Nokia Fabric Services System

A complete data center fabric management, operations and automation platform

Product description

Abstract

As 5G and other cloud-native architectures are being deployed to accelerate the next generation of IT applications, data center operators are facing increased network demands with a budget that is often flat and a staff that is not growing.

To address this emerging reality, networks need to be much more responsive while providing better performance and elastic scalability. Operators also need to make dramatic productivity improvements to their network operations. For network operators to meet these lofty ambitions, the deployment of advanced NetOps automation techniques is a necessity.

The Nokia Fabric Services System was built from the ground up to address these demands. This declarative, intent-based automation and operations NetOps platform delivers agile and scalable network operations for data center and cloud environments.

This document describes the Fabric Services System and explains its foundational elements and key capabilities, including:

- Declarative, intent-based operations and automation across the entire operational life cycle of the data center fabric, from Day 0 design to Day 1 deployment and Day 2+ operations, that simplifies and dramatically reduces effort across many operational tasks
- Observability through “on-change” streaming telemetry from Nokia Service Router Linux (SR Linux); this observability informs and optimizes operational tasks
- A Digital Sandbox, enabling a digital twin capability that emulates the network and is used to dramatically reduce both testing time and network risk
- A Continuous Integration/Continuous Deployment (CI/CD) framework embodying DevOps principles that are applied to the network
- Integration with major cloud management platforms, enabling the automation of network connectivity to support constantly changing application requirements.

Contents

Introduction

With the proliferation of 5G, cloud-native deployments, the hierarchical distribution of compute across many data centers, and the emergence of next-generation consumer and enterprise services, networks are becoming much more complex. With Industry 4.0 and the high-performance use cases that 5G enables, networks are expected to grow to support larger volumes of traffic and large increases in both performance and scalability.

Meeting these demands with existing operational methods and tooling is not sustainable. Network operators need to deploy a methodology that removes the network as an obstacle to agile application development and deployment. At the same time, operators need to dramatically increase operational productivity in the network.

The Nokia Fabric Services System is designed from the ground up to meet these demands by providing a declarative, intent-based framework that allows operators to automate job functions and tasks across the entire life cycle of data center fabric operations, from Day 0 design to Day 1 deployment and Day 2+ operations.

The Fabric Services System is a complete data center fabric management, operations and automation platform. The Fabric Services System functionally provides four key capabilities (see Figure 1):

- **Declarative intent:** Simplifies complexity and automates thousands of tasks with intent-based fabric automation across the entire life cycle of the fabric.
- **Fabric observability:** Provides an unprecedented level of network data access and insight to empower automation and other operational functions.
- **Fabric operations:** Automate operations by programming a complete Continuous Integration/ Continuous Deployment (CI/CD) workflow while lowering the risk of network changes by validating candidate changes with the network's Fabric Services System Digital Sandbox.
- **Fabric integrations:** Integration with major cloud management platforms to ensure that the network stays synchronized with the applications.

Figure 1. Nokia Fabric Services System key capabilities

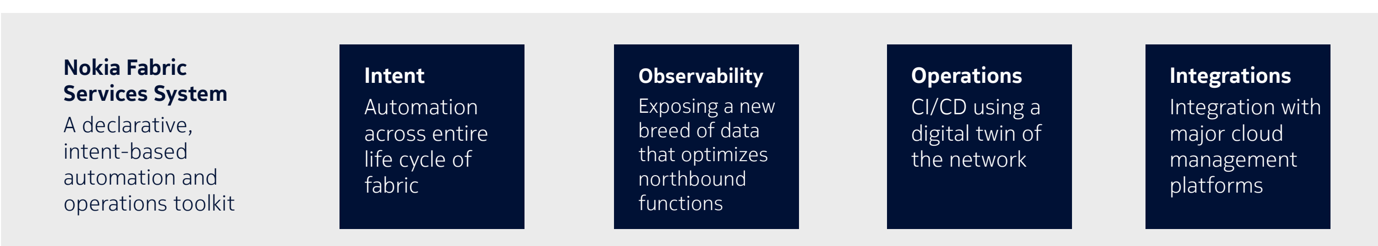

Fabric Services System as part of Nokia Data Center Fabric solution

Since the advent of software defined networking (SDN), the industry has experimented with the separation of the management, control and data planes. For scalable and faster convergence in large-scale data center networks, the architectural approach that has gained industry momentum is a combination of distributed routing (using control planes and data planes) running on data center switches with the necessary network-wide control and automation functions implemented in an external controller. This approach combines the centralized control and programmability of a traditional SDN approach with the higher scalability and convergence of a distributed routing approach. This is the approach we have taken with the Nokia Data Center Fabric solution.

The Fabric Services System is a key component of the Data Center Fabric solution (see Figure 2), which also includes the following products:

- Nokia Service Router Linux (SR Linux): An open, extensible and model-driven network operating system (NOS) based on Linux® that enables scalability, flexibility and efficiency in data center and cloud environments.
- Nokia Data Center Fabric hardware platforms: A portfolio of next-generation leaf and spine switches that deliver massive scalability and performance while aggregating and interconnecting data center and cloud environments.

Figure 2. Nokia Data Center Fabric solution components

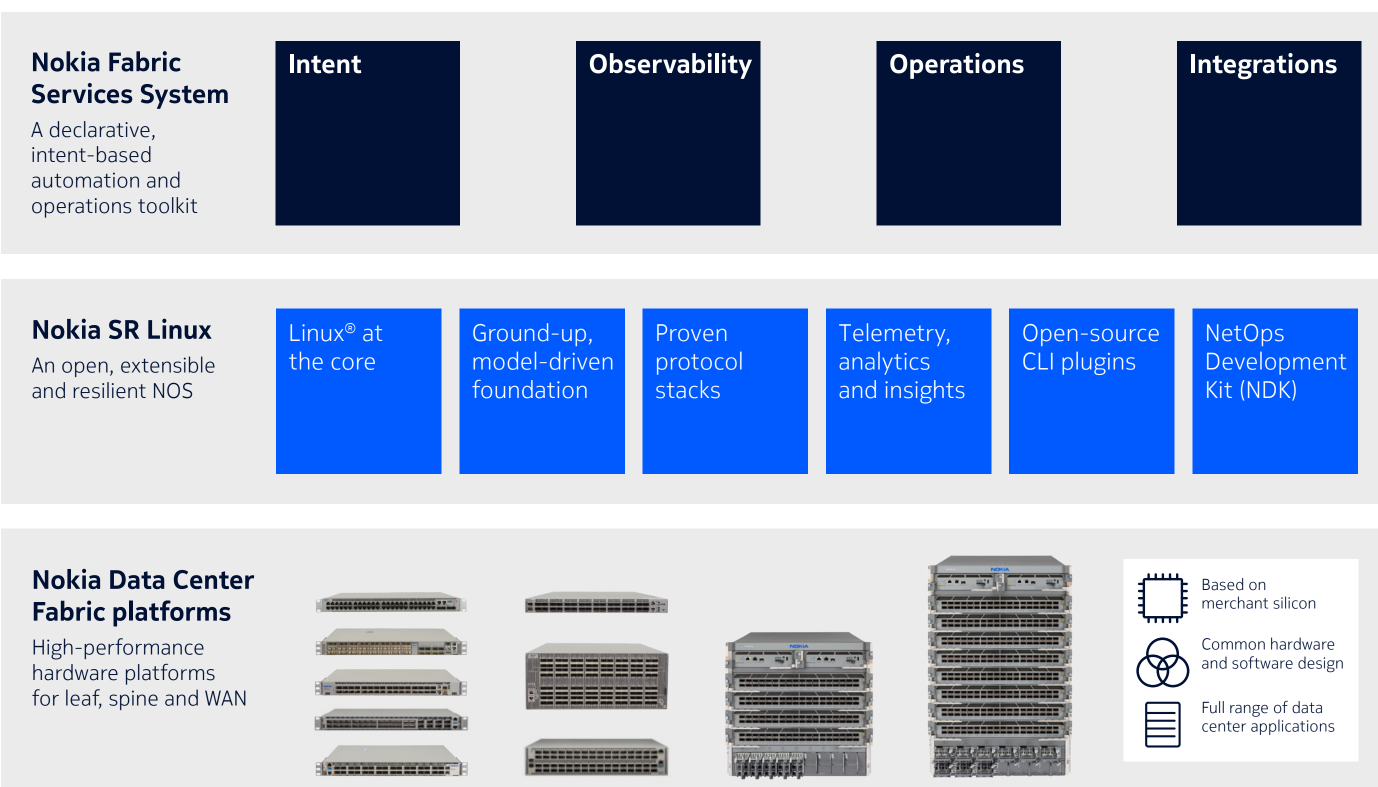

The Data Center Fabric solution enables data center operators to rapidly design and deploy, easily adapt and integrate, and confidently operate and automate data center network fabrics at scale. The solution leverages our:

- Expertise in IP routing and network operations
- Nokia Service Router Operating System (SR OS)
- IP network automation solutions
- Proven track record of building business-critical Ethernet and IP/MPLS networks for communications service providers (CSPs), webscale companies and enterprises globally.

Importance of an open, model-driven, consumable NOS

The declarative, intent-based framework and the automation capabilities of the Nokia Fabric Services System are further enhanced by leveraging a modern NOS that offers an open, model-driven foundation.

The Nokia SR Linux NOS couples each service with its own YANG data model, allowing for broad, deep and efficient network data access from any interface. This modular approach provides open streaming telemetry and network management that use gNMI (gRPC Network Management Interface) to stream network data and configure network devices.

Figure 3 shows the relationship between the Fabric Services System and SR Linux.

Figure 3. The importance of a model-driven NOS

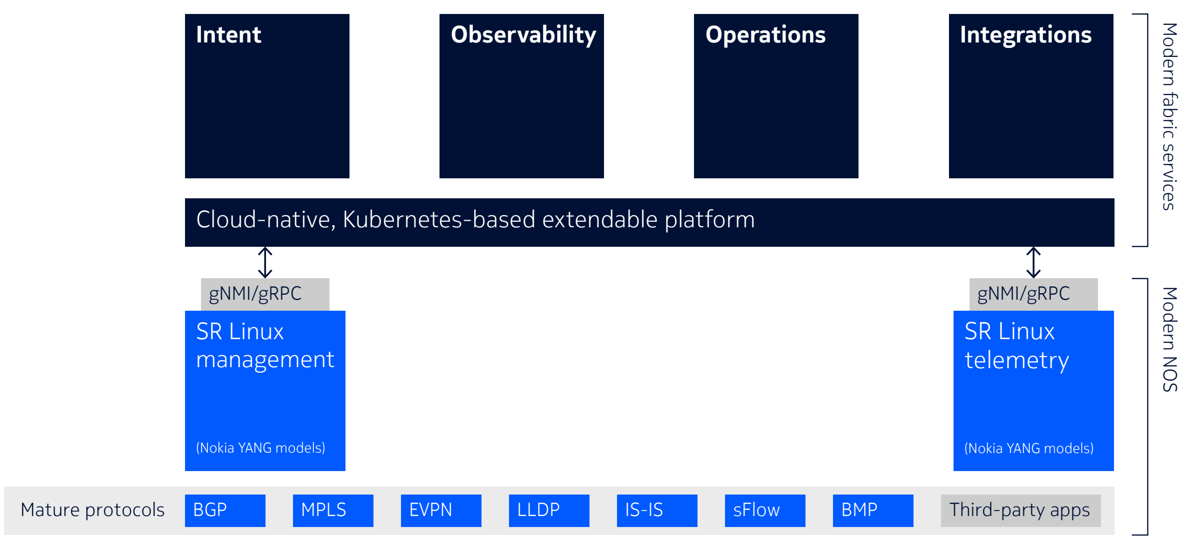

By using this modern approach for the NOS, the Fabric Services System has access to more granular data across the entire fabric in a timely and efficient manner. This data can then be used to understand the state of the network, which is essential for the Fabric Services System to determine if the network is behaving according to its desired intent. In addition, this approach is highly scalable, which is essential in today's networks.

Architecture foundations

The Nokia Fabric Services System adopts a microservices design that is built on top of Kubernetes. The Fabric Services System also adopts many other open-source projects, including Kafka, Keycloak and Apache Zookeeper along with Grafana and Kibana.

As shown in Figure 4, there are several layers of the architecture.

Figure 4. Layers of the Fabric Services System

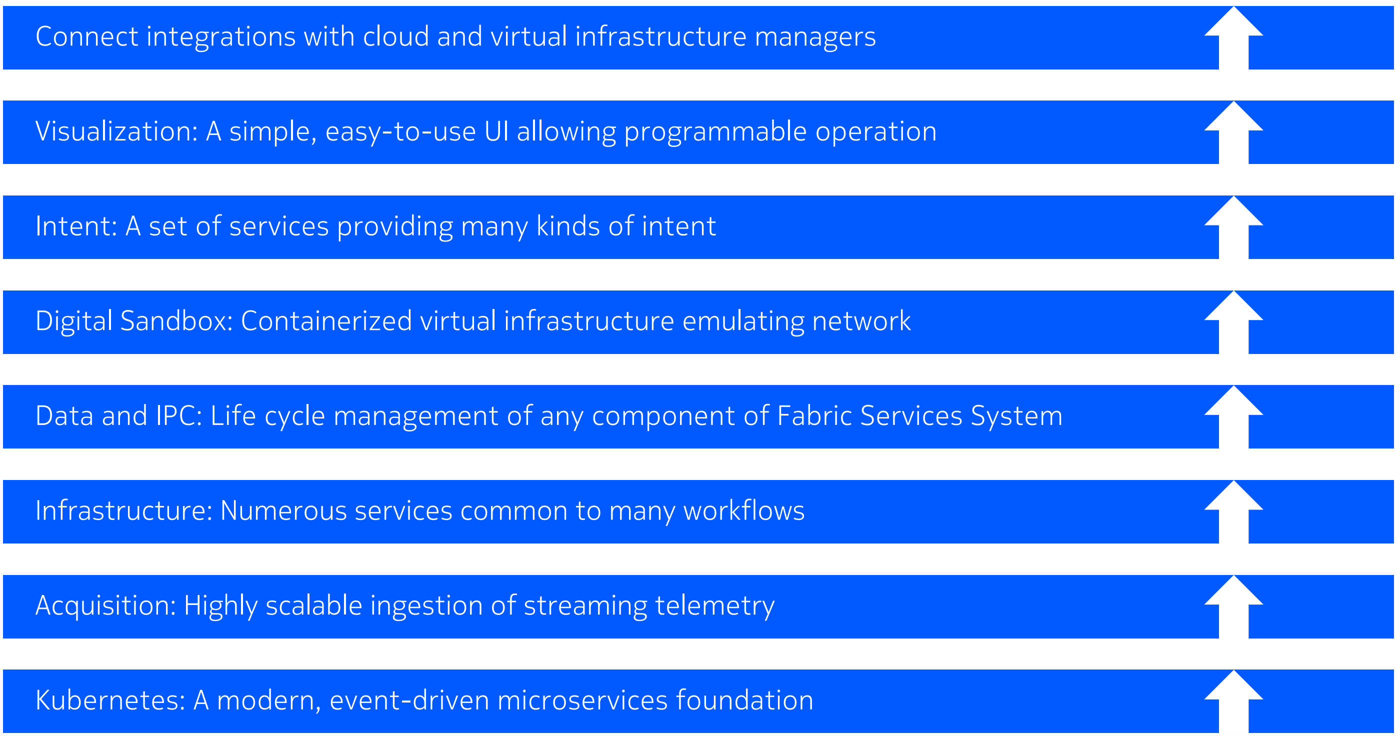

Starting from the bottom of the figure and working up, the layers are:

- **Kubernetes:** Provides an event-driven microservices foundation. Running natively in Kubernetes has numerous benefits, including providing multiple storage backends as well as the ability to define the entire deployment through infrastructure-as-code (IaC) principles using Helm charts as a templating engine.
- **Acquisition:** Purpose-built to ingest the highly scalable streaming telemetry offered by SR Linux and to perform bootstrap and discovery of the fabric. This layer implements a highly scalable gNMI collector service whereas the ConfigSync service is responsible for pushing down configuration based on fabric and workload intents. Other common services in this layer perform initial and ongoing Zero Touch Provisioning (ZTP) of devices, allowing an end user to implement a simple “plug-and-power-up” approach to onboard new devices onto the fabric.
- **Infrastructure:** Contains numerous services common to many different workflows. One service is the label manager, which allows the flexible grouping of objects within the system. This grouping enables policies and intents to be applied to sets of objects instead of to individual objects.

- **Data and inter-process communication (IPC):** Leverages Kafka to distribute work among services and Zookeeper to distribute configuration. Using Kubernetes' kubelet to monitor pods, along with persistent storage via Kafka, IPC allows for any component in the system to be independently life cycle managed–restarted, upgraded, or scaled in or out depending on the offered load. This persistent storage can also be replicated to a secondary site, allowing geo-redundancy between locations.
- **Digital Sandbox:** A containerized virtual infrastructure that emulates the production network by creating its digital twin with exact parity. The Digital Sandbox is used to validate intent across the life cycle of the fabric and is also an essential part of the network validation phase of the CI/CD pipeline process.
- **Intent:** A set of services that provides several types of intent, including fabric intent (Day 0 design), workload VPN intent (Day 1 deployment), and maintenance intent (Day 2+ operations). The intent framework of the Fabric Services System allows operators to define, in an abstract manner, the desired end state of the fabric. By using streaming telemetry to understand the current state, the system can determine any discrepancies from the desired state and implement any required network changes.
- **Visualization:** A simple, easy-to-use UI that allows for complete programmable operation leveraging visualizations using standards-based, open-source solutions such as Grafana and Kibana. Operations that can be performed through the UI can also be performed through REST APIs.
- **Connect:** Performs integration with cloud management platforms, allowing virtual machine (VM) or container "spin up" and "tear down" events to drive network change. This capability enables the data center fabric to react to workload and compute connectivity requirements. Connect integrates using REST APIs and a plugin-based model, enabling seamless, modular and simple integration with cloud management platforms.

Digital Sandbox

One of the key requirements for modern data centers is the ability to make faster, periodic changes while still managing the risk of a change. To enable this capability, the Fabric Services System delivers a cloud-native Digital Sandbox that emulates the data center fabric's live network. The Fabric Services System integrates the Digital Sandbox in all its workflows to provide design validation and change management flexibility, thereby reducing the risk of changes in a dynamic data center environment.

The Digital Sandbox is a virtual network infrastructure (VNI) the operator can use to test and validate any planned network changes. The Digital Sandbox provides a digital twin of the data center fabric, emulating the leaf and spine switches by deploying a containerized Nokia SR Linux instance of each.

The Digital Sandbox leverages on-change telemetry to maintain absolute parity of the network in configuration, routing and state. It also can emulate external Border Gateway (BGP) speakers to generate synthetic traffic.

The Digital Sandbox allows any changes to the production network to be tested and validated before being deployed in the network, greatly reducing risk.

Some of the benefits of Digital Sandbox are:

- **Time and resource savings:** Saves time and resources by quickly and efficiently testing network, configuration and routing scenarios in a virtualized, pre-built environment that is in absolute parity with the network.
- **Reduced risk:** Greatly reduces risk to the network by first validating network changes in a fully emulated environment before deploying the changes in the network.

- **Lower lab expenses:** Reduces the effort and cost of setting up lab environments to test and validate network changes.
- **Reduced power consumption:** Drives a green approach to testing and validation by leveraging a virtualized environment that can be set up and changed in minutes.
- **Ease of use:** A complete virtual infrastructure is built into the Fabric Services System and is fully programmable and easily set up through an intuitive UI.

Declarative, intent-based automation

The Nokia Fabric Services System allows operators to represent the configuration and initial state of the data center fabric in a declarative, intent-based way. With this declarative approach, the desired configuration and state of the fabric can be specified up front in a simplified or abstract manner that defines how the fabric should operate. This desired state, which is stored centrally, represents “the single source of truth” and can be used to iteratively validate the actual state of the network.

This approach provides a strong NetOps foundation that leverages DevOps principles and fits into the bigger movement toward IaC. The Fabric Services System leverages a set of microservices designed to handle various types of intent, including design, deployment and maintenance for Day 0, Day 1 and Day 2+, respectively.

Day 0 intent-driven design

By taking an abstract, intent-based approach for Day 0 design, the data center operator can focus on high-level aspects of the design, identifying the minimal information needed to build a data center fabric. For example, the operator needs to input only a few parameters, such as the number of racks and the number of servers per rack.

Using this information, the system autogenerates the rest of the detailed configuration based on Nokia-certified design templates. The result is a standard BGP-based IP fabric design (e.g., number of racks, number of servers per rack, IPv4/IPv6 addressing, BGP configuration, cable map, etc.) that can be validated using the Digital Sandbox before being deployed to the data center fabric.

Figure 5 shows this workflow and includes how design documentation can also be created through this same intent-based templating system.

Figure 5. Intent-based Day 0 design with the Fabric Services System

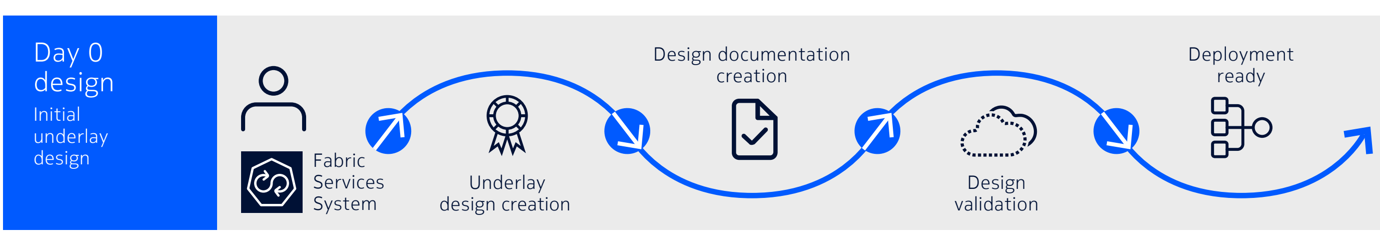

With this intent-based approach, multiple leaf-spine fabrics can be created by easily replicating the first one created or by creating a customized fabric. These fabrics can then be connected using a backbone and the same intent-based approach. The operator simply selects which fabrics to include as its uplink (see Figure 6), and the system autogenerates the configuration needed to support this intent.

Figure 6. Expressing Day 0 design in a modular way

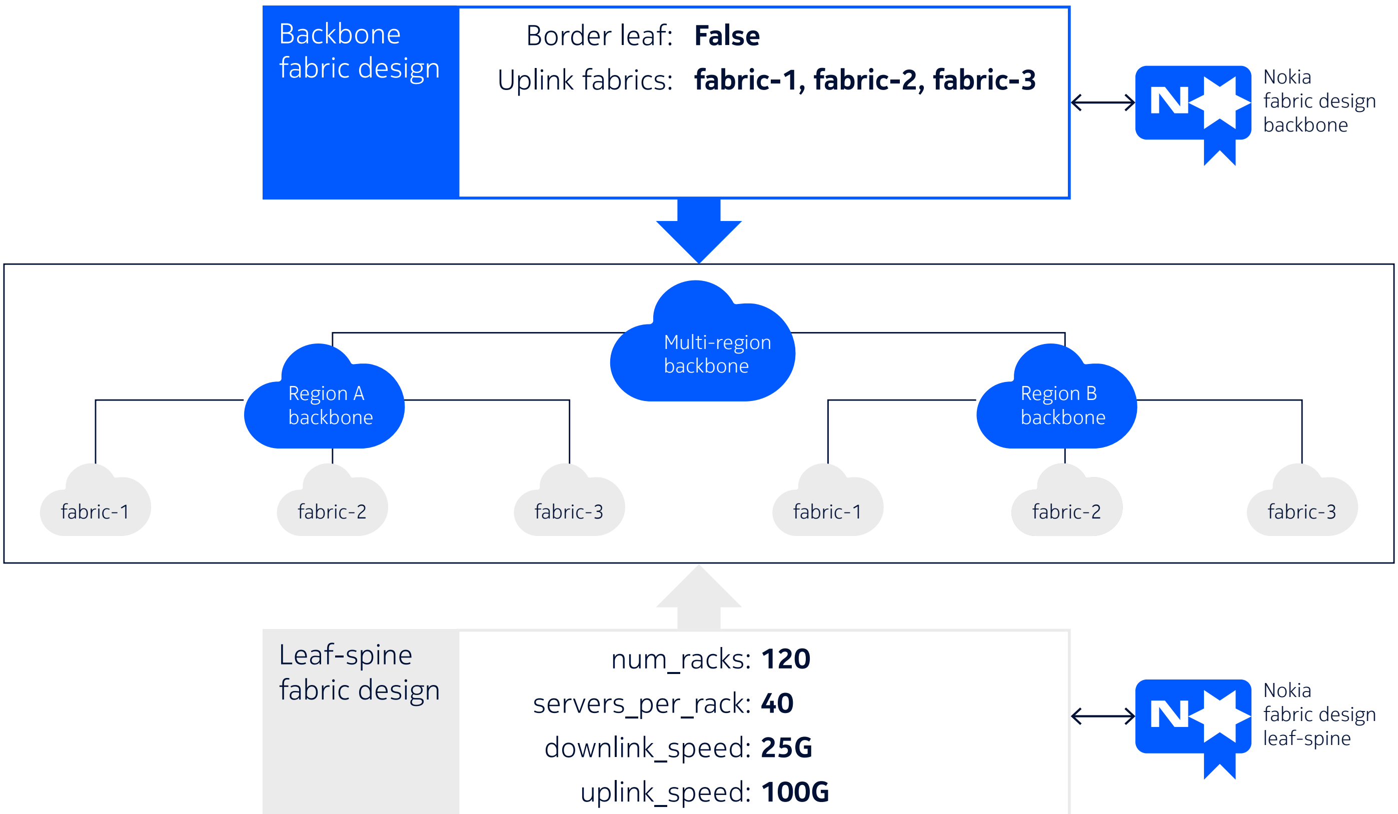

Day 1 intent-driven deployment

For Day 1 deployment, one of the initial tasks performed by the Fabric Services System is fabric discovery and node bootstrap. In addition, the Fabric Services System offers ZTP to turn up new leafs and spines, allowing the adoption of a simple plug-and-power-up approach to onboard new nodes onto the fabric.

After the new nodes are onboarded, the Fabric Services System can then push Day 0's validated design to the fabric, thereby completing deployment of the initial network underlay portion of the fabric.

Day 1 deployment uses workload intent to automate the creation of the required overlay connectivity, to support the initial application workloads that are hosted on attached compute resources. To create this connectivity, the Fabric Services System leverages Ethernet VPN (EVPN)-based implementation, available on SR Linux, to deliver multi-homed Link Aggregation Groups (LAGs) as well as L2 and L3 services within and across the data center fabric.

Workload intent abstracts the complexity of the EVPN configuration by enabling the data center operator to focus on specifying high-level parameters. This high-level intent can be as simple as identifying the set of downlinks an application workload uses to connect to the fabric. Workload intent can be validated using the Digital Sandbox before being deployed into the production network. Figure 7 shows this workflow.

Figure 7. Intent-based Day 1 deployment with the Fabric Services System

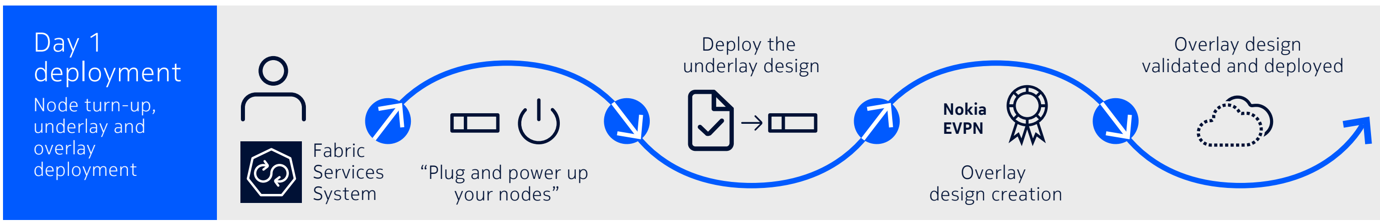

Complexities such as switch-to-switch EVPN and allocation of VXLAN network identifiers, route distinguishers, route targets, Ethernet segment IDs and Ethernet virtual interfaces are all abstracted and left to the Fabric Services System to generate according to the high-level intent parameters specified by the operator. Some of the abstracted complexity is shown in Figure 8.

Figure 8. Hiding the complexity of the overlay design and configuration

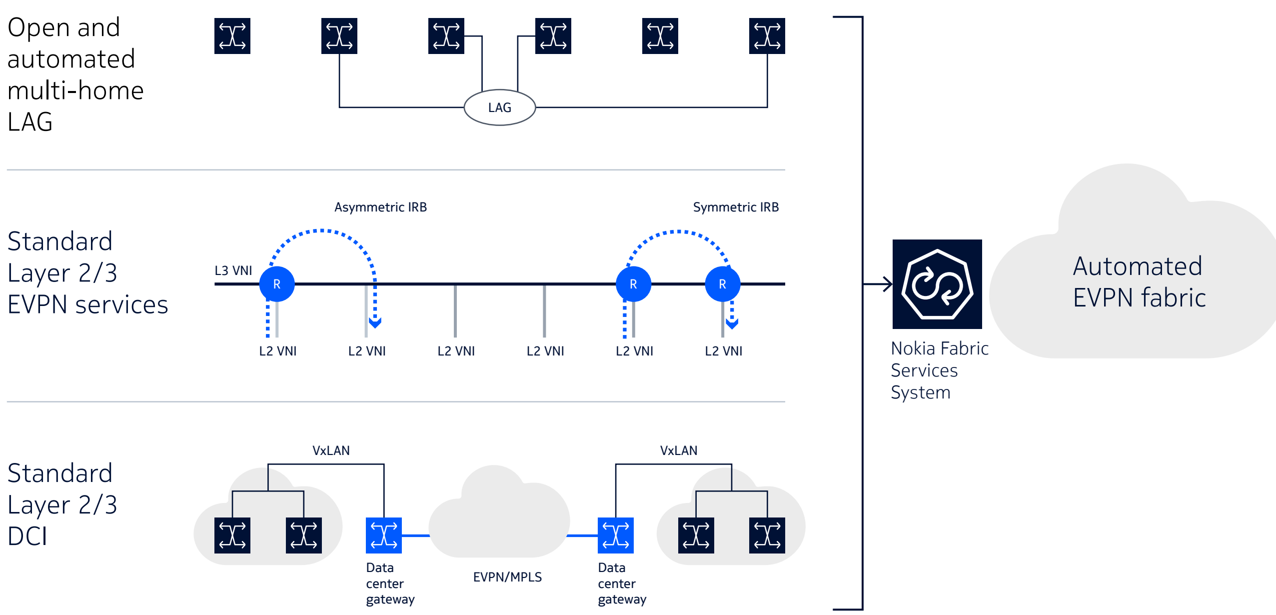

Intent expressed as IaC

As part of using intent to create both the fabric underlay and overlay designs, the system generates machine-readable code that represents the intricate design and configuration details. This code is sent to the network elements (e.g., leaf and spine switches) so that these systems can be turned up and configured. This machine-readable code can be stored in several formats. One popular variant is YAML.

This code represents the desired configuration of the network at any point in time and, similar to software code, it can and will change over time as the network changes (e.g., as new nodes or new features are added). The code is treated like software in that it is version controlled and maintained as a monolithic entity onto itself: infrastructure as code (IaC). This IaC intent represents the single source of truth for the network and is used to continuously validate the actual network state to ensure that the desired state is always met.

With this approach, all configuration and state of the data center fabric is represented in a declarative way in YAML format, supporting IaC while laying a strong foundation for deploying a CI/CD framework for the network infrastructure.

Intent with an IaC approach reduces cost and risk while increasing speed:

- **Cost reduction:** Removing the manual component of configuration reduces human work by many hours, allowing skilled workers to refocus their efforts on other essential data center tasks.
- **Risk reduction:** Automating these tasks eliminates the risk associated with human error—which also decreases downtime and increases reliability.
- **Improved agility:** Intent-based automation improves service agility, allowing operators to test and turn up new services with much more speed and accuracy than before.

Day 2+ intent-driven maintenance

During Day 2+ operations, the Fabric Services System uses maintenance intent (i.e., hardware intent and software intent) to define the desired state of the network in terms of software and hardware. With these two intent types defined, the desired software load and hardware version across the network is defined for each leaf, spine or super-spine switch.

Fabric observability

To operate today's modern data center fabrics, real-time observability information is required to inform various operational tasks. Fabric observability is needed to monitor and provide visibility into the fabric and is achieved by accessing a combination of on-change, multi-dimensional streaming telemetry and log data that represents the network state and is collected directly from the data center fabric.

Multi-dimensional telemetry comes from various sources, including:

- Basic telemetry: Faults, standard statistics, Telecommunications Access Method/Longest Prefix Match (TCAM/LPM), etc.
- The control plane: Link Layer Discovery Protocol/Link Aggregation Control Protocol (LLDP/LACP) state and events, BGP adjacency, BGP routing information base (RIB), forwarding information base (FIB), etc.
- The fabric workload layer: Topology, number of apps, number of flows, etc.

The Nokia Fabric Services System constantly receives this information via the SR Linux gNMI and leverages a purpose-built acquisition layer to ingest this streaming telemetry while scaling as required. The Fabric Services System enables a cloud-native, scale-out collector architecture to ensure that collection capabilities are highly distributed.

Fabric operations

After the data center fabric is designed and deployed, the Day 2+ operations phase begins. In this phase the Fabric Services System compares both design and workload intent (i.e., single source of truth) with all the telemetry data collected from the fabric to both optimize operational tasks and ensure that the network is operating as expected.

Day 2+ intent-driven operations

For Day 2+ operations, the Fabric Services System constantly monitors the fabric by leveraging on-change, multi-dimensional telemetry it receives directly from various sources in the network. The Fabric Services System compares this information with various intents and analyzes the results to find configuration inconsistencies, faults or other deviations that may lead to network issues.

Each inconsistency, fault or other deviation is flagged and presented to the operator to be either accepted or rejected. These inconsistencies can often require a change to the network (e.g., a configuration change or software upgrade) to fix the problem. With the Fabric Services System, the operator can automate the testing and validation of these network changes using the Digital Sandbox. If these changes pass validation, they can be scheduled for automatic deployment into the production network.

This process of automated testing and validation (see Figure 9) dramatically lowers the risk in deploying network changes because it identifies any potential problems before a change is deployed in the network.

Figure 9. Day 2+ intent deviation detection with the Fabric Services System

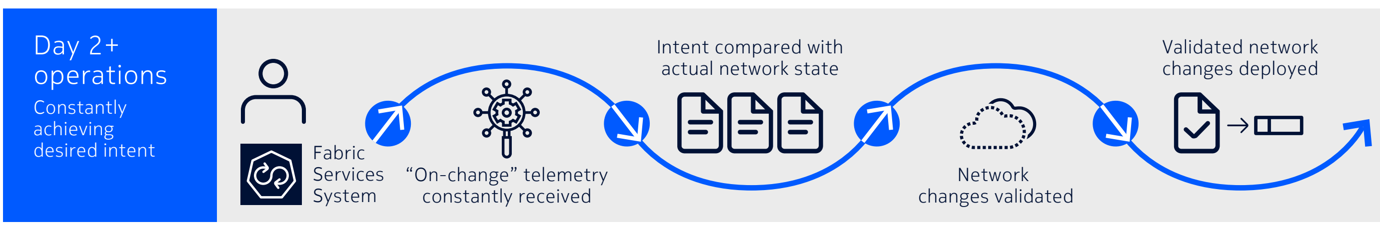

Fabric Services System and the data center fabric's operational life cycle

Figure 10 shows how the Nokia Fabric Service System acts on the data center fabric across all phases of its operational life cycle.

Figure 10. Data center intent-based fabric operations with the Fabric Services System

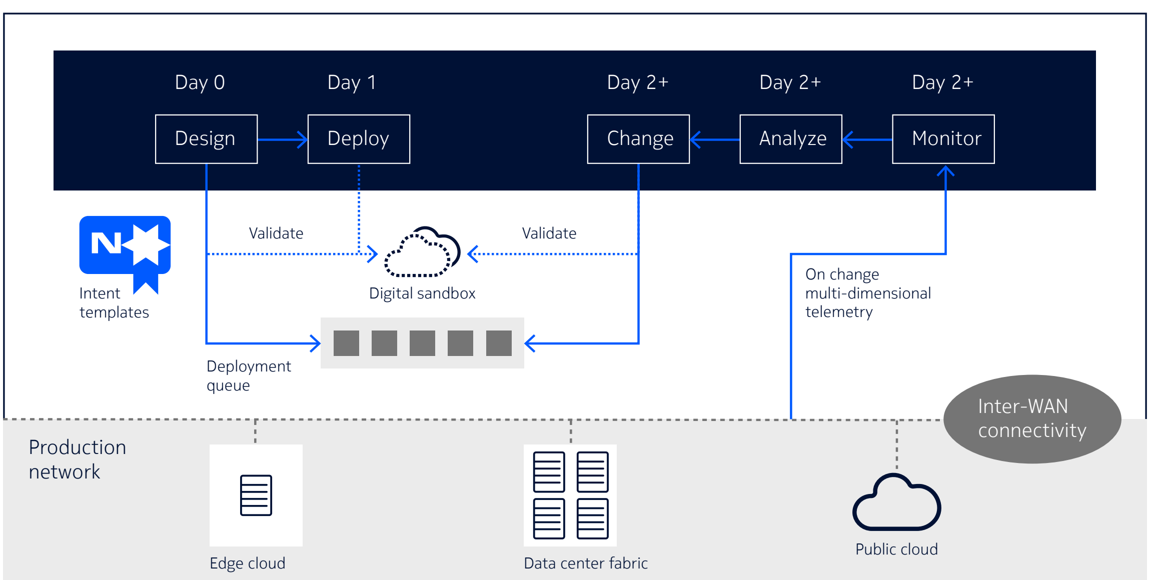

Day 0 fabric designs (e.g., number of racks, number of servers per rack, IPv4/IPv6 addressing, BGP configuration, etc.) and Day 1 network configurations (e.g., Ethernet VPN and Layer 2/Layer 3 [L2/L3] overlays) can be automated with intent-based templates that are first tested and validated in the Digital Sandbox before being sent for deployment in the production network.

During Day 2+ operations, the system constantly monitors and analyzes the network for any deviations from intent. Any resultant network changes can be validated by the Digital Sandbox before being sent for deployment.

Implementing CI/CD for the network

Inherited from the principles of DevOps, one of the emerging requirements for modern data centers is to implement for the network a process that leverages principles from CI/CD. With this capability, operators can make faster, periodic and independent changes to the network, lowering risk and increasing operational agility.

The Digital Sandbox is an integral part of Nokia's approach for CI/CD and is used to trial and validate network changes before deploying them in the production network. Changes can include initial fabric design, initial workload connectivity, software upgrades, introduction of new devices, policy and configuration changes, and failure scenarios.

Figure 11 shows how CI/CD is implemented in the Fabric Services System.

Figure 11. The Digital Sandbox in CI/CD

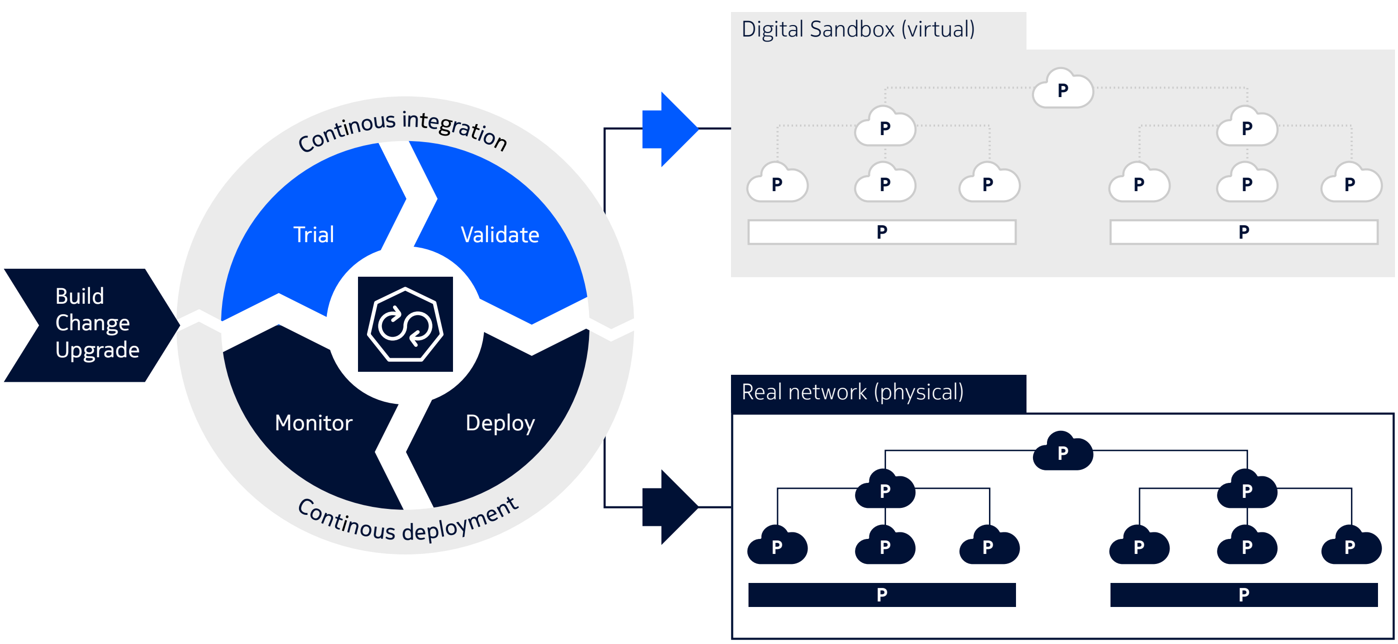

Fabric integrations

The Nokia Fabric Services System provides an open REST API that allows third parties to have full access to the system. A flexible, cloud-native approach enables the Fabric Services System to be integrated with many different customer cloud environments.

Cloud-native architectures, built with microservices and containers, are pushing the limits of network scalability and performance, requiring networks to be much more responsive to changes in applications. Modern data center fabrics need to be synchronized with applications to remove the network as an obstacle to innovation. There needs to be a symbiotic relationship between the applications and the network that serves them.

To tackle this requirement head on, the Fabric Services System has implemented a Connect microservice that allows for integration, via a plugin infrastructure, with cloud management platforms such as OpenStack, VMware, vSphere and Kubernetes. Figure 12 shows how this process is implemented in the Fabric Services System.

Figure 12. The Connect microservice and integration into cloud management platforms

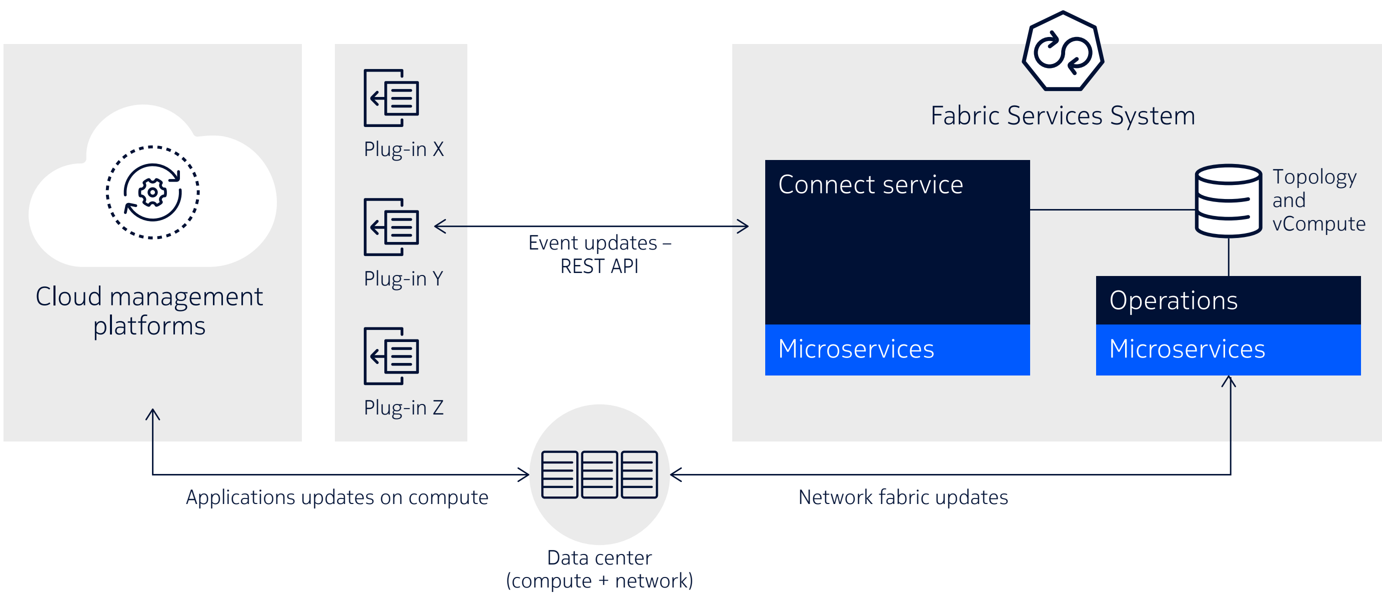

With this integration, any change events to workloads (both virtualized network functions and containerized network functions) are immediately understood by the Connect service. This allows the fabric to react in real time to these events and ensures that L2/L3 fabric connectivity always supports these changes. This type of integration is essential to scale next-generation data center networks.

Conclusion

As the demands on data center networks continue to drive openness, flexibility and agility, the Nokia Data Center Fabric solution was purpose-built to meet these challenges. As part of this solution, the Nokia Fabric Services System delivers declarative, abstract intent where automation and simplification are needed while also delivering detailed insights by monitoring every aspect of the data center fabric. This combination of abstract intent-based automation plus detailed openness and visibility allows the data center operator to perform Day 0 design, Day 1 deployment and Day 2+ configuration, operation, measurement and analysis of a data center fabric.

Learn more

To learn more about the Nokia Data Center Fabric solution:

- Visit the Nokia Fabric Services System web page
- Visit the Nokia Data Center Fabric solution web page
- Visit the Nokia Data Center Fabric Business Case Analysis web page
- Read the Nokia Data Center Fabric solution eBook
- Read the Nokia Service Router Linux product description

Read the data sheets:

- Nokia Fabric Services System
- Nokia Service Router Linux
- Nokia 7250 IXR-10/IXR-6 Interconnect Routers for SR Linux
- Nokia 7220 IXR-D series Interconnect Routers for SR Linux

Abbreviations

API	application programming interface
BGP	Border Gateway Protocol
BMP	BGP Monitoring Protocol
CD	continuous deployment
CI	continuous integration
CLI	command line interface
cSR Linux	Linux containerized Service Router
DCI	Data Center Interconnect
DevOps	development and operations
EVPN	Ethernet VPN
gNMI	gRPC Network Management Interface
gRPC	generalized Remote Procedure Call
IaC	infrastructure as code
IP	Internet Protocol
IRB	integrated routing and bridging
IS-IS	Intermediate System to Intermediate System
LAG	link aggregation
LAN	local area network
LLDP	Link Layer Discovery Protocol
MPLS	multiprotocol label switching
NDK	NetOps Development Kit
NOS	network operating system
SDN	software defined networking
VM	virtual machine
VNI	virtual network interface
VPN	virtual private network
VXLAN	Virtual Extensible LAN
WAN	wide area network
YAML	YAML Ain't Markup Language (data serialization language)
ZTP	Zero Touch Provisioning

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

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

CID:207605