Optical transport networks, automation and high-performance data center connectivity

Built for AI, powered by AI

NOKIA

Is your optical network ready for AI-driven workloads?

Network operators are adding more graphics processing units (GPUs) to their data centers to keep pace with the performance demands of artificial intelligence (AI) applications. These GPUs are often clustered in multiple locations that must then be linked. While optical transport remains the best technology for low latency connectivity, lower power/Gbit transport and high-density solutions, the need to support fast, efficient data transport across geographically distributed locations and resources adds new layers of complexity for operators.

Automation is essential for reducing the effort involved in designing, scaling, deploying and managing the massive numbers of interconnections required for AI. And automation powered by AI has the potential to augment the intelligence of the network to make interconnect operations more productive and improve the user experience for network operators. This paper explores how the massive data flows and distributed demands of AI impact optical networks. It also explains how AI can support innovative network automation that helps operators deliver efficient, secure, intelligent and scalable interconnect transport solutions that meet the demands of a world soon to be driven by AI.

Introduction

As interest in artificial intelligence (AI), machine learning (ML) and generative AI (GenAI) grows, many network operators are outfitting their data centers with compute resources—such as graphics processing units (GPUs)—that will accelerate the parallel computations required by classic AI methods (such as neural networks) and large language model (LLM)-based GenAI.

The computational tasks and performance requirements of a GPU cluster depend on node size, because larger nodes require more GPUs, memory and storage. However, the GPUs may not all be located together. If GPU cluster nodes span multiple racks, multiple floors of a building or geographically dispersed data centers, their locations need to be linked.

Optical transport provides an ideal method for interconnecting GPU clusters. Operators can support these interconnections using point-to-point connections, but it's more likely the data will need to traverse multiple locations across amplified or reconfigurable optical add–drop multiplexer (ROADM)-based line systems.

Multi-location interconnection increases complexity for network operators. They need to manage the interconnections to ensure they comply with the strict requirements of the GPU clusters while maintaining their performance to ensure they operate efficiently in a distributed, clustered manner. Automation can help reduce the complexity involved in designing new connections and then scaling, deploying and managing these networks. As operators introduce automation, they can take advantage of AI to augment the network's intelligence and reduce costs by making their operational practices more productive.

AI traffic as a driver for optical network growth

AI-powered applications require significant volumes of data to train classic AI or GenAI models and ensure they can be used productively across industries. Training and tuning these models to improve their accuracy further increases the demand for high-speed optical connectivity, especially as enterprises connect their private data centers to public data centers.

The rise of AI data flows and high-bandwidth workloads between data centers is a pivotal development. It requires high-performance optical transport technology with network automation and management layers that can optimize interconnections, ensure their resiliency, and minimize downtime and data loss. Zero loss is particularly important during AI model training and inferencing.

GPU nodes can be configured to improve connections between server clusters. As data centers adopt these optimized configurations, it increases demand for optical and fiber innovation. These demands extend into optical transport systems that span metro and regional networks, which require network management along with engineering resources to operate these systems throughout the network lifecycle. AI can play a vital role in helping to manage these networks using automation solutions that can scale optical fiber networks to meet the high processing demands of AI clusters.

AI's role in today's networks

AI is revolutionizing the optical network lifecycle by enabling networks to predict demand, detect and resolve faults, and optimize routes in real time with minimal human intervention. In other words, AI opens a path towards autonomous networks. Composed of an optical system with an intelligent software platform, an autonomous network can **sense** its environment, process that information to **think** about what it senses, and then **act** by choosing the best course of action to achieve specific business objectives.

With high-quality, network-trained data, AI will enhance network automation capabilities by providing predictive analytics, intelligent decision-making and personalized service offerings, creating networks that are not just automated but truly smart.

Automation initiatives are evolving beyond simple robotic-based actions that use scripts or templates to more sophisticated closed-loop operations based on abstractions of requirements called intents. For example, intents can identify service-level agreements (SLAs), route-specific policies or geographical service endpoints to allow the underlying system to orchestrate the operations required to achieve business outcomes. Operators can take advantage of intents without requiring expert knowledge of the optical technology.

Operators that adopt intent-based networking can simplify their network operations by eliminating the need to execute complex manual processes to configure network equipment or react to network issues. They can extend this concept by using AI to enable intents through natural language processing or by using AI within the underlying network to fulfill intents while continuously fine-tuning the system to sustain their business outcomes.

AI will play a key role in moving autonomous networks beyond reactive closed-loop operations. Several standards forums (e.g., Mplify, TM Forum) are defining frameworks that describe the functions in which AI technologies can be integrated to enhance network resiliency. These functions include predictive analysis, intelligent decision-making, self-healing, resource optimization, solution reporting and issue resolution. The use of AI to migrate to autonomous networks will take advantage of closed-loop systems already established within network automation. It will, however, continuously feed the knowledge it acquires back into these systems to ensure the goals associated with the intents are achieved.

With high-quality, network-trained data, AI will enhance network automation capabilities by providing predictive analytics, intelligent decision-making and personalized service offerings, creating networks that are not just automated but truly smart. AI will discover new patterns from the network through unsupervised learning and address many types of operational use cases that are not currently possible without a lot of expert human involvement.

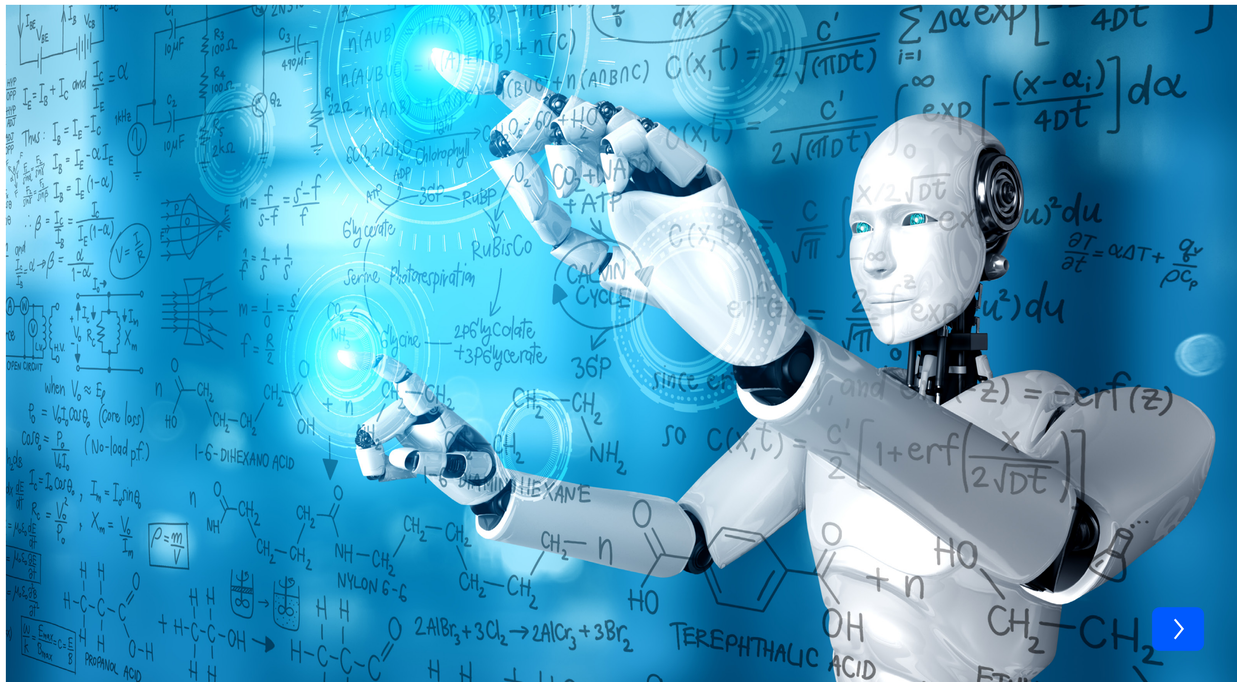

How to build optical transport networks for AI

Understanding the distances between the data centers that include the GPU servers and the length of the optical fiber that interconnects them is essential for planning and designing an optical transport network. Because LLMs can be split and run on multiple GPUs across different node locations, the high capacity of optical transport makes it an ideal choice for interconnection. However, LLM splitting imposes transport planning constraints and requires optical networking that:

- **Ensures high-performance transmission:** AI clusters demand optimal optical performance, including high spectral efficiency, low latency and low energy consumption across various fiber types and link distances.
- **Protects in-flight data:** Security is top of mind as the number of AI data centers grows and interconnects increase. Incorporating at-speed encryption within the optical layer as part of a resilient, multilayer defense-in-depth strategy against intrusive threats will be critical for protecting sensitive data used for AI.
- **Delivers resilience and assurance:** Sourcing real-time telemetry key performance indicators (KPIs) from the network lays the groundwork for building autonomous operations that enable optical networks to sense, think and act in response to network events. For example, autonomous networks can sense optical link degradations or failures, understand how they affect network performance and act to restore traffic on alternative routes.

The speed at which optical layer 1/0 service connections are created determines how quickly data centers can implement GPU resources for AI. This, in turn, affects business outcomes for communications service providers (CSPs), hyperscale and cloud operators, and enterprises alike. The distribution of data centers across metro, regional and long-haul distances increases the complexity involved in managing networks designed for AI scale, performance and SLAs. Simplifying optical multilayer management and performance monitoring across transport segments connectivity data centers becomes vital for maintaining the efficiency of GPU clusters.

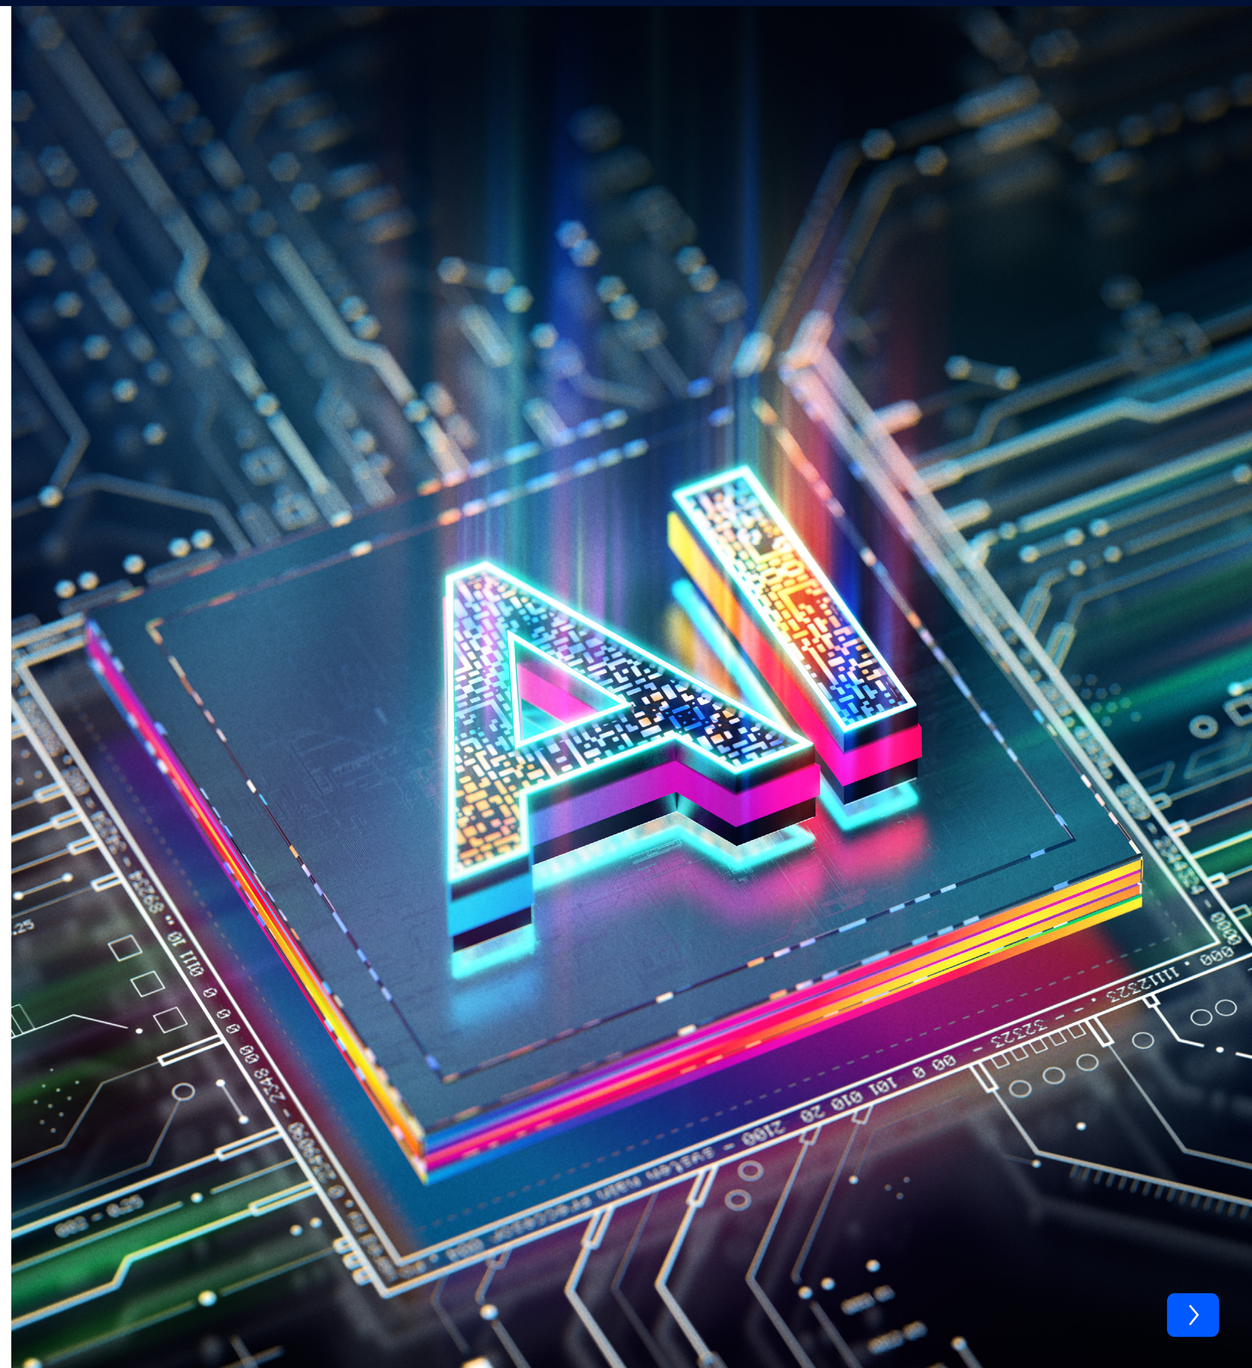

AI-assisted optical operations

The optical network lifecycle is made up of several phases, including planning, commissioning, provisioning, assurance, analysis and optimization.

Within each of these phases, there are many types of operations that a network operator must complete, many of which are time-consuming and error prone. These are areas where the use of AI technology within automation can help simplify operations, including:

- **Optical network design**, which involves creating sites, assigning interconnect rates, span lengths and fiber types, and running feasibility analysis. AI-assisted operations can use a natural language interface to facilitate the design process and improve the user experience. AI can also help network operators optimize optical network performance during the design phase. For example, it can optimize coherent interconnect transmission profiles to achieve the highest capacity without compromising system margins for training or inference cycles.
- **Generating the equipment bill of materials** and preparing the method of procedure and blueprints required to install and interconnect the network equipment at each site. GenAI-assisted operations can make it easier to produce and assemble information and reduce the time required to complete these tasks.
- **Configuring and provisioning the network** for the optical line system and services between endpoint locations. This requires insights on documentation and interfaces, and often results in operational delays. AI-assisted operations can automatically interact with hardware elements and software systems to set up the network.
- **Tracking performance metrics** within the network on a per-link basis to capture a status of capacity consumption for planners, along with service and node availability for network assurance. AI assisted operations can pull data from various network elements and assemble reports on specific metrics through a single pane of glass which can eliminate “swivel-chair” management.
- **Developing troubleshooting tactics** for resolving issues as they are uncovered while documenting them and training other network operators so they can handle similar issues in the future and gain a better understanding of their own networks. AI-assisted operations can diagnose issues as they arise and offer key recommendations to speed up root-cause analysis and notify key stakeholders. AI systems can also retain this knowledge for future occurrences.
- **Establishing maintenance procedures**, conducting repairs and retuning the network configuration are operations that help optimize performance of the network as it ages or when problems arise. AI-assisted operations can adjust the infrastructure to keep the network at optimal conditions, reduce time to repair and extend network asset lifetimes.
- **Augmenting the expertise of operators to accelerate end-customer service delivery**. GenAI can help operators streamline processes, make better decisions and simplify onerous tasks. This type of AI requires LLMs that are trained with measured datasets from the network and large volumes of documentation on use cases, faults and actions that cover various types of transport configurations.
- **Analyzing vast amounts of real-time network KPI data** with predictive AI to mitigate potential disruptions before they impact services, forecast the optical capacity required to support new GPU cluster interconnects or reduce electrical power consumption by placing some interconnects in standby mode when traffic demands are reduced.

Operators can use AI in transport networks to evolve their business, simplify their operations, adhere to strict latency- and performance-aware SLAs for the high-throughput interconnections, and flexibly deploy GPU clusters in multiple locations.

Securing optical networks for AI

The rise of data-driven AI applications—particularly for sensitive applications such as medicine or finance—demands robust security solutions that will protect in-flight data from physical and quantum-based threats. As enterprises shift to cloud platforms and data centers interconnected for AI workloads, there will be more potential vulnerabilities in optical data transmission. Threats to encryption, such as cryptographically relevant quantum computers (CRQC), and to physical infrastructure, such as eavesdropping, could jeopardize the integrity of the network layers, from applications to the optical transport infrastructure.

An approach that combines optical transponders that adhere to robust encryption standards (such as 256-bit Advanced Encryption Standard) with strong encryption keys (with high entropy, such as classical physics- or quantum-based keys) distributed symmetrically will keep in-flight data safe from potential threat actors as it passes between GPU nodes.

This comprehensive approach enables network operators to build a resilient defense-in-depth framework that includes the optical network domain layer. It also begins their migration to post-quantum cryptography (PQC), which will enable them to safeguard critical infrastructure and digital transformations against quantum threats today and in the future.

Overcoming adoption challenges

While AI brings significant promise to optical network management, it also introduces challenges, such as:

- **Autonomy and control:** In developing fully autonomous optical networks, operators must comply with ethical and regulatory frameworks to avoid unintended impacts on sensitive data and SLAs.
- **Explainability:** To build trust in AI systems, operators must be able to demonstrate that these systems have clear, explainable decision-making processes and minimize biases present in training datasets.
- **Data access:** Ensuring access to network KPI data is crucial for training reliable AI models. Post-training, these models must be tested and proven not just effective, but beneficial in real-world deployments.
- **Data integrity and security:** Protecting sensitive data from being harvested and used for malicious purposes is essential, particularly in training AI models.
- **AI governance policies:** Operators must ensure that adoption of AI within optical systems abides by worldwide and country-specific regulatory frameworks for data governance. Providing the necessary insights from AI systems used in automation will allow the algorithms to be monitored, evaluated and updated to prevent flawed or biased decisions as models drift, which can lead to issues with output quality and reliability. Triggering maintenance conditions that require data sets to be retrained may be necessary to ensure AI delivers technological innovation with safety.
- **Power-optimized LLMs:** Operators must carefully select the LLMs they use for network automation to reduce the compute resources required for model storage and inference processing. By optimally using low-scale models to cover a larger application space that may not necessarily be used by automation use cases, they can reduce consumption of electrical power. Operators may require efficient use of one or more smaller-scale LLMs to ensure network automation doesn't hinder their pursuit of long-term environmental, social, and governance (ESG) objectives.

Building networks for AI, using AI, with Nokia

The Nokia edge-to-core portfolio of optical transport solutions, powered by industry-leading coherent optical engines, provides the foundation for scalable, resilient, and quantum-safe inter- and intra-data center network infrastructure. Managing and optimizing these networks at scale in the era of AI requires intelligent automation.

Network operators use Nokia WaveSuite to simplify operations across the entire optical transport lifecycle, taking advantage of AI-driven insights to enhance network planning, streamline operations, and unlock new monetization opportunities. Offering built-in explainable AI, predictive analytics and real-time automation, WaveSuite helps operators accelerate service delivery, optimize network performance and ensure long-term scalability in an era of increasing interconnect complexity.

Industry-leading hardware

Network operators can rely on the Nokia optical transport network portfolio to deliver fast massive interconnect capacity with exceptional performance, security, and spectral efficiency – from the network edge to transcontinental subsea connections.

Powered by in-house, state-of-the-art optical coherent digital signal processors, photonic semiconductor technologies, coherent optical line systems and intelligent pluggables, the portfolio enables:

- Scalable, reliable, and secure data center connectivity
- Lower cost and power per bit of data transport
- Automated and optimized network performance to streamline operations and lower TCO
- Enhanced data control and security to meet sovereignty and mission-critical requirements
- DC optimized compact platforms with the right open APIs

The Nokia optical transport portfolio – which includes the 1830 Global Express (GX) compact modular transport platform, 1830 Photonic

Nokia optical networking enabling solutions – innovation and performance to network the cloud from edge to core.

Automation		**Automation** Enable network scale and monetization with faster, simpler operations while lowering TCO.
Layer 2 Ethernet		**P-OTN switching** Industry-leading network scale supporting efficient, resilient and secure wholesale and business services and Integrated Packet Transport.
Layer 1 OTN and transport		**Coherent transport** Pluggable and super-coherent embedded optics; X-ponder line cards in optimized edge-to-core-to-DC platforms.
	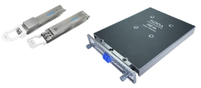	**Optical engines** Pluggable: open, standardized and high-performance QSFP/OSFP; Super-coherent: embedded optics driving the Shannon Limit.
Layer 0 WDM line system	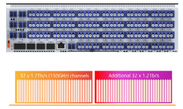	**WDM line systems** Market leading C+L open line systems: C, L, integrated C+L; Complete ROADM options from FOADM to hyper-degree multi-fiber.

Service Switch (PSS), ICE-X multi-haul pluggable coherent transceivers, ICE-D intra-data center optics, and network automation and management solutions – helps network operators drive down cost and power per bit, simplify network operations, and scale their network deployments across geographies – in a sustainable way with built-in intelligence that is quantum-safe.

World-class automation software

Nokia, a global leader in optical network automation technology, offers the WaveSuite automation platform to address key operational challenges across these phases within the transport domain. Built around research by Nokia Bell Labs, the WaveSuite platform integrates AI capabilities designed to improve productivity throughout the lifecycle while making AI explainable to the user to help build trust in the technology.

WaveSuite optical automation addresses the scaling challenges identified above by enabling:

- Efficient and simplified network operations that shorten the time to provision, configure, deploy, and manage optical networks.
- Automated network planning to optimize network resources, ensuring optimal performance between metro, regional or transcontinental data centers.
- Simplified and accelerated service turn-up for new business models—such as GPU as-a-service network slices—through service virtualization and network abstraction to ensure better compliance with GPU interconnection SLAs.
- Trending analytics to proactively plan and schedule maintenance activities and resolve potential faults within the optical network before they impact the interconnection network.

Network operators rely on the WaveSuite network automation platform to improve time to market and reduce capital expenditure. WaveSuite helps ensure their networks deliver the needed bandwidth securely and with performance optimized for interconnect delivery.

Beyond scaling networks to meet new data center interconnect challenges, WaveSuite's capabilities are categorized into four key areas:

- **Manage:** Simplify network planning and scaling to enable secure user access and efficient resource management.
- **Operate:** Streamline capacity upgrades and expand network resources through automated and efficient business processes.
- **Optimize:** Utilize insights from network health KPIs to quickly adapt to changing demands, manage real-time events and ensure maximum service availability.
- **Monetize:** Develop innovative and dynamic service offerings, such as network-as-a-service, to enhance customer satisfaction, adhere to SLAs and accelerate time to market.

Network operators can take advantage of AI to transform these categories into intelligent, interactive frameworks. This will help them improve operational efficiency and achieve their service delivery goals through AI technologies such as classic algorithms, and generative AI.

While network operators are focusing on developing, implementing and servicing optical transport networks that can meet the capacity, resiliency, and performance requirements for AI data center interconnects, they can also take advantage of AI technology to support their network operation tasks across the entire network lifecycle. Within WaveSuite, AI will promote the following transport domain operational use cases in each of its framework pillars:

- **Manage:** Use capacity forecasting to support hands-free design, configuration and phased network rollouts. This will enable operators to streamline decision-making and respond to changing conditions in real time.
- **Operate:** Enhance knowledge sharing and troubleshooting through NLP-powered guides, live network queries and role-based network management. Trained with embedded documentation, GenAI will enhance the overall user experience and accelerate root-cause resolution, reducing costs and improving network quality. Further, explainable AI mechanisms will foster trust of the AI-driven recommendations.
- **Optimize:** Employ pattern recognition and fiber sensing to identify and assess network behaviors, detect physical infrastructure threats and mitigate service disruptions before they occur. This will improve SLA adherence and minimize the need for human intervention. AI will complement existing system-level closed-loop operational tasks and reduce energy consumption by identifying standby network equipment and optimal low-energy routes, developing a more sustainable network.
- **Monetize:** Help service operators develop end-subscriber-based, end-to-end optical domain slicing with service-level characteristics tailored for their specific applications. This includes specific connectivity requirements for training and inferencing the LLMs within the GPU cluster.

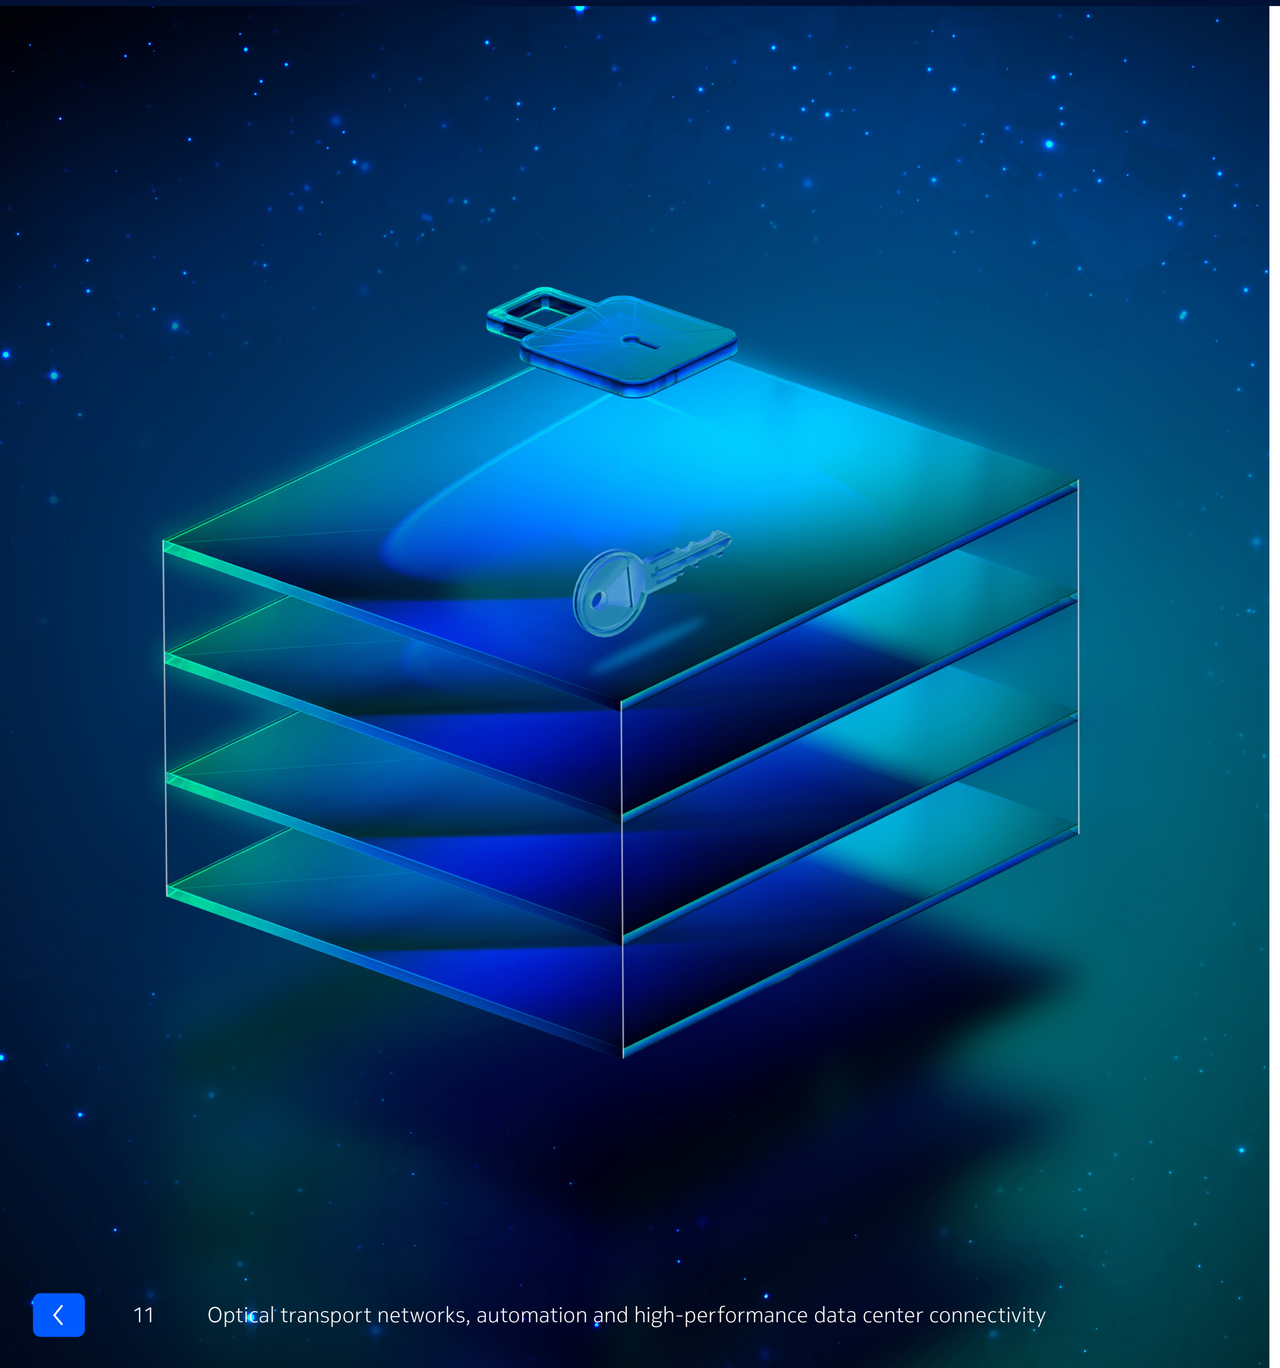

Network operators that offer interconnects to clustered GPU nodes can use AI capabilities within the transport domain to upsell performance by automatically reconfiguring service characteristics (e.g., latency, diversity and protection) based on optimized performance. GenAI can generate service intent requests by using natural language expression to expand order fulfillment within the transport domain.

Securing everything with quantum-safe networks

Nokia Quantum-Safe Networks (QSN) mitigate these risks by incorporating advanced symmetric cryptography and quantum-resistant high-entropy key sources into the optical transport layer. These features, powered by Nokia Bell Labs research and Nokia quantum partners, safeguard data against classical and quantum attacks. They also ensure long-term security against future decryption attempts (harvest now, decrypt later scenarios).

Overcoming obstacles to AI in network operations

The Nokia WaveSuite platform addresses concerns with AI adoption—including autonomy and control, explainability, data access, data integrity and security, governance, and power consumption — through advanced methodologies, secure policies and Nokia Bell Labs research.

Features such as NLP-based user prompts, decision-tree insights gleaned from product documentation and role-based controls ensure operators maintain oversight while benefiting from AI-powered automation to improve performance and reduce downtime.

By integrating security best practices and adopting tools such as Nokia Digital Twins and Optical Professional Services to develop business process modelling, WaveSuite supports the transition to trusted and adaptive partial or fully autonomous networks. This enables service providers to build networks that think, sense and act so they can focus on driving business growth rather than managing the network.

Future trends and opportunities

The integration of AI in optical networks will continue to drive advancements and create opportunities across the industry, including areas such as:

- **Standards development:** Organizations such as the IETF, Mplify and TM Forum are drafting proposals to standardize traditional and GenAI integration. Nokia actively contributes to these discussions, ensuring future solutions align with industry needs.
- **Monetization:** GenAI-powered intent-driven optical connectivity driven by operations support system (OSS) integration can enable premium services, such as bandwidth-on-demand, which boost revenues for business-critical applications that require enhanced performance.
- **Predictive intelligence:** AI can facilitate advanced fault detection and event classification, reducing downtime and improving SLA adherence.
- **Scaling challenges:** As private/public data center interconnects grow, AI network automation can address operational complexities, particularly amid shortages in the skilled workforce.
- **Pluggables and IP/optical convergence:** As pluggable digital coherent optics (DCOs) deliver higher transmission rates and longer reaches, their use for metro and regional interconnects will grow. This convergence between the IP and optical domains for interconnection will require AI-based network management to provide better coordination and data exchange. AI will also facilitate intent-based human interaction using natural language.
- **Energy optimization:** AI data centers place large demands on electrical infrastructure. Service providers can use AI to identify and adapt to energy-intensive network hotspots, applying deep learning algorithms and policy-based optimization strategies to minimize impact on the already constrained electrical grid.

Conclusion

AI is a disruptive technology that is changing the way industries work. But it is also a technology that will continue to fuel the rapid growth of the transport network that supports it. With the massive number of optical interconnects required for AI, data center connectivityand other high-bandwidth services, the transport network can reach unprecedented scales. Network operators need the same technology used within automation to transform the way they interact with and manage the network lifecycle. This technology must also enable them to stay focused on their business objectives over time.

The Nokia optical networking solutions portfolio, WaveSuite automation platform, and Optical Professional Services deliver efficient, secure, intelligent, simplified and scalable solutions tailored to AI-driven demands and high-performance data center connectivity. By adopting these technologies, service providers can unlock new opportunities, enhance productivity, and expand their addressable market reach, all while building trust and ensuring sustainable growth in an increasingly digitally connected world.

Abbreviations

AI	artificial intelligence
CRQC	cryptographically relevant quantum computers
DCO	digital coherent optic
ESG	environmental, social and governance
GenAI	generative artificial intelligence
GPU	graphics processing unit
IETF	Internet Engineering Task Force
KPI	key performance indicator
LLM	large language model
Mplify	Mplify Alliance
ML	machine learning
NLP	natural language processing
PQC	post-quantum cryptography
PSE	Photonic Service Engine
PSS	Photonic Service Switch
QSN	quantum-safe networks
ROADM	reconfigurable optical add-drop multiplexer
SLA	service-level agreement
TM Forum	Telemanagement Forum

Nokia OYJ
Karakaari 7
02610 Espoo
Finland

Tel. +358 (0) 10 44 88 000

CID: 214737 (September)

nokia.com

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.