Bell Labs Consulting

Data center fabric reliability study

In collaboration with Futurum

Application note

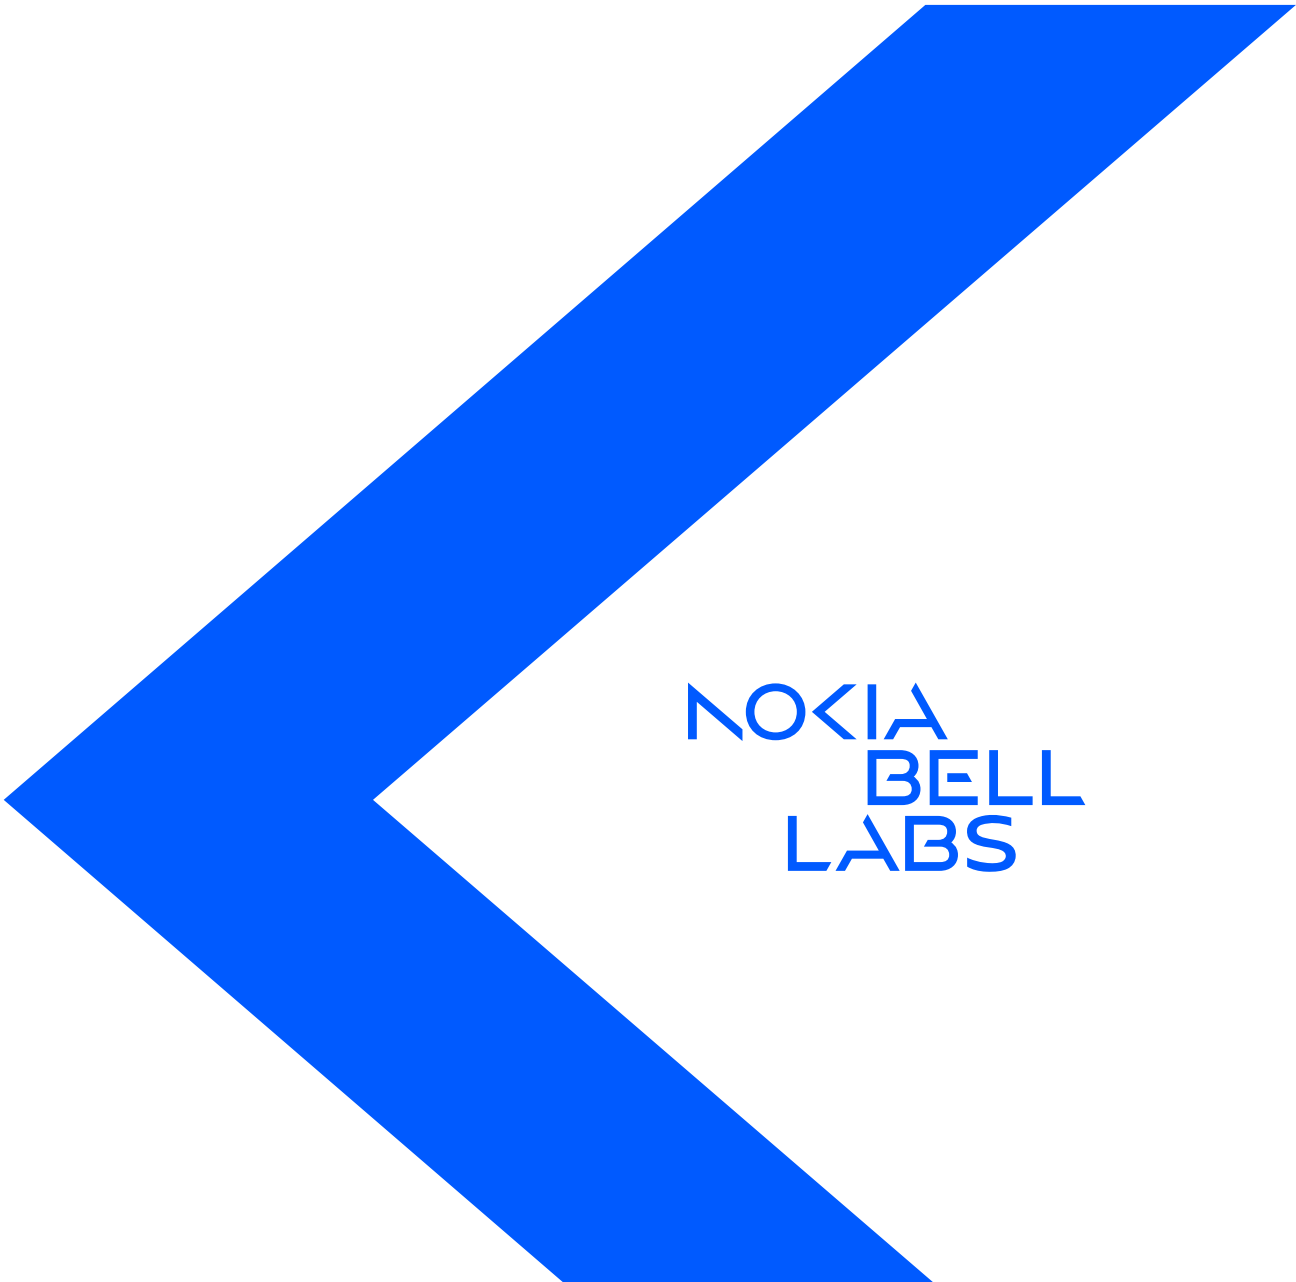

Abstract

Data centers are mission-critical hubs that must now handle intensive, artificial intelligence (AI)-driven workloads that require instantaneous processing and real-time insights. Superior network reliability is mandatory for supporting traditional and AI applications and ensuring uninterrupted digital interactions.

Nokia addresses this need for reliability with a “Human Error Zero” strategy that aims to eliminate network-level mistakes stemming from vendor software bugs or user misconfigurations. While the industry agrees on the importance of reliability, there is limited information that quantifies and compares the reliability impact of legacy versus modern data center switching solutions, including hardware, software and network operations.

The data center fabric reliability study described in this document addresses this gap with a comprehensive analysis and model that identify factors that have the most significant impact on network reliability. A collaboration between Bell Labs Consulting and Futurum, it quantifies the impact of downtime from operational and financial perspectives. The study also provides actionable guidance for enhancing reliability by adopting the Nokia Data Center Fabric solution, which delivers resilience through modernization and advanced network operations.

Contents

1 Executive summary

In today's data-driven economy, every transaction, decision and customer interaction depends on always-on networks. Data center networks must deliver unmatched reliability to support traditional enterprise applications and data-intensive AI workloads. Even brief interruptions can cause data corruption, transaction failures, lost revenue and damage to customer trust. All of this makes high availability a business-critical requirement rather than an optional feature.

1.1 Study scope and approach

The data center fabric reliability study described in this document:

- Helps quantify the reliability, operational and financial impacts of legacy switching solutions versus those of a modern, best-of-breed data center fabric.
- Uses a model developed by Bell Labs Consulting to evaluate how hardware, software, configuration and operational factors within a data center leaf–spine architecture impact network reliability.
- Compares a present mode of operation (PMO) represented by a "legacy" baseline with outdated hardware, manual processes and limited automation with a future mode of operation (FMO). The FMO is represented by a "best practice" baseline based on the Nokia Data Center Fabric solution, which uses a "Human Error Zero" strategy implemented through the Nokia SR Linux network operating system (NOS) and the Nokia Event-Driven Automation (EDA) platform to eliminate problems caused by vendor bugs and user misconfigurations.

1.2 Key findings

The study's key findings reveal four significant reliability gains with the FMO:

1 Massive improvements in network availability and downtime

The PMO achieves only 99.981736% (3.7 nines) availability with 96.1 minutes of annual downtime, while an FMO with Nokia SR Linux and EDA reaches 99.999235% (5.1 nines) availability with 4 minutes of annual downtime.

The Nokia Data Center Fabric solution delivers:

- **23.9x less or a 96% reduction** in downtime
- **Downtime of 4 minutes per year, compared to 96 minutes per year** with the legacy PMO
- **Availability of 5.1 nines, compared to 3.7 nines** with the legacy PMO

2 Significant reductions in downtime for all phases of the data center fabric operations lifecycle

The FMO with SR Linux and EDA minimizes downtime through support for advanced automation tools and capabilities including an integrated digital twin, intent-based provisioning, an event handling system (EHS), pre- and post-checks, and enhanced telemetry.

The Nokia Data Center Fabric solution delivers:

- **Up to 95% reduction** in downtime for configuration and provisioning tasks
- **Up to 99% reduction** in downtime for operations and monitoring tasks
- **Up to 99% reduction** in downtime for maintenance tasks

3 Measurable cost reductions

The study demonstrates nonlinear but substantial financial benefits. Using moderate and aggressive scenarios, the FMO with SR Linux and EDA yields major reductions in penalty and operational costs, revenue losses and reputational damage.

The Nokia Data Center fabric solution delivers:

- **Up to 60% reduction** related to penalty costs + operational costs
- **Up to 53% reduction** related to revenue loss during downtime + impacted customer churn
- **Up to 44% reduction** related to brand value degradation + churn due to brand perception

4. Real-world savings

A case study of a midsize financial institution translates these reductions to dollar savings.

If the company chose to implement the Nokia SR Linux and EDA solution, it would realize:

- **~$US94M savings** related to penalty cost
- **~$US66M savings** related to revenue loss
- **~$US207M savings** related to brand value

1.3 Operational benefits

The study reveals that the Nokia Data Center Fabric solution delivers several key operational benefits, including:

- Enabling network operations teams to meet “mean time to innocence” goals and removing the network as a blame source for outages.
- Reducing alert fatigue (being overwhelmed by sheer volume of alerts and errors) and stress which improves morale and workplace satisfaction.
- Lowering the barrier to automation and delivering faster, more predictable network operations and greater agility.

Adopting the Nokia Data Center Fabric solution with SR Linux and EDA transforms high availability from aspiration to reality, cuts human error, minimizes downtime, and generates measurable operational and financial gains.

2 Objectives

2.1 Data center applications and network reliability

Data centers now need to support a mix of workloads, including traditional enterprise applications and data-intensive AI and machine learning applications. While each of these workloads imposes specific requirements on the network, they all share one foundational requirement: network reliability.

For a recent study, Futurum[1] asked organization leaders to rate decision criteria for building their next data center network. Network reliability topped the list, with 86% of respondents identifying it as the most important factor (Figure 1).

1 “The Data Center Networking Imperative: Key Trends Driving the Next Era of Data Centers”. Futurum, 2025. https://pages.nokia.com/t0212u-the-data-center-networking-imperative-key-trends-driving-the-next-era-of-data-centers.html#single-form-section-title (accessed Oct. 21, 2025.

Figure 1: Decision criteria for building your next data center network

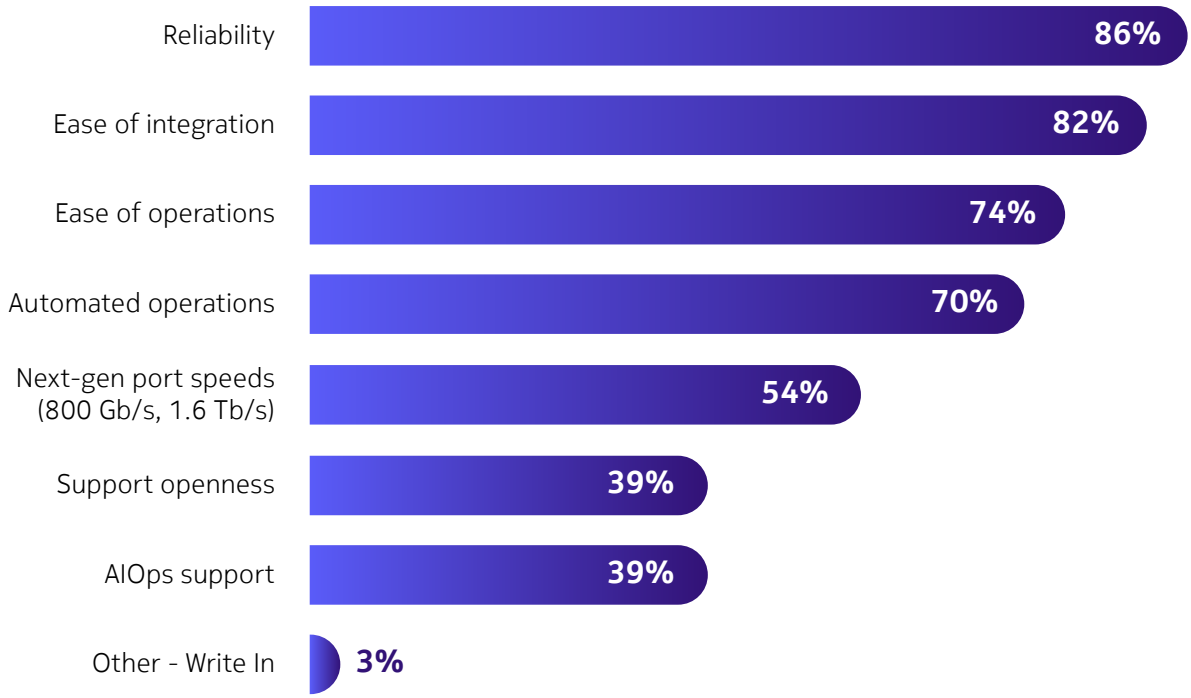

Traditional applications such as web services, databases, enterprise resource planning (ERP) systems, storage and virtualization environments require steady throughput, predictable latency and continuous access to ensure service continuity for end users and business stakeholders. The smallest network interruptions can cause data corruption and transaction failures, which result in loss of customer trust.

AI applications and workloads add stringent new requirements for data center networking. AI accelerators (specialized high-performance compute processors) need to process massive amounts of data across tens to thousands of nodes in real time. A single network fault can drastically impact the process of teaching an AI system how to learn information and make decisions and predictions. It can also waste precious compute cycles. Network reliability at this scale must deliver deterministic performance and rapid fault isolation, and prevent costly downtime.

In short, network reliability is a must-have for data center operators seeking to gain a competitive advantage.

2.2 Defining network reliability

In the data center context, the notion of reliability often refers to the entire data center infrastructure. This document focuses specifically on network reliability.

Network reliability refers to the network's ability to remain operational. For data center switching networks, reliability is about the hardware and software used on the switches, as well as the way the network is operated. Network operations tasks related to Day 0 design, Day 1 deployment and Day 2+ operations have a direct and significant impact on network reliability.

Availability is one of the most widely used measures of reliability. It represents the percentage of time a network or system is up and ready to use.

Downtime is inversely proportional to availability: Higher availability means lower downtime. For example, 99.999% availability equals roughly 5 minutes of downtime per year, while 99% availability equals roughly 3 days and 15 hours (~5,220 minutes).

Most unplanned downtime in data center networks is linked to operational causes, such as misconfigurations, manual errors, delayed fault responses, inefficient monitoring or maintenance tasks.

Data center network reliability is therefore related to switch hardware and software reliability and the way the network is designed, deployed and operated.

2.3 Reliability study objectives

A data center fabric is a network of leaf and spine switches that work together to provide a resilient and scalable infrastructure for connecting traditional and AI-based applications installed on servers. Each switch also runs NOS software that provides the intelligence, routing capabilities, telemetry and other programmable capabilities required to operate the switch. The network is typically managed by data center fabric management and automation tools or platforms.

Human Error Zero is Nokia's mission and strategy to eliminate errors from data center networks, particularly those that stem from vendor mistakes (such as software bugs) and user mistakes (such as misconfigurations). While data center network reliability is widely recognized as critical, there remains a need to quantify the degree to which different networking solutions enhance or undermine overall reliability.

The data center fabric reliability study featured in this document addresses this need by focusing solely on how reliability is influenced, measured and improved within the data center fabric. Bell Labs Consulting developed the comprehensive reliability study model, which:

- Analyzes the areas that drive reliability in data center fabrics. It models the impacts of hardware, software and, most importantly, network operations tasks on network reliability.
- Quantifies how modernized operations and intelligent automation can reduce human error and downtime.
- Provides essential guidance on how downtime reduction can translate to operational efficiency and potential financial savings.
- Validates how existing data center switching infrastructures can be enhanced by implementing a best-of-breed solution (i.e., the Nokia Data Center Fabric) that offers resilience through modernization and automation.

3 Methodology

3.1 Reliability study approach

The reliability study model compares the PMO and FMO. The PMO represents an existing solution baseline that lacks essential capabilities related to network reliability, while the FMO represents the recommended solution baseline, which provides the capabilities required to deliver enhanced network reliability.

Unless stated specifically, the results and findings featured in this document relate to comparison between a "legacy" baseline for PMO and the "best practice" baseline for FMO (Figure 2). The criteria and related capabilities for each of these baselines are described in the next section.

The best practice baseline represents the Nokia Data Center Fabric solution. The solution includes Nokia data center switches, the open, extensible and resilient Nokia SR Linux NOS, and the Nokia Event-Driven Automation (EDA) management and operations platform.

The Bell Labs Consulting model (hereafter referred to as "the model") applies two FMO scenarios to the best practice baseline: FMO 1, which is based on implementing SR Linux and FMO 2, which is based on implementing SR Linux and EDA. As depicted in Figure 2, each FMO enables a set of capabilities that help deliver enhanced network reliability.

Figure 2: Reliability study approach

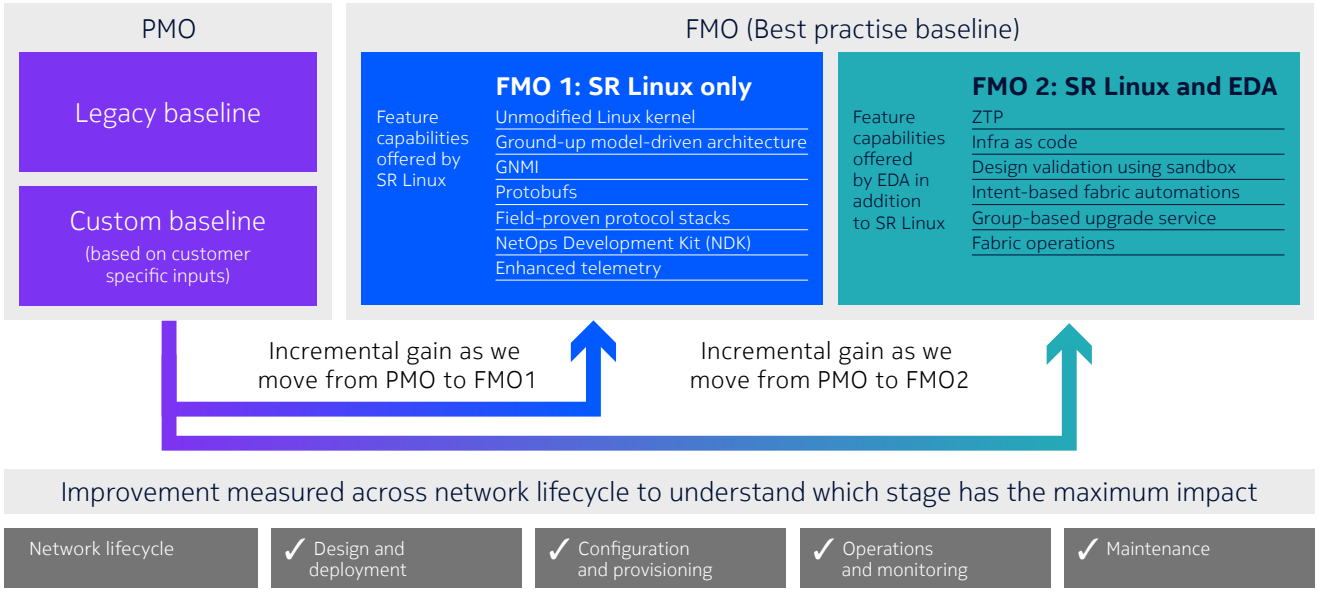

The model can use the legacy baseline or a custom baseline for the PMO. This means it can be tailored to a specific customer's current network reliability metrics based on its existing vendors' products and solutions.

The reliability study models the contributions of hardware and software and the operational impact of better tools to enhance network reliability. Unreliable network operations can result in significant downtime. As shown at the bottom of Figure 2, the model analyzes specific tasks across the network operations lifecycle that have an impact on reliability. It maps the typical data center fabric lifecycle—design, deployment, configuration, provisioning and ongoing operation (including upgrades and maintenance)—into use cases and then into the model.

This approach helps quantify how each area influences overall availability and downtime, which makes the results tangible for network operations staff. For example, if the downtime associated with a maintenance task is reduced by a certain percentage, the operations team can immediately understand the benefit. The model then compares the calculated network availability and the downtime between the PMO, FMO 1 and FMO 2 to quantify the improvement in network reliability.

3.2 Comparing the legacy and best practice baselines

Table 1 compares the legacy (PMO) and best practice (FMO) baselines. It highlights the advantages of adopting a Nokia Data Center Fabric-based FMO that delivers resilience through modernization and advanced network operations.

- The FMO leverages next-generation data center switching solutions with cutting-edge hardware, NOS and automation platforms. In contrast to the PMO's reliance on outdated architectures and protocols, the FMO supports modern technologies such as Linux, Kubernetes and cloud-native approaches.
- The FMO adopts a Human Error Zero approach, utilizing modern leaf-spine architectures with IP/EVPN/VXLAN fabrics to minimize operational errors. It prioritizes quality and ensures high reliability with switches built on the latest merchant silicon, a resilient and open NOS and a 1:1 ratio of developers to test engineers.
- The FMO provides a comprehensive management and automation toolkit, incorporating advanced features such as digital twin technology to enhance reliability, predictability and operational efficiency across the network lifecycle. The PMO lacks the capabilities required for reliable, guaranteed operations, which negatively impacts network reliability.

Table 1: Legacy (PMO) vs best practice (FMO) baselines

Criteria	Legacy	Best practice (Nokia approach)
Data center switching vendor profile	• Vendors that failed to keep up and enhance their product and solution offerings • Older switch architectures and protocol mechanisms (e.g., switch stacking) • Port speeds of 10 Gb/s and below, in need of upgrades to 100 Gb/s and beyond • History of applying temporary approaches and fixes with existing products	• A next-generation data center switching solution with leading-edge innovation in hardware, NOS and automation platforms • Based on modern technologies and innovations (e.g., Linux, cloud native, microservices, digital twins, Kubernetes) to implement a truly modern data center switching solution • Port speeds up to 800 Gb/s and beyond in support of traditional and AI applications and use cases • A ground-up approach to address the challenges faced by data center networking teams and operations staff
Design approach and philosophy	• Early origins in layer 2-centric designs and spanning tree approaches with primary focus on the lower and middle ranges of the data center switching market • Little or no consideration or adoption of layer 3 data center network topologies • Mix of disparate data center networking solution components from acquisitions and spin-ins that lead to complex management and operations	• A Human Error Zero approach that aims to eliminate all human errors caused by vendor product issues and network operations • Built for modern leaf–spine architectures that implement IP/EVPN/VXLAN based fabrics • A fresh start helped align design approach to applying modern technology approaches and innovations

Criteria	Legacy	Best practice (Nokia approach)
Hardware and software quality and reliability	• Not a core part of the go-to-market approach and philosophy • Reliance on early generations of proprietary and merchant silicon in data center switches • NOS utilized is near or beyond end-of-life • Poor to average track record of product reliability • Customers express frustration related to quality issues • Typical developer-to-test-engineer ratio of 2.5–3.5:1	• Quality-first approach to system design where reliability is not an afterthought • Switches based on the latest generation of merchant silicon • A new NOS that is truly open, extensible and resilient • 20-plus years of expertise in building business- and mission-critical hardware and software products • Customer positivity rating of ~95%; sales and marketing customer positivity rating of ~98%[2] • Software design and test teams work collaboratively and are staffed appropriately with a 1:1 ratio of developers to test engineers
Network operations capabilities that deliver reliability	• Significant reliance on CLI for provisioning, monitoring, incident troubleshooting and all operations • Limited NOS alert capabilities (e.g., SNMP only, no streaming telemetry support) • Lack of management solutions that support essential automation capabilities • Limited and/or proprietary automation tooling that is overcomplicated, closed or not user friendly	• Offers a modern management and operations platform • Includes several built-in and unique features (e.g., digital twin, event handling system) specifically designed to improve reliability and predictability • Provides a comprehensive management and automation toolkit for all phases of the network operations lifecycle • Maintains an open solution architecture to leverage best-of-breed open-source tooling

3.3 Model approach

The network reliability study analyzes the impact of the best-effort baseline SR Linux (FMO 1) and SR Linux with EDA (FMO 2) solutions on improving end-to-end data center fabric reliability compared to the legacy PMO.

The model uses a simplified leaf–spine architecture (Figure 3) with a pair of spine switches and a pair of leaf switches, each in a 1+1 active–active redundant setup. This setup provides load balancing and fast failover. Reliability targets are set at business-critical availability (~99.999%) for the fabric.

2 Nokia Customer Survey Program report (IP Networks business unit)

Figure 3: Reference configuration applied to model

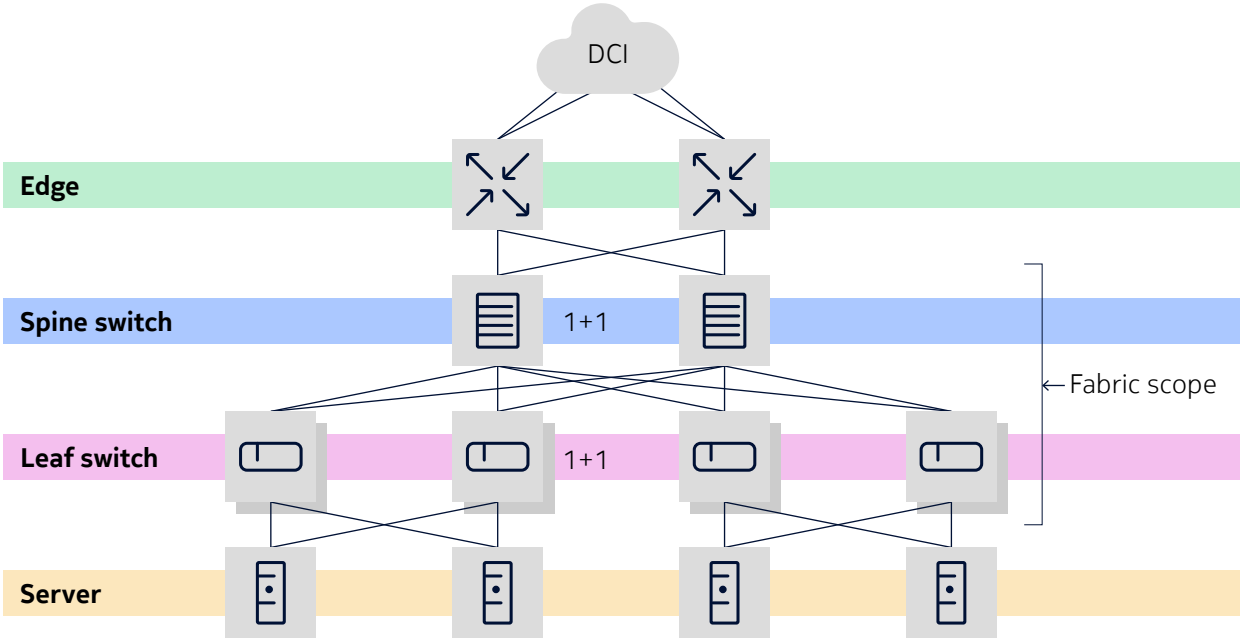

Reliability benefits are quantified and compared using availability and downtime values. Reliability is calculated for the data center fabric (leaf, spine layers and their network connectivity), excluding edge routers and broader data center components. The reliability study also translates reliability improvements into business and financial impacts. This makes it easy to see the performance difference between the legacy PMO and the two FMO options.

3.4 Model parameters and assumptions

The data center fabric reliability model (Figure 4) represents each switch as three distinct components—hardware, software (NOS) and configuration—each susceptible to independent failure. For example, a software bug or an erroneous routing table can impair a fully operational switch.

The model includes active and backup switches in a 1+1 active–active redundant configuration for spine and leaf switches. It also accounts for common cause failures (CCFs), where software bugs or configuration errors such as faulty routing tables simultaneously affect the active and backup switches.

Additionally, the model accepts input data (the downtime frequency and the duration) and transforms them to mean time to failure (MTTF), the average time expected between failures, and mean time to restore (MTTR), the average time it takes to restore service after a failure occurs.

Figure 4: Model parameters for calculating network reliability

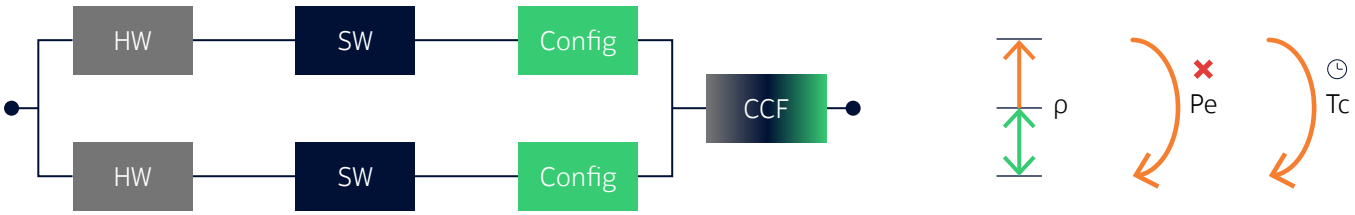

Hardware
MTTFh: Data collected over a period of time across vendors: tens of years, decreased by poor LCM
MTTRh: Typically 1 (on-site) or 4 (off-site), increased by poor operations

ρ: Overutilization
Pe: Protection error probability
Tc: Failover delays

Software
MTTFs: Upgrades and bugs, typically 1-4/year
MTTRs: Upgrades and bugs, 10 minutes or hours (bugs), increased by poor operation

Configuration
MTTFc: Operator/NEM dependent
MTTRc: Operator/NEM dependent

MTTF: Mean Time to Failure
MTTR: Mean Time to Restore

Common Cause Failures (either HW, SW or Config)
MTTFcc: Operator/NEM dependent
MTTRcc: Operator/NEM dependent

The model also incorporates three operational factors that impact reliability:

1. Overutilization, where loading switches beyond 50% in an active–active setup prevents full traffic recovery during failover.
2. Protection error probability, where misconfigurations cause the backup switch to malfunction.
3. Failover delay, where non-zero failover times, even a few seconds, caused by reinitialization or reauthentication reduce availability, particularly for high-availability components such as switches.

All these factors—hardware, software, configuration and operations—are modeled for spine and leaf switches. Hardware failure rates are sourced from the Nokia Systems and Network Reliability Engineering team. Software, configuration and operational parameters are derived from anonymized industry data and best practices observed over the past decade. This ensures a comprehensive assessment of network reliability for the data center fabric.

The model then compares the calculated network availability and downtime for the PMO, FMO 1 and FMO 2 to quantify the improvement in network reliability. It also provides a breakdown of the resulting improvement across the four network lifecycle phases.

FMO 1 measures how SR Linux alone improves end-to-end data center fabric reliability versus the PMO. FMO 2 adds EDA to SR Linux and quantifies the extra reliability gain versus the PMO.

Nokia asserts that the best practices baseline for FMO will lead to a rethinking of the way data center fabrics are designed, deployed and operated. By driving the Human Error Zero vision and embedding reliability into the design of data center networks, this approach significantly reduces the frequency of outages. Reliability and uptime improve dramatically, turning high availability from aspiration to reality. Advanced network operations tools cut manual effort, lower the chances of human error, and make operations tasks more predictable and reliable. For these reasons, the model assumes and applies a ~90% gain in operational reliability for the FMO.

3.5 Modeling network operations lifecycle use cases

The data center network lifecycle includes Day 0 design, Day 1 deployment and Day 2+ operations phases. Each phase includes network operations, job functions and tasks.

- Day 0 involves planning, architecture design and validation to ensure the network meets intent before physical implementation.
- Day 1 involves initial rollout, device provisioning, configuration activation, zero-touch provisioning (ZTP) and interoperability in multivendor setups.
- Day 2+ involves ongoing day-to-day operations, real-time monitoring, issue detection and adaptive responses to maintain performance.

The reliability study model analyzes network operation tasks that affect reliability, mapping the typical data center fabric lifecycle to specific use cases and then to the model.

The model applies the following uses cases to analyze and quantify how network downtime is impacted by the absence or presence of critical capabilities associated with supporting each use case. It evaluates data center fabric reliability by comparing capabilities related to these uses cases within the legacy and best practice modes of operation.

Digital twin for design and deployment

An integrated digital twin is a continuously synchronized virtual replica of the live network. It can simulate changes, failures and performance conditions without impacting real traffic, thereby boosting reliability, reducing human error and increasing operational confidence.

A network digital twin helps data center teams simplify design and deployment by enabling capabilities such as:

- Topology modeling and validation: Build spine–leaf or super-spine–leaf designs and verify scalability, latency and redundancy before the hardware arrives.
- Policy and routing verification: Simulate Ethernet VPN (EVPN), Border Gateway Protocol (BGP), quality of service (QoS) and security policies in normal and failover conditions.
- Resilience testing: Inject link failures, node reboots or partitions to confirm equal-cost multi-path (ECMP) routing, fast reroute and other redundancy mechanisms.
- Accelerated rollout: Once validated, push templates and configurations directly from the digital twin to production for a first-time-right deployment.

Once the digital twin has been validated, the same configuration and templates can be automatically pushed to production, ensuring consistency between the designed and deployed networks. This shortens deployment times and enables teams to identify and address issues early, leading to improved network reliability.

Digital twin for configuration and provisioning

A network digital twin helps data center teams with configuration and provisioning by enabling capabilities such as:

- Pre-deployment validation: Test new configurations in the twin before applying them to the production environment.
- Change impact simulation: Visualize how routing tables, paths or services will change with new configurations or firmware upgrades.

- Automated compliance: Continuously compare the intended versus actual state to detect drift, misconfigurations or security violations.
- Infrastructure as code (IaC) and continuous integration/continuous deployment (CI/CD) integration: Embed the twin in automation pipelines so every change is automatically tested before rollout.

In this case, the digital twin minimizes the risk of misconfigurations, reduces change-related outages and enforces consistency across the fabric.

Event handling system for operations and monitoring

An EHS detects, interprets and automatically reacts to network events that could impact performance or availability, preventing minor issues from becoming outages. It continuously receives real-time telemetry data (link utilization, packet drops, latency, interface errors) and automatically triggers predefined actions when a threshold breach or anomaly is identified (e.g., port saturation, link flap, high packet loss).

The EHS prevents service degradation by reacting instantly to early warning signs, keeping latency and loss within service-level agreement (SLA) limits. It reduces downtime through faster detection and repair with minimal manual effort. It also improves operational consistency by using repeatable, automated responses instead of variable human decisions.

Upgrade and pre-traffic drain for maintenance

The upgrade and pre-traffic-drain capability automates and orchestrates planned maintenance and upgrades so service stays up and the network remains reliable, eliminating the typical risks of packet loss, session drops or configuration drift.

This capability automatically reroutes live traffic away from the target device or interface before any work begins. The drained node is placed in a controlled state that isolates it from the data plane while maintaining control plane visibility, preventing any impact on neighboring devices. A centralized, automated workflow locks the device only when it's ready to reboot, cutting hold times and shortening the maintenance window. Additionally, if every upgrade step can be simulated in the integrated network digital twin, it identifies incompatibilities or failures ahead of the live change.

The upgrade and pre-traffic-drain capability can enable zero-loss maintenance (i.e., no packet drops or session resets during upgrades). It reduces downtime by providing faster and more predictable windows and quicker service restoration. When combined with a digital twin it enables safe, predictable changes where software and configuration updates work as expected.

4 Results and findings

4.1 Availability improvements and downtime reduction

Figure 5 illustrates the availability metrics for a network fabric using three configurations: PMO, SR Linux (FMO 1) and SR Linux and EDA (FMO 2). The PMO is based on the legacy baseline while FMO 1 and FMO 2 are based on the best practice baseline. (For reference, Table 1 provides a comparison between these baselines.)

The availability metrics in Figure 5 correspond to a pair of spines and a pair of leaf switches. In the leftmost column, the PMO configuration shows the spine and leaf switches achieving an availability of 99.99%, which corresponds to approximately 48 minutes of downtime per year for each at a device level. When combined at the fabric level, the multiplicative effect of spine and leaf availability results in 99.981736% (3.7 nines) availability, equating to about 96.1 minutes of cumulative annual downtime.

Figure 5: Availability improvements and downtime reduction

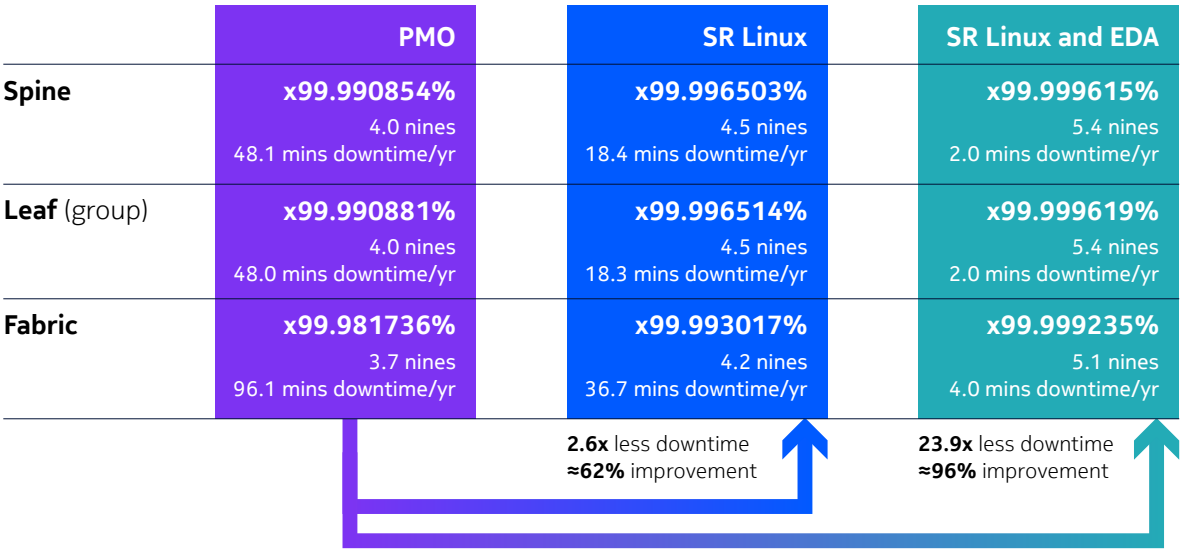

Availability improves significantly with the SR Linux solution. Both spine and leaf switches reach 99.9965% availability, which reduces downtime to roughly 18.3–18.4 minutes per year—a 2.6x reduction compared with the PMO. At the fabric level, this translates to 99.993017% (4.2 nines) availability, or approximately 36.7 minutes of cumulative downtime annually.

The SR Linux and EDA configuration further enhances reliability, achieving 99.9996% availability for both spine and leaf switches, with downtime reduced to about 2 minutes per year per device. At the fabric level, this results in 99.999235% (5.1 nines) availability, which is equivalent to just 4 minutes of cumulative downtime annually. This represents a 23.9x reduction in downtime compared with the PMO baseline, showcasing the substantial reliability gains with SR Linux and EDA.

FMO delivers massive improvements in availability and downtime

- 23.9x less downtime, a 96% reduction
- Downtime of 4 minutes per year, compared to 96 minutes per year with the legacy PMO
- Availability of 99.999235%, compared to 99.981736% with the legacy PMO

4.2 PMO availability is inadequate for critical networks

An availability of 99.99% for spine and leaf routers in the PMO configuration might appear robust. However, this equates to a downtime of ~48 minutes per year. Also, individual availability numbers must account for the multiplicative effect of this availability across the entire system, or, in this case, across the leaf–spine layers.

In the PMO network example, the spine and leaf switches can be likened to a track and train that must work together. While each switch may appear individually reliable, a 99.99% chance of working equates to a downtime of ~48 minutes per year. When the switches are combined, the chance of the whole system working drops to 99.981736%, which equates to a significant downtime of ~96 minutes per year.

This multiplicative impact shows substantial system-level downtime in a fabric, where the spine and leaf switches must function together.

This effect worsens as other critical elements such as edge routers, servers, power supplies and air conditioning are factored in because each has its own downtime risks. Assuming each additional component also has ~99.99% availability, the multiplicative impact further degrades system-wide availability, potentially pushing downtime to several hours annually. This demonstrates that the availability of the PMO is inadequate for business- and mission-critical environments.

4.3 Downtime reduction for operations use cases

Figure 6 contains two charts that compare downtime and its reduction across network operations use cases for the PMO SR Linux, and SR Linux and EDA configurations. These use cases are mapped to the design and deployment, configuration and provisioning, operation and monitoring, and maintenance phases of the lifecycle.

The lefthand chart shows downtime in minutes per year, while the righthand chart highlights the percentage reduction in downtime for SR Linux (blue) and SR Linux and EDA (teal) relative to PMO.

For PMO, downtime is highest during configuration and provisioning (~35 minutes per year) and operation and monitoring (~30 minutes per year), with lower values for design and deployment (~10 minutes per year) and maintenance (~20 minutes per year).

SR Linux significantly reduces downtime across all phases. SR Linux and EDA together achieve near-zero downtime, delivering superior reliability.

Figure 6: Downtime reduction for operations use cases

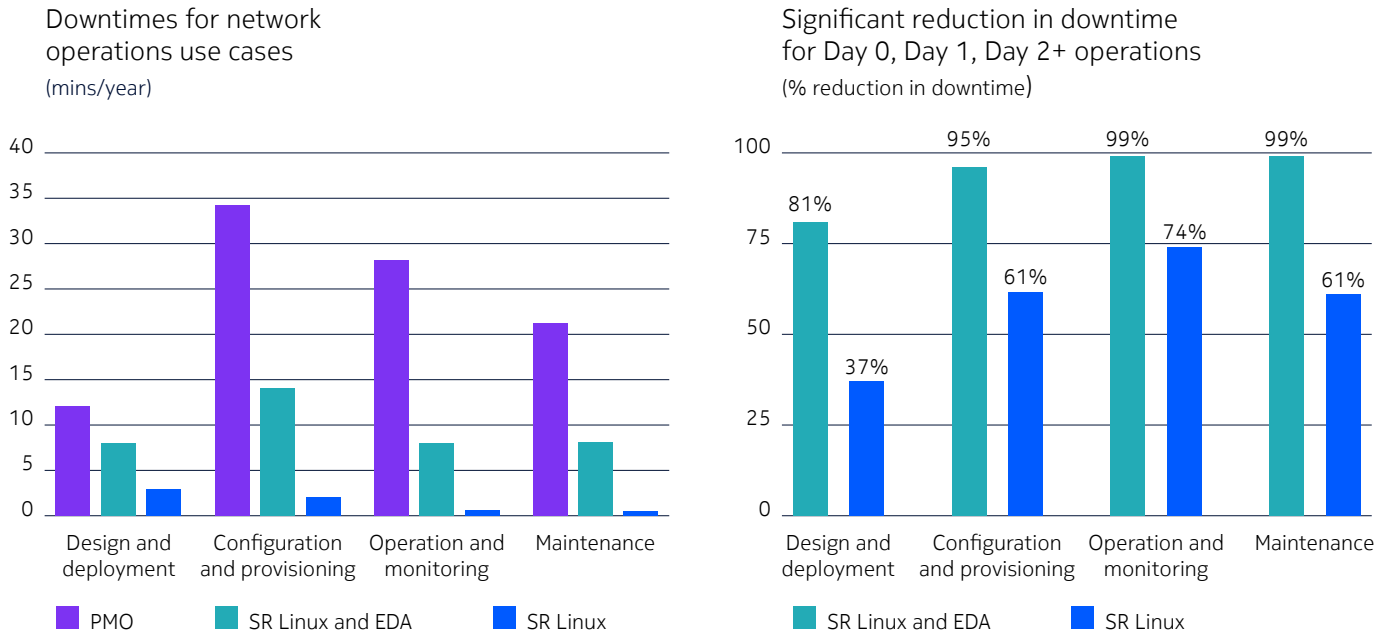

The righthand chart quantifies these improvements further: The SR Linux and EDA FMO cuts downtime by 81% for design and deployment, 95% for configuration and provisioning, 99% for operation and monitoring, and 99% for maintenance. This consistent improvement across the network lifecycle underscores the benefits of SR Linux and EDA.

Table 2 compares the reliability improvements by operations lifecycle phases for the PMO and FMO (Nokia SR Linux and EDA scenario). It also summarizes the capabilities supported by the FMO scenario to help realize these improvements.

Table 2: Reliability improvements by operations lifecycle phase

Lifecycle phase	PMO	FMO (Nokia SR Linux and EDA)
Design and deployment	Manual, static design using templates and diagrams leads to configuration drift. Production networks diverge from intended design, causing unexpected behavior and outages.	• EDA digital twin: A like-for-like virtual replica of the production fabric enables full design validation using real intent inputs before deployment. • Intent-based modeling: Engineers define abstract network states (e.g., topology, policies, failover behavior) that are automatically translated into validated configurations. • No translation errors: Configurations tested in the digital twin are deployed directly to production—no copy/paste or manual edits.
Configuration and provisioning	Manual per-device scripting, inconsistent templates and human interpretation cause configuration errors—a top contributor to downtime.	• Reliability gains: Reliability improves significantly when the design phase (planning and modeling) and the deployment phase (configuration and provisioning in production) are tightly aligned and mirror each other with high precision. • Dry-run validation before every change: EDA automatically checks syntax, policy conflicts and topology consistency. • Modular, model-driven architecture (SR Linux): Incremental, non-disruptive updates via open APIs. • Group-based provisioning: EDA supports parallel, automated configuration across fabric nodes. • Closed-loop feedback: Post-deployment telemetry ensures ongoing alignment with intent.
Operations and monitoring	Monitoring is reactive, so engineers manually correlate logs, respond to alerts and fix issues after degradation begins.	• Event handling system (EHS): The EHS provides pre-programmed, real-time automated responses to anomalies (e.g., port saturation, packet drops). • Proactive correction: The system pushes routing policies, rebalances traffic or disables failing interfaces before impact. • Closed-loop automation: The system supports continuous telemetry plus detection of patterns that can help trigger remediation.
Maintenance	Manual traffic drain, device isolation and long upgrade windows introduce failover delays, protection errors and service disruption.	• Graceful traffic drain + maintenance mode: Devices are isolated with zero packet loss. • Minimized upgrade window: The device remains operational until it is ready to reboot. • Pre-tested in digital twin: All upgrade steps are validated in a production-identical environment. • Group-based upgrades: The system supports parallel, automated updates across the fabric.

The FMO delivers significant reductions in downtime for all phases of the data center fabric operations lifecycles

- Up to 95% reduction for configuration and provisioning tasks
- Up to 99% reduction for operations and monitoring tasks
- Up to 99% reduction for maintenance tasks

4.4 Areas that drive reliability in data center fabrics

Figure 7 compares the downtime in minutes per year for the PMO, FMO 1 and FMO 2 configurations for five types of failure root causes that impact reliability in data center fabrics: hardware failures, firmware upgrades, configuration failures, CCFs and protection failures.

Figure 7: Areas that drive reliability in data center fabrics

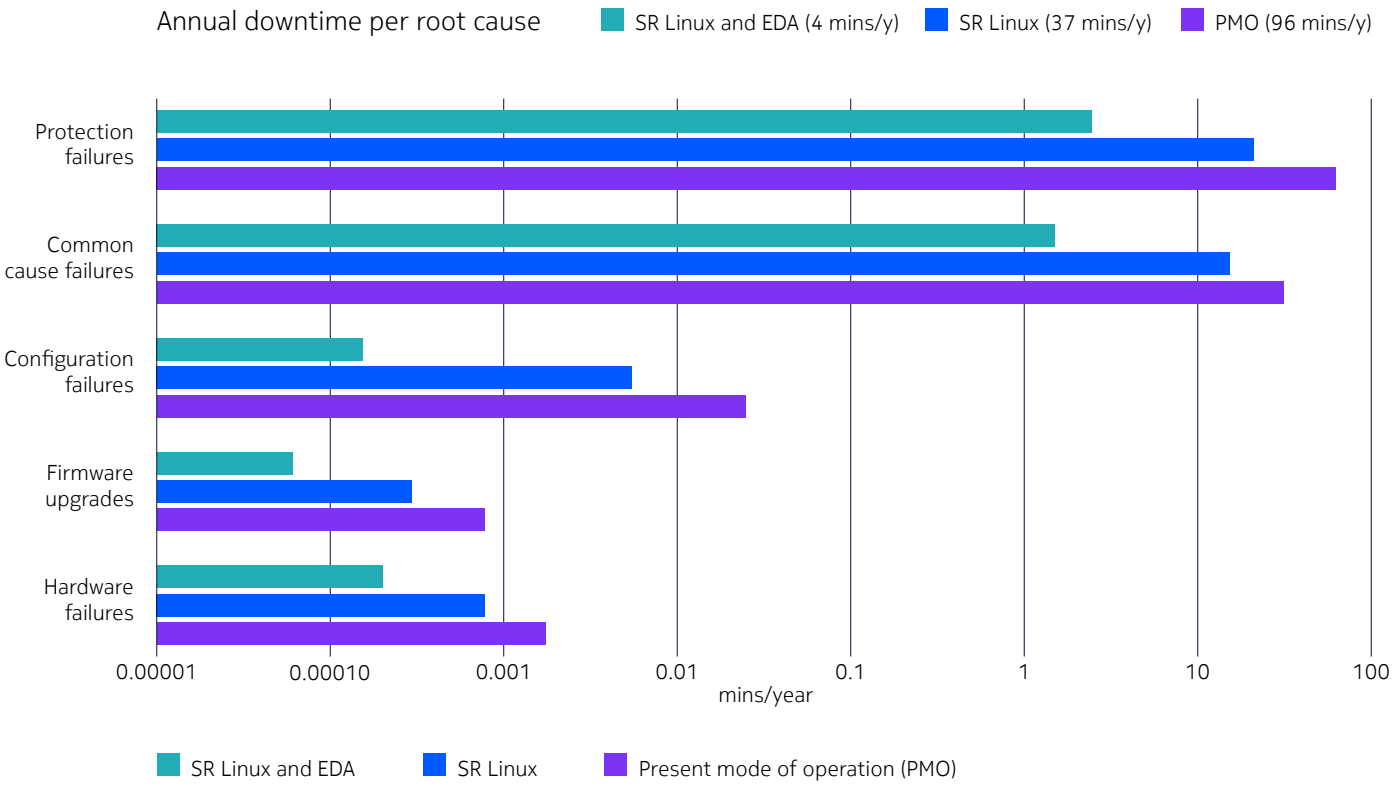

- Hardware failures are caused by physical defects in switch components such as power supplies, ports and processors. These failures stem mainly from manufacturing quality and component lifespan, not from operational practices.
- Firmware upgrades can introduce downtime because of reboot cycles, compatibility issues or upgrade errors. The downtime duration and frequency depend on planning and testing. Poor preparation raises the risk of failure.

- Configuration failures in routing tables, VLANs or policy settings can impair a fully operational switch and cause downtime. Manual mistakes or a lack of validation increase risk, while tools such as ZTP help mitigate risk.
- A CCF involves a single fault (e.g., a software bug or misconfigured routing table) that simultaneously affects the active and backup switches, defeating redundancy and causing fabric-wide downtime.
- Protection failures occur when the backup switch does not take over cleanly in a 1+1 active–active redundancy setup. Causes may include overutilization (> 50% load) that prevents full traffic recovery, misconfigurations or bugs in the backup switch, or failover delays (reinitialization or reauthentication).

Except for hardware failures, all other root causes relate to operations. Network operations play a critical role in downtime, as most failures (except hardware) stem from operational factors such as configuration errors, CCFs and protection issues.

The model compares the legacy PMO with FMO 1 (SR Linux alone) and FMO 2 (SR Linux and EDA), showing up to 90% operational reliability gains for FMO because of reduced human error and advanced tools.

Hardware failure rates come from the Nokia Systems and Network Reliability Engineering team. Based on historical data and reliability analysis, the FMO improvement for hardware failure rates is capped at 20%. An improvement of up to 90% is seen for operational parameters—overutilization, protection errors and failover delays—because of enhanced operational practices enabled by the Nokia Data Center Fabric solution used for the FMO setups.

4.5 Enhanced reliability with a quality-first approach

The Nokia data center switching solution delivers leading-edge innovation through modern hardware, a Linux-based open and extensible NOS, and a cloud-native automation platform that leverages microservices, digital twins, and Kubernetes. This next-generation solution supports port speeds up to 800 Gb/s and beyond and is optimized for traditional and AI workloads.

Designed from the ground up for leaf-spine IP/EVPN/VXLAN fabrics, the Nokia solution adopts a Human Error Zero philosophy to eliminate errors from vendor products and operations. Reliability is embedded from the start with a quality-first approach, using latest-generation merchant silicon and a no-compromises approach to developing highly reliable software and systems.

Backed by over 20 years of expertise in mission-critical systems, Nokia ensures rigorous development with a 1:1 developer-to-test-engineer ratio (compared to a 2.5-to-3.5:1 ratio for the rest of the industry).

Nokia frequently collects customer feedback as part of its commitment to quality. Two key metrics are the percentage of positive responses across products and services, and sales and marketing. Nokia has achieved 95% and 98% positivity, respectively for these two metrics.

4.6 How Nokia SR Linux improves reliability

SR Linux boosts network reliability through a combination of open-source foundations and purpose-built features. By running an unmodified Linux kernel, SR Linux can receive vendor-agnostic security patches and leverage the extensive testing performed by the global Linux community, reducing exposure to known vulnerabilities.

- A ground-up, model-driven architecture eliminates translation layers, cutting complexity and the risk of configuration or transformation errors.
- The single gNMI API for set/get operations and streaming telemetry further simplifies operations, minimizing misconfigurations and human error.

- All inter-process communications use protobufs, providing language-agnostic, structured messaging that delivers a roughly 2x performance gain and makes extensions easier to develop.
- The platform also incorporates field-proven protocol stacks and a programmable CLI that lower vendor lock in, enable custom network applications and allow fabric-wide orchestration through the Nokia EDA platform. These capabilities enhance automation and stability.
- The EHS in SR Linux is capable of performing preprogrammed actions based on port saturation or packet drops. It automatically triggers corrective actions to prevent network degradation and mitigate issues (e.g., by switching ports or pushing a routing policy).
- Enhanced, granular telemetry offers high-performance, programmable monitoring that enables early detection of anomalies before they impact service.

Together, these capabilities create a more robust, stable and easily managed data center fabric, which directly translates into higher availability and reduced downtime.

For a complete list of Nokia SR Linux value-add capabilities and how they improve data center fabric reliability, please refer to Table A1.

4.7 How Nokia EDA improves reliability

The combination of Nokia SR Linux and EDA dramatically boosts network reliability by automating and validating every stage of the fabric lifecycle. EDA provides several key capabilities that improve reliability in data center networks:

- EDA supports pre-checks and post-checks as core validation mechanisms. This ensures safe and predictable changes by simulating and verifying configurations before and after deployment. Pre-checks help identify potential issues before applying changes to the live network. Post-checks help confirm the desired network configuration and state, detecting partial failures and preventing partial deployments by adopting a feature called atomic transactions, which enables an “all-or-nothing” approach to making configuration changes.
- A core component of EDA is its integrated digital twin—a virtual, containerized replica of the production network. Before any change reaches the live network, the EDA digital twin enables full design validation. This cuts CAPEX/OPEX risk by 80%–90% and prevents configuration-related failures.
- ZTP abstracts support intent-based onboarding and handle secure key generation and distribution, eliminating manual configuration errors and the downtime they cause.
- Treating the infrastructure as code lets operators define declarative fabric intents that are repeatable, version-controlled and easy to compare, while leveraging familiar DevOps tools and talent pools.
- Intent-based fabric automation builds on ZTP to deliver zero-human-error deployments that ensure consistent, error-free rollouts. The group-based upgrade service reduces interruptions by orchestrating parallel, automated switch upgrades within a minimal maintenance window.
- The EHS in SR Linux is capable of preprogrammed actions based on port saturation or packet drops. It automatically triggers corrective actions to prevent network degradation and mitigate issues (e.g., by switching ports or pushing a routing policy).
- The EDA platform’s unique operational visibility provides precise capacity planning and rapid root-cause analysis, which help prevent over-provisioning and accelerate fault remediation. Together, these capabilities create a highly automated, repeatable and observable fabric that minimizes human error, reduces downtime and enhances overall network availability.

For a complete list of Nokia EDA value-add capabilities and how they improve data center fabric reliability, please refer to Table A2.

4.8 Financial assessment

Financial impact analysis

The three bar charts in Figure 8 illustrate the potential financial benefits of implementing Nokia SR Linux and, SR Linux and EDA to reduce costs associated with network downtime in data center fabrics. This analysis is based on industry-validated assumptions and incorporates key business metrics such as:

- Penalty costs, including regulatory fines or compliance penalties.
- Operating costs incurred in recovering from outages.
- Revenue losses from unavailable services or transactions.
- Customer churn resulting from downtime.
- Reputational loss quantified through brand value degradation and delayed effects such as challenges in customer acquisition.

Figure 8: Financial benefit assessment – Cost reductions

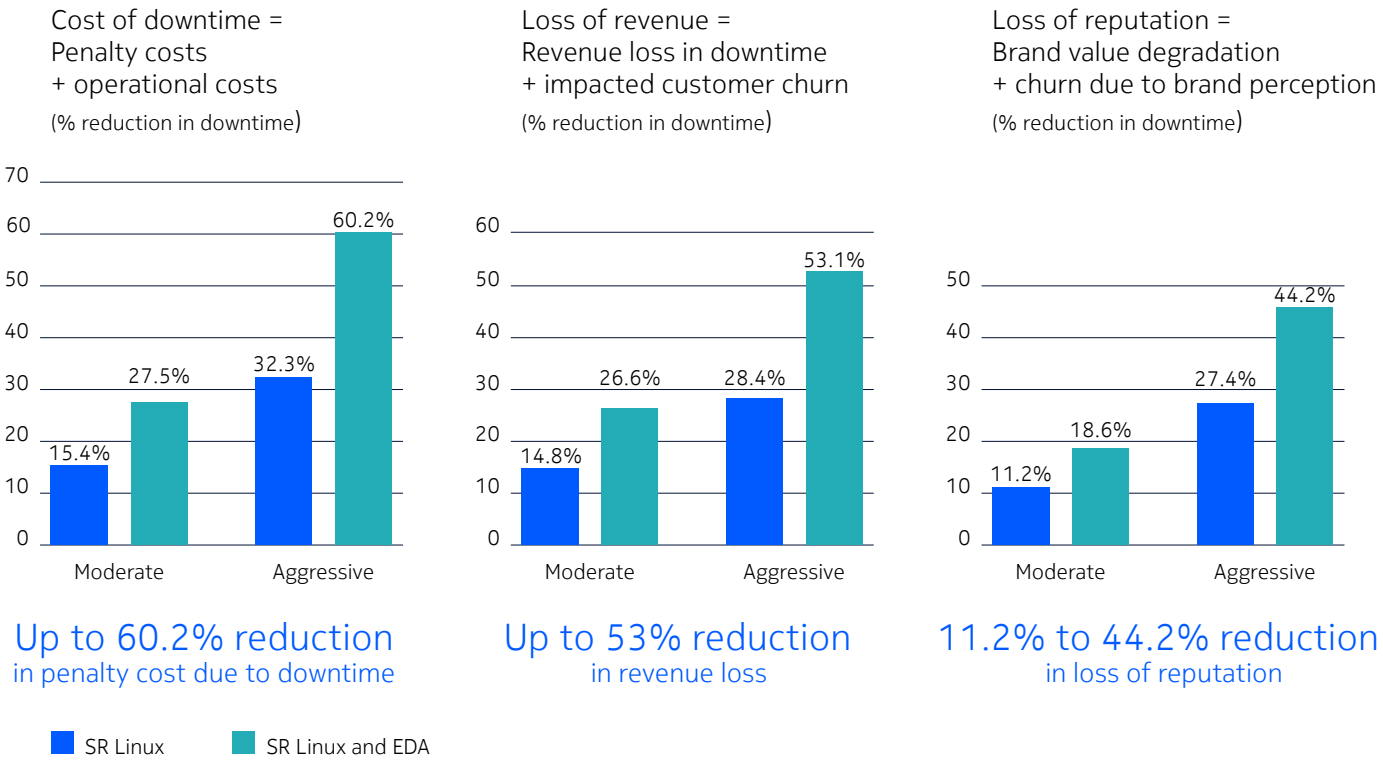

The charts consider sector-specific variations—such as e-commerce, healthcare or finance—where the impacts differ because of multiple factors but may still result in penalties. These include the thresholds defined, whether the downtime occurs during peak hours or non-peak hours, the ripple effect on consumer services, regulatory requirements and transaction sensitivity.

To capture this range, Figure 8 depicts two scenarios: “moderate” (conservative assumptions) and “aggressive” (more optimistic, high-impact cases). The relationship between downtime reduction and financial savings is nonlinear. For instance, while the SR Linux and EDA deployment achieves a downtime improvement of up to 96%, financial reductions are lower because of threshold-based regulations, service interdependencies and cascading effects that amplify even brief outages.

The cost of downtime (penalty costs + operational costs) chart shows percentage reductions in absolute value. In the moderate scenario, SR Linux yields a reduction of about 15%, while the combination of SR Linux and EDA achieves around 27%. In the aggressive scenario, the reductions increase to approximately 32% for SR Linux and 60% for SR Linux and EDA, a potential 60.2% overall reduction in penalty costs due to minimized downtime.

In the loss of revenue (revenue loss during downtime + impacted customer churn) chart, the moderate scenario shows a roughly 15% reduction for SR Linux and 27% for SR Linux and EDA. These reductions increase to 28% and 53% respectively in the aggressive scenario, demonstrating a potential 53% reduction in revenue loss risk through enhanced reliability.

The third chart quantifies loss of reputation (brand value degradation + churn due to brand perception), where moderate scenario reductions are about 11% for SR Linux and 18% for SR Linux and EDA. These reductions increasing to 27% and 44% in the Aggressive case, resulting in 11.2%–44.2% mitigation of reputational financial impacts.

SR Linux and EDA provide significant cost reductions

- Up to 60% related to penalty costs + operational costs
- Up to 53% related to revenue loss during downtime + impacted customer churn
- Up to 44% related to brand value degradation + churn due to brand perception

Quantifying savings: Impact on a midsize financial institution

Figure 9 presents a case study that details the financial benefits of transitioning from a legacy data center solution to Nokia SR Linux and EDA. The study is based on 2024 data for a real-world midsize financial institution.

This sector was chosen because of its critical reliance on service availability, driven by strict regulatory requirements, direct consumer interactions, and the immediate impact of network connectivity on financial transactions, revenue generation and brand reputation.

The company’s profile includes a market cap of $US9.4 billion, a brand value of $US0.47 billion, 500,000 depository customers (households/businesses), 30,000 health savings account (HSA) bank business customers and annual revenue of $US4.179 billion. It is supported by two data center sites, 1,200 full-time equivalent (FTE) IT workers, 800 applications and an annual TCO including downtime of $US156 million. The data centers serve 3,530,000 annual customers (retail and HSA accounts).

The company’s current data center solution is at par with the legacy (PMO) baseline. Our assessment shows that if the company chooses to migrate from its current solution to a Nokia solution with SR Linux and EDA, its network availability will improve substantially from 3.7 nines to 5.1 nines.

Figure 9 summarizes the financial benefit assessment for a midsize financial institution.

Figure 9: Case study – Financial benefit assessment

Company	**Midsize Financial Institution**
Year of reference	**2024**

Financial details

Market cap	9.400B USD
Brand value	0.470B USD
Depository customers – households/businesses	500,000
HSA bank account holders	3,000,000
HSA bank business customers	30,000
Revenue	4.179B USD

Data center details

# of data center sites	2
IT staff FTEs	1,200
Application count (core + noncore)	800
Annual customers (retail + HSA acc)	3,530,000
Annual TCO (including downtime)	156M USD

Annual financial impact

	SR Linux	**SR Linux and EDA**
Reduction in penalty cost	~$US24M–$US50M	~$US43M–$US94M
Reduction in revenue loss	~$US18M–$US33M	~$US35M–$US66M
Mimizing loss of brand value	~$US34M–$US78M	~$US87M–$US207M

The assessment compares financial impacts for two scenarios—FMO 1 (SR Linux alone) and FMO 2 (SR Linux and EDA)—against the current legacy solution, which aligns with a 3.7 nines availability. Migrating to SR Linux and EDA boosts network availability to 5.1 nines, yielding substantial financial benefits.

As explained in the previous section, since financial impacts vary due to multiple factors the case study uses two scenarios (moderate and aggressive) to reflect the range of savings. The range of savings listed below reflects these scenarios.

- Penalty costs, which may include regulatory fines and operational recovery efforts, are projected to decrease $US24–$US50 million with SR Linux (FMO 1) and $US43–$US94 million with SR Linux and EDA (FMO 2), reflecting a potential $US94 million reduction in the aggressive scenario.
- Revenue loss, driven by service outages and customer churn, is expected to drop $US18–US$33 million with SR Linux and $US35–$US66 million with SR Linux and EDA. This translates to a 0.4%–1.6% positive impact on the company's $US4.179 billion annual revenue because of reduced downtime and churn.
- Minimizing brand value and reputation loss, which can escalate to billions of dollars in severe outages (e.g., past case studies cite billion-dollar write-offs). The savings related to brand value loss ranges from $US34–$US78 million with SR Linux to $US87–$US207 million with SR Linux and EDA, showcasing significant risk mitigation.

These estimates, derived from industry averages, highlight the nonlinear relationship between downtime reduction (e.g., 96% with SR Linux and EDA) and financial savings, influenced by regulatory thresholds, transaction sensitivity and ripple effects.

While the current legacy solution incurs high costs, the improved reliability with SR Linux and EDA offers a benchmark for a midsize company. Fully quantifying the benefits with precise dollar values requires customer-specific data on current performance, costs and revenue losses, tailored to the legacy setup and planned migration.

Nokia SR Linux and EDA deliver real-world savings for a midsize financial network

- ~$US43M–$US94M savings related to penalty cost*
- ~$US35M–$US66M savings related to revenue loss
- ~$US87M–$US207M savings related to brand value loss

* For moderate and aggressive scenarios

5 Summary

In today's data-driven, information-hungry world, reliable networks are the backbone of every digital transaction and service. The data center fabric reliability study shows that moving from legacy equipment (99.981736% availability) to Nokia's modern solution—SR Linux alone (99.993017% availability) or SR Linux and EDA (99.999235% availability)—can cut annual downtime from 96 minutes to just 4 minutes, delivering up to a 23-fold reduction.

This reliability boost translates into concrete financial gains—up to 60% lower penalty and operational costs, 53% less revenue loss and 44% reduced reputational damage—while empowering operations teams, reducing alert fatigue and accelerating business agility.

With Nokia SR Linux and EDA, operators can embrace the Human Error Zero approach now and turn high availability from an aspiration into a measurable, cost-saving reality.

6 Learn more

To learn more about how SR Linux and EDA can bring new levels of reliability, simplicity and adaptability to data center switching and cloud infrastructures, visit the Nokia Data Center Fabric page.

For an analysis of data center fabric operations functions, tasks and work efforts for all phases of the data center fabric operations lifecycle, refer to the Bell Labs Consulting white paper, "Data center fabric business case analysis (BCA)."

7 Abbreviations

AI	artificial intelligence
BGP	Border Gateway Protocol
CAPEX	capital expenditure
CCF	common cause failure
CI/CD	continuous integration/continuous deployment
CLI	command-line interface
DCI	data center interconnect
ECMP	equal-cost multi-path
EDA	Event-Driven Automation
EHS	event handling system
ERP	enterprise resource planning
EVPN	Ethernet VPN
FMO	future mode of operation
GNMI	gRPC Network Management Interface
IaC	infrastructure as code
IP	Internet Protocol
IXR	Interconnect Router
LCM	lifecycle management
MTTF	mean time to failure
MTTR	mean time to restore
NDK	NetOps Development Kit
NEM	network element manager
NOS	network operating system
OPEX	operating expense
PMO	present mode of operation
QoS	quality of service
SLA	service-level agreement
SNMP	Simple Network Management Protocol
VXLAN	Virtual Extensible LAN
ZTP	zero-touch provisioning

8 Appendix A: Nokia SR Linux and Nokia EDA differentiation

Table A1: Nokia SR Linux differentiation

SR Linux value adds	Feature description	How it improves solution reliability	Applicability
Unmodified Linux kernel	• Support for vendor-agnostic security patching and reuse of existing Linux tools	Continuous testing performed by global Linux community	D0, D1, D2+
Ground-up model-driven architecture	• No translation layer	Reduces complexity and transformation error	D0, D1, D2+
GNMI	• A single API for set/get and streaming telemetry	Reduces operational complexity and minimizes misconfigurations	D0, D1, D2+
Protobufs	• Open-to-the-core architecture enabled by protobuf-driven IPCs delivers 2x performance improvement with open and polyglot extensibility	Enables efficient, structured and language-agnostic communication between system components	D0, D1, D2+
Field-proven protocol stacks	• Minimum vendor dependence for CLI enhancements • Optimized, repeatable workflows • Tuning system to help with debugging • Open and familiar implementation (Python) • Orchestration of CLI plugins at a fabric level enabled by the Nokia EDA platform	Ensures robustness, stability and interoperability, which are core pillars of reliability	D0, D1, D2+
Programmable CLI plugins	• Minimum vendor dependence for developing network applications • NDK enables customers to develop extensions to NOS (SR Linux) and core components, including new protocols • Orchestration of NDK apps at a fabric level enabled by the Nokia EDA platform	Enables automation	D1, D2+
Enhanced telemetry	• Granular and programmable • High performance	Enables early detection of issues	D0, D1, D2+
Event Handling System	• Pre-programmed actions based on specified criteria	Ensures faster detection and repair	D2+

Table A2: Nokia EDA differentiation

SR Linux and EDA value adds	Feature description	How it improves solution reliability	Applicability
ZTP	• Abstract, intent-based ZTP • Automated secure key generation and distribution	Minimizes downtime risks and configuration mistakes through automation	D0, D1
Infrastructure as code – declarative intents	• Representing IaC enables repeatable and reliable automation • Infrastructure (fabric intents) as code makes it easy to compare differences between multiple versions of fabric intents • Wide-ranging development tools and expertise reused, making it easy to hire skillset on tools or generic DevOps	Repeatability	D1, D2+
Fabric design	• Customers can create various flavors of fabric designs and visualize them with detailed cable connections and per-node configurations	Reduces man-hours required for fabric design and improves operations	D0
Design validation using digital twin	• Customers can validate their designs with a digital twin of the real network, which provides considerable CAPEX and OPEX savings because 80%-90% of scenarios can be validated without real hardware	Reduces the risk of configuration errors and failures in the real environment	D0
Intent-based fabric automation	• Fits into the ZTP architecture and does not require detailed networking knowledge—zero human errors and fewer man-hours	Ensures consistent and error-free deployment	D1
Group-based upgrade service	• Enables within a small maintenance window upgrade of all switches in a fabric • Parallel and automated upgrades enable upgrades via smaller maintenance window, less downtime etc. mapping to minimum downtime to the business	Reduces downtime	D2+
Fabric operations	• Unique operational visibilities deliver accurate capacity planning, reducing CAPEX by avoiding over-provisioning • Unique operational insights enable faster root cause analysis, which reduces man-hours and downtime	Error protection	D2+
Event handling system	• Pre-programmed actions based on specified criteria	Faster detection and repair	D2+

About Bell Labs Consulting

Bell Labs Consulting empowers clients to harness future technologies for enduring business success. By integrating our expertise across business and services strategy, network and cloud infrastructure, and operations and transformation, we craft solutions rooted in empirical evidence, ready for immediate application.

Bell Labs Consulting engages with leading global clients to inspire and drive the industry's transformation through the accelerated adoption of future technologies. Bell Labs Consulting builds actionable, client-specific, leading edge solutions. We leverage our deep industry knowledge, our broad experience, and analytical tools to identify strategic opportunities, support decision-making, optimize deployments and operations, and support our clients to succeed in their business and digital transformation.

Bell Labs Consulting is part of Nokia Bell Labs, the world-renowned research arm of Nokia. Nokia Bell Labs invented many of the foundational technologies that underpin information and communications networks and the corresponding digital ecosystems. This research has produced ten Nobel Prizes, five Turing Awards and numerous other awards.

To contact Bell Labs Consulting on reliability or any other technology strategy challenge, visit https://www.nokia.com/bell-labs/bell-labs-consulting/.

About Futurum

The Futurum Group is a leading independent research, advisory, media and analyst organization focused on analyzing emerging and market-disruptive technologies. Through its portfolio of companies—including Futurum Research, Signal65, Six Five Media, Tech Field Day and Visible Impact—the firm provides data-driven insights, market intelligence and strategic content that help global enterprises navigate the intersection of technology innovation and business transformation.

Futurum and Solutional work together as trusted research and content partners, combining Futurum's analyst, media, and amplification capabilities with Solutional's deep subject-matter expertise and production excellence. Together, they deliver high-impact research, creative assets and multichannel amplification programs that bring complex technology narratives to life for technical and business audiences alike.

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, which is celebrating 100 years of innovation.

With truly open architectures that seamlessly integrate into any ecosystem, our high-performance networks create new opportunities for monetization and scale. 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

Document code: SR1549500EN (November) CID214998