Nokia Fabric Services System

Release 23

The Nokia Fabric Services System delivers scalable, intent-based automation for all phases of data center fabric operations, including Day 0 design, Day 1 deployment and Day 2+ monitoring, analysis and change.

The Nokia Fabric Services System is a key component of the Nokia Data Center Fabric solution, which also includes the Nokia Service Router Linux (SR Linux) network operating system (NOS) and the Nokia Data Center platforms.

Feature summary

- Intent-based, declarative automation across the entire operational life cycle from Day 0 design to Day 1 deployment and to Day 2+ operations, simplifying many operational tasks while reducing network risk
- “On change” streaming telemetry from SR Linux, to inform and optimize operational tasks while performing continuous intent validation
- A Digital Sandbox that emulates the network and is used to reduce both testing time and network risk
- Implements a Continuous Integration / Continuous Deployment (CI/CD) methodology embodying DevOps principles that are applied to the entire network

Overview

With the proliferation of 5G, edge clouds and cloud-based deployments, networks are becoming 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. Networks are also expected to support large increases in 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. This is the essence of NetOps.

The Nokia Fabric Services System is a complete data center fabric management, operations and automation platform. It provides an intent-based NetOps operational toolkit that addresses operators’ emerging challenges by providing the operational tools needed to thrive in the NetOps era.

The Fabric Services System implements a modern and innovative cloud-native foundation to automate all phases of network operations. From Day 0 design to Day 1 deployment and Day 2+ operations (monitor, analyze, change), all operational tasks can be executed through intent-based automation.

Any network change during any of these phases can be first validated by the Fabric Services System's fabric network emulation tool called the Digital Sandbox. This validation dramatically increases change accuracy while reducing network risk.

Figure 1 shows the operational phases of the Fabric Services System and their relationship to the production network and the Digital Sandbox.

Figure 1. Data center intent-based fabric operations with the Fabric Services System

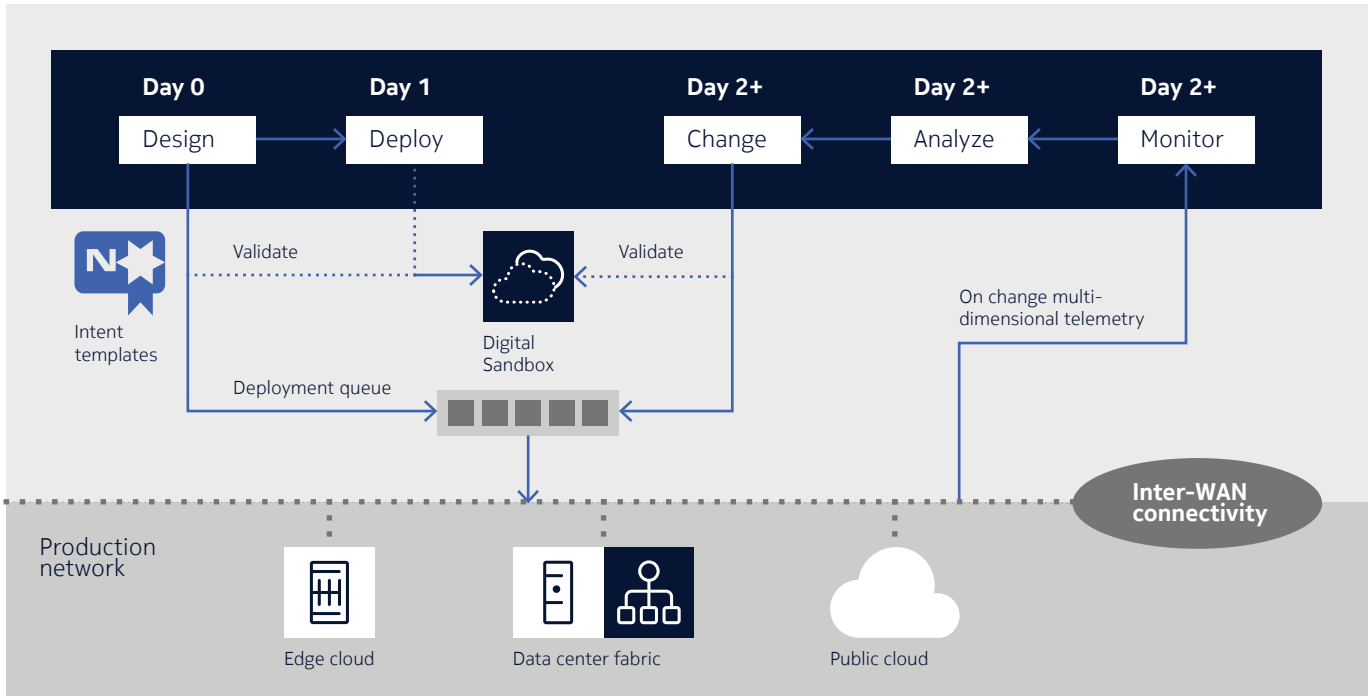

Day 0 fabric designs (e.g., number of racks, number of servers per rack, IPv4/IPv6 addressing, Border Gateway Protocol [BGP] configuration) and Day 1 network configurations (e.g., Ethernet VPN and Layer 2/Layer 3 [L2/L3] overlays) can be automated with intent-based templates that are first tested and 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.

Key aspects of the Fabric Services System

Design aspects

The Fabric Services System is purpose-built for the next generation of data center and cloud environments. Its event-driven, cloud-native microservice design is built on top of an open, extendable Kubernetes foundation.

Because the Fabric Services System relies on a Kubernetes framework, it also leverages a large set of community-stabilized, open-source projects, including Kafka, Kibana, Keycloak and ZooKeeper.

To accommodate the high-speed, low-latency messaging required to support these microservices, a high-speed messaging bus is implemented.

The Fabric Services System also provides a Digital Sandbox that is represented by a set of containerized SR Linux instances that emulate the data center fabric. End users can leverage the Digital Sandbox to trial and validate any changes planned for deployment in the production network, reducing risk to the fabric while speeding up testing and validation.

Setting up the data center fabric with intent

The 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 approach, the desired configuration and state of the fabric can be specified up front in a simplified or abstract manner and can be constantly reused to validate the network. This approach provides a strong NetOps foundation that leverages DevOps principles and fits into the bigger movement toward infrastructure as code (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 may only need to input the number of racks and the number of servers per rack. Using this information, the system auto-generates 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) that can be validated using the Digital Sandbox before being deployed to the data center fabric.

The configuration can also be customized, including the attributes generated by these design templates.

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 zero touch provisioning (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 Nokia’s Ethernet VPN (EVPN)-based implementation, available on SR Linux, to deliver multihomed Link Aggregation Groups (LAGs) as well as L2 and L3 services within and across the data center fabric.

Workload intent simplifies EVPN configuration by enabling the data center operator to focus on high-level intent. 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.

Fabric operations

After the data center fabric is designed and deployed, the Day 2+ operations phase begins. Because new workloads can still be created during this phase, workload intent can also be leveraged here. Other types of intent are also supported in this phase, including design intent and maintenance intent. These intents allow the network operator to define, in an abstract manner, the desired end state of the fabric.

Monitoring the network for intent deviations

For Day 2+ operations, the Fabric Services System constantly monitors the fabric by leveraging on change multidimensional 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 in its Digital Sandbox. If these changes pass validation, they can be automatically scheduled for deployment into the production network.

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

Multidimensional telemetry

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 through a combination of on change multidimensional streaming telemetry and log data that represents the network state and is collected directly from the data center fabric.

Multidimensional 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 Fabric Services System constantly receives this information via SR Linux's gRPC Network Management Interface (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.

Data visualization

The Fabric Services System performs data correlations to provide insights. In addition, the system interfaces with a pluggable alerting infrastructure to provide alerts to the data center operator.

Fabric integrations

The Fabric Services System provides an open REST application programming interface (API) that allows third parties to have full access to the system. A flexible, cloud-native approach enables the Fabric Service System to be integrated into many different customer cloud environments.

The Fabric Services System Connect microservice allows for integration via a plugin infrastructure with cloud management platforms such as OpenStack, VMware vSphere, Red Hat OpenShift and Kubernetes. 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 assures that L2/L3 fabric connectivity always supports these changes. This type of integration is essential to scale next-generation data center networks.

Digital Sandbox

The Digital Sandbox is a virtual network infrastructure (VNI) upon which the operator can test and validate any planned network changes. After validation, the changes can be deployed to the physical production network with reduced risk.

Digital twin

The Digital Sandbox provides a digital twin of the data center fabric, emulating the leaf and spine switches by deploying a containerized SR Linux instance of each; this allows the Digital Sandbox to build an exact replica of the fabric.

The Digital Sandbox leverages on change telemetry to maintain absolute parity of the network. It also can emulate external 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.

Validation across entire operational life cycle

Both fabric design intent and workload intent can be validated by the Digital Sandbox. However, its use extends well beyond Day 0 design and Day 1 deployment.

Although data center fabrics are built once, onboarding new workloads and making policy changes for existing workloads are constant and frequent tasks for data center operators.

These changes are all candidates to be tested and validated by the Digital Sandbox. In addition, software upgrades and configuration changes are frequent and need to be tested and validated.

Implementing the principles of DevOps CI/CD for the network

Inherited from the principles of DevOps, one of the emerging requirements for modern data centers is to implement a process that leverages principles from Continuous Integration/Continuous Deployment (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 2 shows the Digital Sandbox in action.

Figure 2. Fabric Services System Digital Sandbox

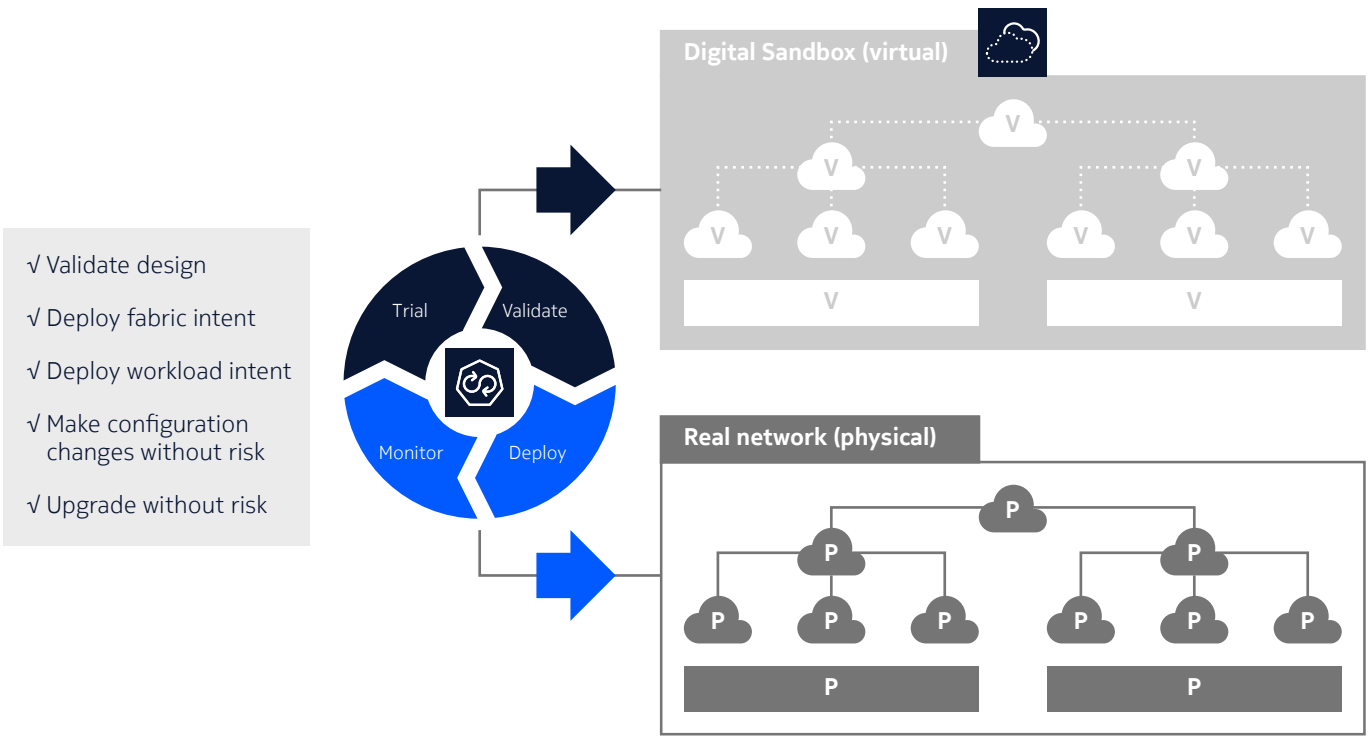

Feature and protocol support

The Nokia Fabric Services System supports, but is not limited to, the following features and protocols.

Overall platform features

Platform

- Multiservice K8s-based platform
- Label factory: Configurable interface; node associations by label; automatic system-generated labels and associations
- Role-based access control (RBAC) with users, groups and roles; services demarcation; static resource grouping
- Open API-compliant REST APIs
- Swagger documentation
- Web-based UI; browser-based notifications for alarms
- SR Linux image management
- Intent alert log; log/event dashboard
- Local IP Address Management; autonomous system number allocations
- Per-region deployment queues; commit transactions

Supported hardware platforms

- Nokia 7250 IXR-6/10 Interconnect Routers for SR Linux
- Nokia 7220 IXR-D1/D2/D2L/D3/D3L/D5 Interconnect Routers for SR Linux
- Nokia 7220 IXR-H2/H3 Interconnect Routers for SR Linux

Fabric

- Cable map; design generation from intent
- Underlay intent-deviation handling
- Certified design templates
- Flexible leaf-spine template with chassis/fixed systems
- Backbone template with chassis/fixed systems
- Optional border leaf
- Breakouts on 7220 IXR-D3 (access only)
- Dynamic Host Configuration Protocol (DHCPv4)-based ZTP (including relay)
- Static inventory: Manual import of devices without ZTP
- External Border Gateway/interior BGP (EBGP/iBGP)-based fabric
- Multi-region (multi-tenancy) support
- Brownfield import capability - Ability to import unmanaged nodes into the Fabric Services System

Workload

- L2 and L3 workloads supported by EVPN virtual routing and forwarding (VRF) instances
- Automatic VXLANv4, VNI and EVPN instance (EVI) allocation
- Supports LAG + LACP
- EVPN L2 multi-homing all active with up to 4 nodes; auto-generation of LAG IDs; Ethernet Segment Identifier (ESI)
- Provider edge/customer edge (PE-CE) BGP

Cloud Integration features

- End-to-end automation with the following cloud management systems: OpenStack, VMware vSphere, Red Hat OpenShift and Kubernetes
- Application network management from within the cloud management system for L2 and L3 networking
- Advanced use case support through integrated workflow, including BGP PE-CE, QoS, and ACLs
- Integration with OpenStack with VirtIO, SR-IOV and OVS-DPDK interfaces
- Integration with VMware vSphere with dvSwitch, port groups, and SR-IOV interfaces
- Integration with Red Hat OpenShift and Kubernetes with multiple CNI plugins and SR-IOV

Policy features

QoS

- Differentiated Services Code Point (DSCP) rewrites
- DSCP-to-Fiber Channel mapping

Access control lists (ACLs)

- L3/L4 match criteria
- Permit, deny, log
- Prefix lists via match groups

Digital Sandbox features

- Kubernetes-based fabric topologies using containerized SR Linux
- Integration with fabric and workload intent framework using uniform policies
- Design validation before deployment in production network

Operations features

- Fault management
- Node replacement policies with single or multiple nodes

Learn more

To learn more about the Data Center Fabric solution, see the web page

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

701750 (September) CID212306