TE LSP admin-tag) for SR-MPLS TE-Policy. A tag can be used to dictate service to TE-Policy binding based on user-defined constraints.

SR Linux supports the collection of LSP statistics. LSP statistics offer real-time utilization information, allowing for efficient optimization of network paths during congestion.

Non-Stop Routing (NSR) is a high-availability feature that enables a router to continue forwarding traffic without disruption during a control plane switchover.

For more information, see the "Uncolored SR-MPLS TE-Policy" chapter of the *SR Linux Segment Routing Guide*.

10.12 Colored TE policy

A Colored SR-MPLS TE policy describes a source-routed path from a head-end router to an endpoint, and it determines which traffic flows should be steered through that path; the colored attribute allows up to 32 programmed segment lists per candidate path, allowing ECMP of traffic across parallel segment lists.

SR Linux supports the collection of segment list (LSP) statistics for colored TE Policies.

With Non-Stop Routing (NSR) enabled, the TE-Policy segment lists and candidate paths in colored TE-Policy remain intact in case of any faults in the Traffic Engineering Manager process.

For more information, see the "Colored TE policy" chapter of the *SR Linux Segment Routing Guide*.

10.13 Service Activation Testhead (SAT)



Note: This feature is supported on 7250 IXR-6e, 7250 IXR-10e, 7250 IXR-X1b, and 7250 IXR-X3b platforms.

SAT is a point-to-point out-of-service test methodology defined in the ITU-T Y.1564 standard. It is used to validate Ethernet-based service configuration and performance before the service is delivered to a customer. The test ensures that the network can support the service in accordance with the agreed-upon service level agreements (SLAs).

ITU-T Y.1564 defines two classes of tests:

- Service configuration tests: These are short-duration tests that verify the accuracy of service configuration. They include:
 - committed information rate (CIR)
 - CIR combined with peak information rate (CIR-PIR)
 - policing (traffic is tested at 125% of the configured PIR)
- Service performance test: This typically long-duration test evaluates key performance metrics of the service.

The SAT tool administers tests based on configured transmission rates and measures performance metrics including:

- throughput
- frame loss ratio (FLR)

- frame delay (FD), minimum, maximum, and average
- frame delay variation (FDV), minimum, maximum, and average

Additional statistics include the number of transmitted and received packets, start and end times, and remaining test duration. Statistics are available during the execution of the test.

See "Service Activation Testhead (SAT)" chapter in the *SR Linux OAM and Diagnostics Guide* for detailed information.

10.14 IPFIX

IPFIX enables SR Linux to provide detailed visibility into network traffic by exporting flow-based telemetry information to external analytics and monitoring systems. Based on the IETF IPFIX (Version 10) standard, the feature captures traffic flows observed on network interfaces, aggregates packet statistics, and exports structured flow records using a flexible, template-driven format.

By transforming packet-level activity into summarized flow data, IPFIX allows operators to monitor application behavior, analyze traffic patterns, detect anomalies, and support capacity planning and billing use cases without impacting forwarding performance.

IPFIX on SR Linux supports scalable flow observation through configurable sampling, efficient flow caching, and standards-based export mechanisms. Flow records include key attributes such as source and destination addresses, ports, protocols, timestamps, and traffic counters, enabling deep operational and security insights across data center, edge, and service provider networks.

See "IPFIX" chapter in the *SR Linux OAM and Diagnostics Guide* for detailed information.

11 Standards and protocol support

**Note:**

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

11.1 Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

11.2 Border Gateway Protocol (BGP)

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*

RFC 4724, *Graceful Restart Mechanism for BGP – helper mode*

RFC 4760, *Multiprotocol Extensions for BGP-4*

RFC 4893, *BGP Support for Four-octet AS Number Space*

RFC 5492, *Capabilities Advertisement with BGP-4*

RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*

RFC 7606, *Revised Error Handling for BGP UPDATE Messages*

RFC 7705, *Autonomous System Migration Mechanisms and Their Effects on the BGP AS_PATH Attribute*

RFC 8212, *Default External BGP (EBGP) Route Propagation Behavior without Policies*

11.3 Bridging and management

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1AX, *Link Aggregation*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1Qbb, *Priority-based Flow Control*

11.4 Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ip-aliasing-04, *EVPN Support for L3 Fast Convergence and Aliasing/Backup Path*

draft-ietf-bess-evpn-ipvpn-interworking-16, *Interconnecting EVPN and IPVPN Domains*

RFC 8214, *Virtual Private Wire Service Support in Ethernet VPN (EVPN)*

RFC 8365, *A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)*

RFC 8560, *Seamless Integration of Ethernet VPN (EVPN) with Virtual Private LAN Service (VPLS)*

RFC 9014, *Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks*

RFC 9047, *Propagation of ARP/ND Flags in an Ethernet Virtual Private Network (EVPN)*

RFC 9135, *Integrated Routing and Bridging in Ethernet VPN (EVPN) – Asymmetric and Symmetric IRB Procedures, Mobility Procedure*

RFC 9136, *IP Prefix Advertisement in Ethernet VPN (EVPN)*

RFC 9161, *Operational Aspects of Proxy ARP/ND in Ethernet Virtual Private Networks*

RFC 9251, *Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)*

RFC 9625, *Optimized Inter-Subnet Multicast forwarding for Ingress Replication*

11.5 gRPC Remote Procedure Calls (gRPC)

authz.proto reference 0.3.0, *gNSI Authz Service*

certz.proto reference 0.4.0, *gNSI Certz Service*

credentialz.proto reference 0.4.0, *gNSI Credentialz Service*

factoryreset.proto version 0.1.0, *gNOI Factory Reset Service*

file.proto version 0.1.0, *gNOI File Service*

gnmi.proto version 0.10.0, *gNMI Service Specification*

gnmi_ext.proto, *gNMI Commit Confirmed Extension*

gnmi_ext.proto, *gNMI Config Subscription Extension*

gnmi_ext.proto, *gNMI Depth Extension*

gribi.proto version 1.0.0, *gRIBI Service Specification*
healthz.proto version 1.3.0, *gNOI Healthz Service*
os.proto version 0.1.1, *gNOI OS Service*
packet_link_qualification.proto version 1.1.0, *gNOI Packet Link Qualification Service*
system.proto version 1.2.0, *gNOI System Service*
PROTOCOL-HTTP2, *gRPC over HTTP2*

11.6 Intermediate System to Intermediate System (IS-IS)

ISO/IEC 10589:2002 Second Edition, *Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS – MT ID #0 and MT ID #2*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering – except IP Reachability TLV 135*

RFC 5306, *Restart Signaling for IS-IS*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 5316, *ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering – IPv4 TE Router ID Sub-TLV 11 and IPv6 TE Router ID sub-TLV 12*

RFC 6119, *IPv6 Traffic Engineering in IS-IS – IPv6 TE Router ID TLV 140*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 7981, *IS-IS Extensions for Advertising Router Information*

RFC 8202, *IS-IS Multi-Instance – single topology*

RFC 8570, *IS-IS Traffic Engineering (TE) Metric Extensions – Min/Max Unidirectional Link Delay Sub-TLV 34*

RFC 8919, *IS-IS Application-Specific Link Attributes*

RFC 9885, *Multi-Part TLVs in IS-IS*

11.7 Internet Protocol (IP) general

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 2131, *Dynamic Host Configuration Protocol* – static address allocation

RFC 2132, *DHCP Options and BOOTP Vendor Extensions* – sections 3.1, 3.2, 3.3, 3.5, 3.8, 3.14, 3.17, 9.2, 9.6

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 3046, *DHCP Relay Agent Information Option*

RFC 3162, *RADIUS and IPv6*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*

RFC 4252, *The Secure Shell (SSH) Authentication Protocol*

RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*

RFC 4254, *The Secure Shell (SSH) Connection Protocol*

RFC 4649, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option*

RFC 5424, *The Syslog Protocol*

RFC 8415, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)* – sections 17, 17.2, 17.2.1, 17.2.2, 17.2.3, 18.2, 22.1, 22.2, 22.3, 22.4, 22.6, 22.7, 22.13

RFC 8907, *The Terminal Access Controller Access-Control System Plus (TACACS+) Protocol*

11.8 Internet Protocol (IP) version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 1191, *Path MTU Discovery* – router specification

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1812, *Requirements for IPv4 Routers*

RFC 5227, *IPv4 Address Conflict Detection*

11.9 Internet Protocol (IP) version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 3587, *IPv6 Global Unicast Address Format*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration – router functions*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

11.10 Label Distribution Protocol (LDP)

RFC 3037, *LDP Applicability*
RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol – helper mode*
RFC 5036, *LDP Specification*
RFC 5443, *LDP IGP Synchronization*
RFC 5561, *LDP Capabilities*
RFC 5919, *Signaling LDP Label Advertisement Completion*
RFC 7552, *Updates to LDP for IPv6*

11.11 Multiprotocol Label Switching (MPLS)

RFC 3031, *Multiprotocol Label Switching Architecture*
RFC 3032, *MPLS Label Stack Encoding*
RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services – E-LSP*
RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*
RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
RFC 5036, *LDP Specification*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

11.12 Open Shortest Path First (OSPF)

RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart – helper mode*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 5187, *OSPFv3 Graceful Restart – helper mode*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

11.13 Path Computation Element Protocol (PCEP)

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*
RFC 8231, *Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE*
RFC 8408, *Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages*
RFC 8664, *Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing*

11.14 Pseudowire (PW)

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*

11.15 Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*
RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
RFC 2597, *Assured Forwarding PHB Group*
RFC 3140, *Per Hop Behavior Identification Codes*
RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*

11.16 Segment Routing (SR)

RFC 8402, Segment Routing Architecture
RFC 8660, Segment Routing with the MPLS Data Plane
RFC 8667, *IS-IS Extensions for Segment Routing*
RFC 9350, *IGP Flexible Algorithm* – IS-IS

11.17 Simple Network Management Protocol (SNMP)

draft-blumenthal-aes-usm-04, *The AES Cipher Algorithm in the SNMP's User-based Security Model* – CFB128-AES-192 and CFB128-AES-256
IANAifType-MIB revision 202004020000Z, *ianaifType*
IANA-RTPROTO-MIB revision 201604250000Z, *ianaRtProtoMIB*
IEEE8023-LAG-MIB revision 200006270000Z, *lagMIB* – selected OIDs
LLDP-MIB revision 200505060000Z, *lldpMIB* – selected OIDs
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 2578, *Structure of Management Information Version 2 (SMIPv2)*
RFC 2579, *Textual Conventions for SMIPv2*
RFC 2790, *Host Resources MIB* – selected OIDs
RFC 2856, *Textual Conventions for Additional High Capacity Data Types*
RFC 2863, *The Interfaces Group MIB* – selected OIDs
RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*
RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
RFC 3413, *Simple Network Management Protocol (SNMP) Applications*
RFC 3414, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)* – selected OIDs
RFC 3430, *Simple Network Management Protocol (SNMP) over Transmission Control Protocol (TCP) Transport Mapping*
RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types* – selected OIDs
RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*
RFC 4001, *Textual Conventions for Internet Network Addresses*
RFC 4293, *Management Information Base for the Internet Protocol (IP)* – selected OIDs

RFC 7630, *HMAC-SHA-2 Authentication Protocols in the User-based Security Model (USM) for SNMPv3*

11.18 Timing

GR-1244-CORE Issue 3, *Clocks for the Synchronized Network: Common Generic Criteria*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*

ITU-T G.811, *Timing characteristics of primary reference clocks*

ITU-T G.8261, *Timing and synchronization aspects in packet networks*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*

ITU-T G.8262.1, *Timing characteristics of an enhanced synchronous Ethernet equipment slave clock (eEEC)*

ITU-T G.8264, *Distribution of timing information through packet networks*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*

ITU-T G.8275.2, *Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network*

RFC 3339, *Date and Time on the Internet: Timestamps*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

11.19 Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) – server, unauthenticated mode*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

RFC 8545, *Well-Known Port Assignments for the One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP) – TWAMP*

RFC 8762, *Simple Two-Way Active Measurement Protocol – unauthenticated*

RFC 8972, *Simple Two-Way Active Measurement Protocol Optional Extensions – unauthenticated*

RFC 9503, *Simple Two-Way Active Measurement Protocol (STAMP) Extensions for Segment Routing Networks – excluding Sections 3, 4.1.2 and 4.1.3*

11.20 Virtual Private LAN Service (VPLS)

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

11.21 Yet Another Next Generation (YANG)

RFC 6991, *Common YANG Data Types*

RFC 7950, *The YANG 1.1 Data Modeling Language*

RFC 7951, *JSON Encoding of Data Modeled with YANG*

11.22 Yet Another Next Generation (YANG) OpenConfig Modules

openconfig-aaa.yang version 1.0.0, *OpenConfig AAA Model*

openconfig-acl.yang version 1.3.3, *OpenConfig ACL Model*

openconfig-aft-network-instance.yang version 0.3.1, *OpenConfig AFT Network Instance Model*

openconfig-aft-summary.yang version 0.2.0, *OpenConfig AFT Summary Model*

openconfig-bgp-policy.yang version 8.1.0, *OpenConfig BGP Policy Model*

openconfig-defined-sets.yang version 1.0.0, *OpenConfig Defined Sets Model*

openconfig-flexalgo.yang version 0.1.0, *OpenConfig Flexible Algorithms Model*

openconfig-gnsi-acctz.yang version 0.4.0, *OpenConfig Acctz Model*

openconfig-gnsi-authz.yang version 0.4.0, *OpenConfig Authz Model*

openconfig-gnsi-certz.yang version 0.6.0, *OpenConfig Certz Model*

openconfig-gnsi-credentialz.yang version 0.7.0, *OpenConfig Credentialz Model*

openconfig-gnsi-pathz.yang version 0.3.0, *OpenConfig Pathz Model*

openconfig-gribi.yang version 0.1.0, *OpenConfig gRIBI Model*

openconfig-if-aggregate.yang version 2.4.5, *OpenConfig Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.15.0, *OpenConfig Interfaces Ethernet Model*

openconfig-if-ethernet-ext.yang version 0.1.1, *OpenConfig Interfaces Ethernet Extensions Model*

openconfig-if-ip.yang version 3.5.1, *OpenConfig Interfaces IP Model*

openconfig-if-sdn-ext.yang version 0.2.0, *OpenConfig Interfaces SDN Extensions Model*

openconfig-interfaces.yang version 3.8.0, *OpenConfig Interfaces Model*

openconfig-isis-policy.yang version 0.8.0, *OpenConfig IS-IS Policy Model*

openconfig-lacp.yang version 2.1.0, *OpenConfig LACP Model*

openconfig-lldp.yang version 0.2.1, *OpenConfig LLDP Model*

openconfig-local-routing-network-instance.yang version 1.1.0, *OpenConfig Local Routing Network Instance Model*

openconfig-metadata.yang version 0.1.0, *OpenConfig Metadata Model*

openconfig-network-instance.yang version 4.6.0, *OpenConfig Network Instance Model*

openconfig-p4rt.yang version 1.0.0, *OpenConfig P4Runtime Model*

openconfig-platform.yang version 0.31.0, *OpenConfig Platform Model*

openconfig-platform-controller-card.yang version 0.2.0, *OpenConfig Platform Controller Card Model*
openconfig-platform-cpu.yang version 0.1.1, *OpenConfig Platform CPU Model*
openconfig-platform-fabric.yang version 0.1.0, *OpenConfig Platform Fabric Model*
openconfig-platform-fan.yang version 0.1.1, *OpenConfig Platform Fan Model*
openconfig-platform-healthz.yang version 0.1.1, *OpenConfig Platform Health Model*
openconfig-platform-integrated-circuit.yang version 0.3.1, *OpenConfig Platform Linecard Integrated Circuit Model*
openconfig-platform-linecard.yang version 1.2.0, *OpenConfig Platform Linecard Model*
openconfig-platform-pipeline-counters.yang version 0.5.1, *OpenConfig Platform Pipeline Counters Model*
openconfig-platform-port.yang version 1.0.1, *OpenConfig Port Model*
openconfig-platform-psu.yang version 0.2.1, *OpenConfig Platform PSU Model*
openconfig-platform-software.yang version 0.1.1, *OpenConfig Platform Software Model*
openconfig-platform-storage.yang version 0.1.0, *OpenConfig Platform Storage Model*
openconfig-platform-transceiver.yang version 0.16.0, *OpenConfig Transceiver Model*
openconfig-qos.yang version 0.11.2, *OpenConfig QoS Model*
openconfig-routing-policy.yang version 3.5.0, *OpenConfig Routing Policy Model*
openconfig-sampling.yang version 0.1.0, *OpenConfig Traffic Sampling Model*
openconfig-sampling-sflow.yang version 1.1.0, *OpenConfig Traffic Sampling sFlow Model*
openconfig-system.yang version 2.3.0, *OpenConfig System Model*
openconfig-system-controlplane.yang version 0.2.0, *OpenConfig System Controlplane Model*
openconfig-system-grpc.yang version 1.1.0, *OpenConfig System gRPC Model*
openconfig-system-utilization.yang version 0.1.0, *OpenConfig System Utilization Model*
openconfig-terminal-device.yang version 1.9.2, *OpenConfig Terminal Device Model*
openconfig-vlan.yang version 3.2.2, *OpenConfig VLAN Model*

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)