

Nokia Service Router Linux

Release 21

Nokia Service Router Linux (SR Linux) is an open, extensible and resilient network operating system (NOS) designed to enable superior scalability, flexibility and efficiency in webscale, service provider, enterprise data center and cloud environments.

Overview

Evolving application needs and the increasing adoption of distributed data centers is driving the need to rethink how data center networks are designed, deployed and operated. Nokia SR Linux was designed to solve challenges in modern data center networks, where the primary challenges are lack of scalability, inflexibility and the need for operational simplification.

Nokia SR Linux is a truly open and modular NOS. It implements a ground-up, model-driven architecture combined with field-proven routing protocol stacks to create a unique foundation that delivers superior openness, extensibility and resiliency.

This innovative design foundation supports superior telemetry, unrivaled flexibility for implementing and customizing network tools and third-party applications, and plug-and-play integration—all critical features for modern data center networking.

SR Linux is a key component of the Nokia Data Center Fabric solution, which also includes the Nokia Fabric Services System and the Nokia Data Center platforms.

Architecture and building blocks

SR Linux was developed in close collaboration with some of the world's largest data center operators. Its design choices were motivated by the need to provide high levels of openness, programmability and flexibility to meet growing DevOps and agility requirements. It allows operators to scale their data center networks while retaining flexibility and simplicity of operations.

Linux-based NOS

Openness is best achieved with an underlying open operating system such as Linux®. Nokia SR Linux is a Linux-based NOS that builds on the Linux kernel to provide a set of loosely coupled services that come together to provide the functional blocks of a NOS, along with simple interfaces for their consumption.

SR Linux uses an unmodified Linux kernel as the foundation on which to build a suite of network applications (Nokia applications as well as customer applications). This provides many benefits, including reliability, portability and ease of application development. Using an unmodified kernel speeds the delivery of security patches to critical system components, and leveraging a common Linux distribution provides access to equally field-proven management and infrastructure applications.

Ground-up, model-driven design

Openness cannot be an afterthought and must be designed from the foundation.

SR Linux implements a truly open architecture built around model-driven management and modern interfaces for its consumption (see Figure 1).

Our core design approach opens up the infrastructure by adopting emerging technologies, to allow services to expose their functionality in the form of generalized Remote Procedure Call (gRPC) services and protocol buffers (protobufs). This gives network applications the ability to simplify their internal schemas based on data modeling languages, and then expose these schemas directly for consumption by northbound interfaces: any object, any interface.

Modular, state-sharing and extensible architecture

Nokia SR Linux applications share state with each other via a publish/subscribe (pub/sub) architecture (see Figure 2). The Nokia pub/sub architecture is implemented using protobufs, gRPC and the Nokia Impart Database (IDB).

gRPC provides an efficient, secure communication channel for applications and allows native extensions through third-party applications.

IDB is a lightweight database that is optimized for handling high volumes of messages while protecting against any one application slowing down the whole system. IDB provides a reliable and scalable delivery mechanism, guaranteeing delivery of updates to subscribers.

SR Linux supports a variety of data center networking chipsets through the Nokia eXtensible Data Path (XDP). The Nokia XDP serves as a hardware abstraction layer that speeds time-to-market for multiple data center networking chipsets.

Figure 1. Nokia SR Linux ground-up, model-driven design for true openness

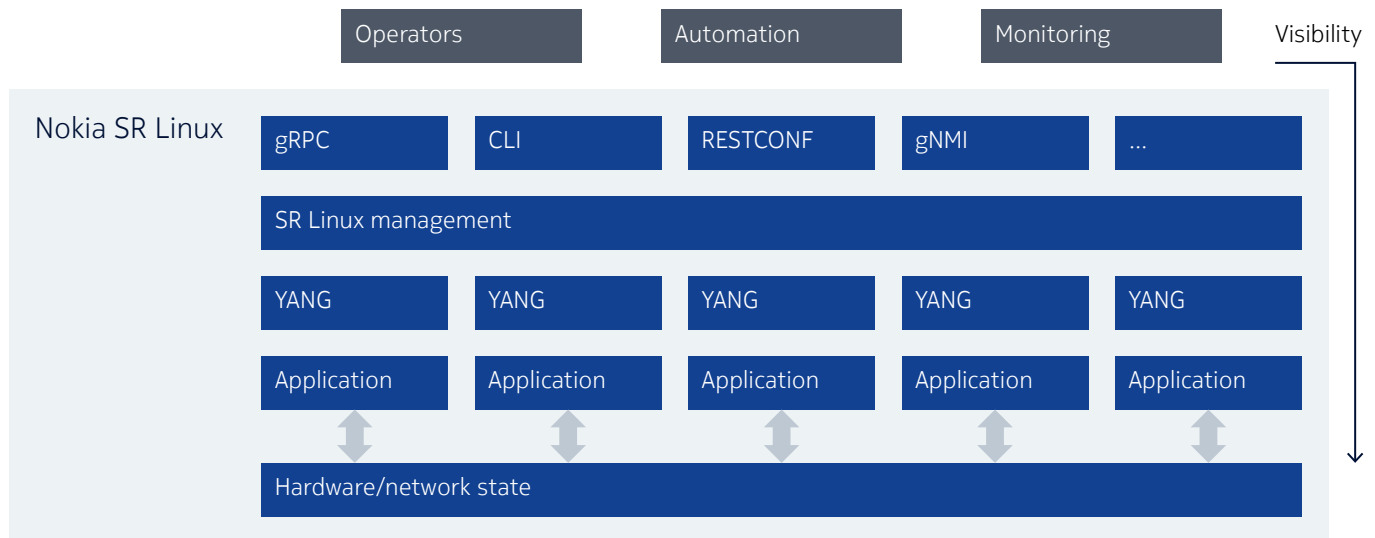
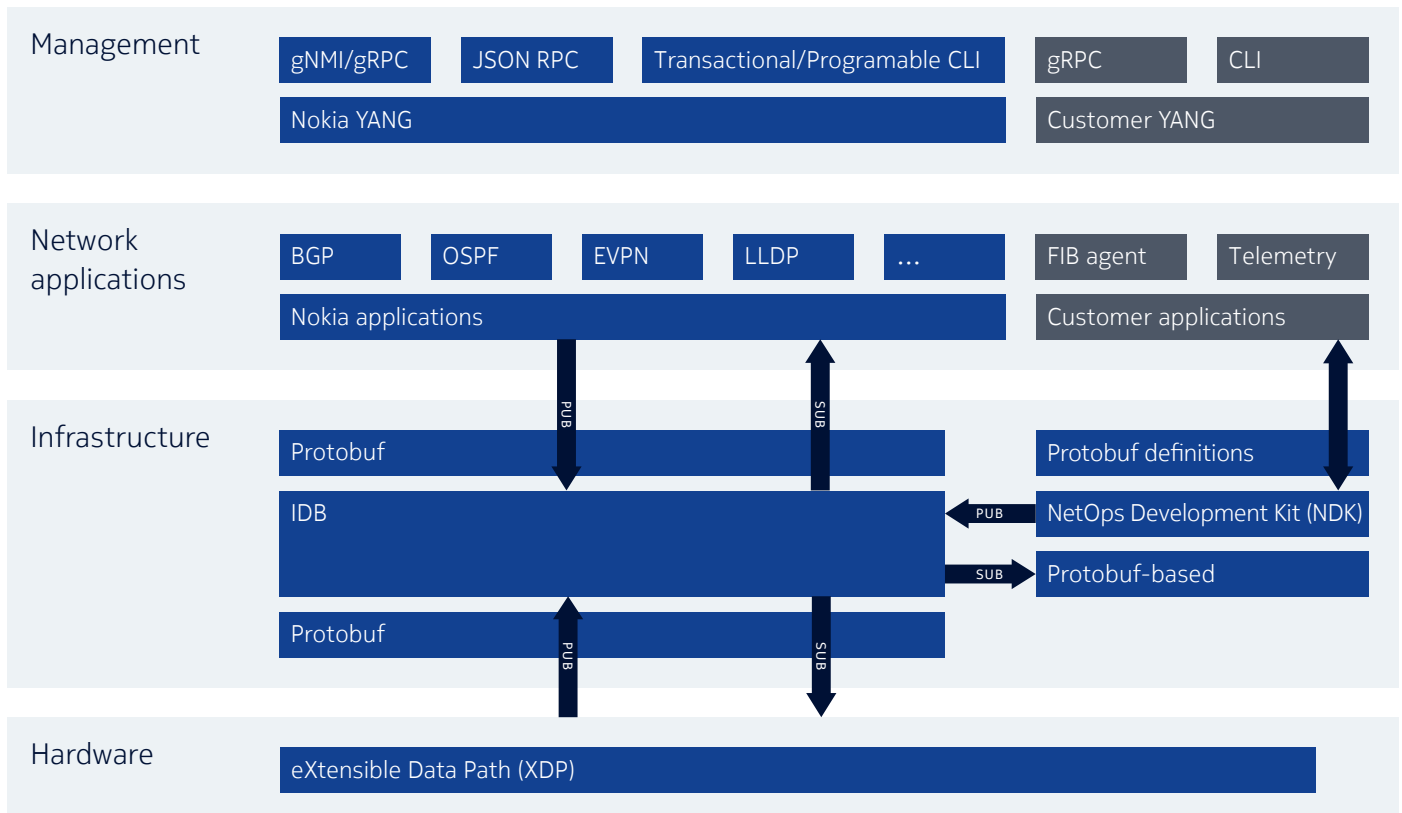


Figure 2. SR Linux modular, state-sharing and extensible architecture



Field-proven protocol stacks

A robust NOS is characterized by operating system software stability and its ability to support field-proven, scalable and reliable software feature sets. Data centers are increasingly adopting leaf-spine fabric designs using enhanced IP routing with Multiprotocol Border Gateway Protocol (MP-BGP), Ethernet VPN (EVPN) and Virtual Extensible LAN (VXLAN) protocols.

SR Linux was designed using field-proven protocol stacks from the Nokia Service Router Operating System (SR OS). The SR OS has a strong pedigree, with over 1.3 million routers deployed in more than 1,600 IP networks, including the internet backbone and some of the largest service provider, webscale and enterprise networks worldwide.

By using these field-proven protocol stacks, data center operators can immediately benefit from the robustness, stability, scalability and interoperability of a proven operating system.

High-availability features

Hardware and software high availability is crucial for ensuring maximum system uptime in data center and cloud environments. SR Linux-enabled hardware platforms support redundant fan and power configurations as well as hot-swappable, redundant control and fabric modules in modular chassis-based hardware platforms.

SR Linux supports warm reboot features that can be used to perform a soft reset or trigger an in-service software upgrade (ISSU). Nonstop forwarding (NSF) capabilities minimize data plane outages. NSF capabilities also enable data center hardware platforms to continue to forward packets. While leveraging control plane graceful restart, peers can continue to pass traffic.

Mature EVPN implementation

The SR Linux EVPN implementation leverages the field-hardened Nokia SR OS and delivers the same benefits of proven EVPN resiliency, stability, scalability and interoperability as the extensively deployed SR OS.

Our enhanced EVPN feature set supports both Layer 2 and Layer 3 EVPN connectivity with VXLAN data-plane encapsulation, EVPN host route mobility, plus all-active and single-active EVPN multi-homing.

All-active EVPN multi-homing supports up to four leaf nodes per multi-homed server, compared to an active-active implementation, which supports only two leaf nodes per multi-homed server.

MPLS and segment routing for data center WAN connectivity

SR Linux leverages proven MPLS protocol stacks from the Nokia SR OS. MPLS capabilities include support for Label Distribution Protocol (LDP) and segment routing (along with either Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First Interior Gateway Protocols (OSPF IGPs), allowing BGP to use either LDP- or segment routing-based MPLS shortcuts between data center locations. This functionality allows SR Linux to seamlessly connect IP-routed domains within data centers across the MPLS domain in the WAN.

Scalable streaming telemetry

SR Linux is designed to meet the demands of a model-driven world where visibility—and the scalability and granularity of that visibility—are paramount. SR Linux delivers an open, extensible and performant infrastructure that allows the retrieval of fine-grained system state, setting of configuration, and a scalable interface to support more granular data with push-based streaming.

SR Linux was built with an open, scalable telemetry framework at its core, internally using gRPC, gRPC Network Management Interface (gNMI) and protobufs. Because SR Linux is natively model-driven, it is immediately ready for streaming telemetry without requiring any translation layers.

Superior CLI programmability

Every SR Linux application (including third-party network applications) supports its own YANG model, which can be loaded into the system. With this design, the YANG data model is defined first, and from it, operators can derive the command line interface (CLI), the application programming interfaces (APIs), and the show-output formats related to software capabilities.

In addition to the gNMI interface, SR Linux includes an advanced, Python™-based CLI and a JSON-RPC API for management. The CLI provides a flexible framework for accessing the system's underlying data models and is based on quality-of-life (features that improve the operator experience) embraced by DevOps communities.

Operators can leverage CLI plugins to completely customize the way the CLI operates, plugging in Linux commands or pulling the state/configuration from various locations and combining these with system state/configuration to allow advanced logic. This capability streamlines the adoption of SR Linux because the interface can be customized.

Enhanced NDK

Nokia SR Linux allows third-party network applications to be fully integrated into the system with the same functionality as Nokia applications. This includes consistent configuration using YANG, telemetry support, life cycle management and visibility of system resources.

The Nokia NetOps Development Kit (NDK) enables application developers to leverage SR Linux's underlying model-driven architecture, with a simple, clean, decoupled integration. This allows all applications in the system to support data modeling, transactional configuration and—most important—massively scalable streaming telemetry.

The NDK uses gRPC and protobufs to provide maximum flexibility for languages supported and backwards compatibility. This approach differs from others, which are restricted to certain languages, versions and/or libraries.

The NDK allows data center operators to teach the network their language and respond directly to business demands—all without worrying about the scalability or functionality of the routing stack or the underlying infrastructure.

With the mindset to build for openness, common infrastructure utilized by Nokia applications is exposed, allowing a uniform operating model and visibility deep into the heart of the system.

Nokia Fabric Services System

The Nokia Fabric Services Platform complements and extends the capabilities provided by our SR Linux architecture foundation. The platform provides a modern, flexible toolkit that delivers automation at scale for all phases of data center fabric operations, including Day 0 design, Day 1 deployment and Day-2+ configuration, operation, measurement and analysis of a data center fabric.

The Fabric Services Platform enables a full turnkey data center solution for automation, intent-based approaches, telemetry collection and analytics. The platform also includes a unique capability (the Digital Sandbox) for emulating the live network in software that assists with planning, testing and troubleshooting data center fabric designs.

The Fabric Services Platform offers significant improvements over existing fabric management systems, which are rigid in their monolithic application architecture and limited in features, flexibility and scalability.

The system leverages Kubernetes as the base for a highly extendable, containerized, microservices-based platform. This includes containerized platform functions, full access to existing Kubernetes utilities and applications, and the option for operators to include customized or third-party applications within the solution.

Nokia Data Center platforms

The Nokia portfolio of Data Center platforms addresses the needs of modern data centers. The portfolio offers a broad range of high-performance platforms for data center leaf-spine deployments. Both modular, chassis-based platforms and fixed-form-factor platforms are available.

The portfolio includes:

- Nokia 7250 IXR-10/IXR-6 Interconnect Routers for data center fabrics
- Nokia 7220 IXR-H series Interconnect Routers for data center fabrics
- Nokia 7220 IXR-D series Interconnect Routers for data center fabrics

The 7250 IXR and 7220 IXR products implement SR Linux.

Software features

SR Linux supports, but is not limited to, the following software features.

Open Linux support

- Support for unmodified Linux kernel
- Access to Linux tools, patching and packaging
- Containerized SR Linux
- Linux control groups (cgroupsv2)

Layer 2 features

- Dot1q and untagged sub-interfaces
- Ethernet IEEE 802.1Q (VLAN) with support for jumbo frames
- Link aggregation: Link Aggregation Group (LAG) and Link Aggregation Control Protocol (LACP)
- Link Layer Discovery Protocol (LLDP) on all interfaces
- MAC storm control
- Media access control (MAC) loop prevention
- Virtual routing and forwarding (VRF): MAC-VRF

Layer 3 features

- IPv4/v6 routing
- BGP with iBGP/eBGP: Support for IPv4/v6, including:
 - Core Prefix independent convergence
 - 4-byte autonomous system number
 - Route reflector
 - Dynamic BGP
 - eBGP multi-hop
- IS-IS v4/v6
- OSPFv2 and OSPFv3
- Static routes for IPv4/v6
- Equal cost multi-path with consistent and resilient hashing and configurable hash fields
- IPv6 flow label hashing
- VRF: Multiple VRF support
- Maintenance modes
- Bi-directional forwarding detection (BFD), micro BFD (mBFD)
- Interfaces: Loopback interfaces, Integrated Routing and Bridging (IRB)
- Routing policy: Layer 3/Layer 4 access control list (ACLs) with validation; accept, reject and log actions

MPLS and segment routing¹ (SR)

- Interface LDP over IPv4
- SR-ISIS over IPv4
- BGP shortcuts over LDP
- BGP shortcuts over SR-ISIS
- MPLS QoS via EXP to forwarding class mapping
- MPLS ACL filters
- ICMP tunneling
- ICMP extensions for MPLS

¹ Supported on Nokia 7250 IXR-6/IXR-10 platforms

Network virtualization²

- EVPN with VXLANv4 encapsulation
- EVPN Layer 2 and Layer 3 connectivity
- EVPN all-active multi-homing; single active multi-homing for Layer 2 and Layer 3
- EVPN host route mobility
- Provider edge customer edge (PE-CE) BGP path attribute propagation in EVPN

QoS

- Intelligent packet classification, including IPv4, IPv6 match-criteria-based classification
- Queuing/scheduling:
 - Strict priority
 - Weighted Round Robin (WRR)
 - Weighted Random Early Detection (WRED)
 - Explicit Congestion Notification (ECN)
- QoS classification based on DiffServ Code Point (DSCP)

System management and automation

- Native model-driven architecture, configuration candidates, exclusive mode, checkpoints, rollbacks
- Management interfaces: gNMI, JSON, CLI (transactional, Python CLI, CLI plugins)
- Per-user configurable options for CLI
 - Local Authentication, Authorization and Accounting (AAA) with Role Based Access Control (RBAC)
 - Terminal Access Controller Access Control System (TACACS+) AAA via privilege levels
- Access to common Linux utilities: Bash, cron, and Python

² Supported on Nokia 7220 IXR D2, D2L, D3 and D3L platforms

- Telemetry:
 - Subscription-based telemetry for modeled data structures, either changed or sampled
 - sFlow
 - Logging infrastructure
- Python-based Zero Touch Provisioning (ZTP)
- Address management: Dynamic Host Configuration Protocol (DHCP) v4/v6 relay
- DHCP v4/v6 server with static allocations
- Interactive mirroring
- Unified Forwarding Tables (UFT) profiles
- NDK:
 - gRPC and protobuf-based interface for tight integration
 - Leverages SR Linux model-driven architecture
 - Direct access to other application functionality, e.g., forwarding information base (FIB), Link LLDP and BFD
 - Native support for streaming telemetry

Resiliency

- Support for redundant fan and power configurations in data center hardware platforms
- Support for hot-swappable, redundant control and fabric modules
- Warm reboot³ to perform soft reset or trigger an in-service software upgrade (ISSU)
 - Nonstop forwarding (NSF)
 - Graceful restart client for BGPv4/v6

Security

- Distributed and aggregated ACLs and policers for control and management plane
- Mirroring from interface/sub-interface or ingress ACL
- Mirroring to Switch Port Analyzer³ (SPAN) and Encapsulated Remote SPAN (ERSPAN)
- IPv6 Router Advertisements (RA) guard

Supported hardware platforms

- Modular platforms:
 - Nokia 7250 IXR-10/IXR-6 Interconnect Routers for data center fabrics
- Fixed configuration platforms:
 - Nokia 7220 IXR-D series Interconnect Routers for data center fabrics
 - Nokia 7220 IXR-H series Interconnect Routers for data center fabrics

Learn more

To learn more about the Data Center Switching Fabric solution, see the [web page](#)

³ Supported on Nokia 7220 IXR D2, D2L, D3 and D3L platforms



About Nokia

We create the critical networks and technologies to bring together the world's intelligence, across businesses, cities, supply chains and societies.

With our commitment to innovation and technology leadership, driven by the award-winning Nokia Bell Labs, we deliver networks at the limits of science across mobile, infrastructure, cloud, and enabling technologies.

Adhering to the highest standards of integrity and security, we help build the capabilities we need for a more productive, sustainable and inclusive world.

For our latest updates, please visit us online www.nokia.com and follow us on Twitter [@nokia](https://twitter.com/nokia).

Nokia operates a policy of ongoing development and has made all reasonable efforts to ensure that the content of this document is adequate and free of material errors and omissions. Nokia assumes no responsibility for any inaccuracies in this document and reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

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.

© 2021 Nokia

Nokia Oyj
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: 20785944 (December) CID207598