NOKIA | Lenovo

Nokia and Lenovo partner to seamlessly automate AI workloads

Modern data center networks that meet high-performance demands

Modern data center networks that meet high-performance demands

Your data center networks are already stressed to the limit to meet the scalability, capacity and performance required to deliver the mission and business-critical applications and services that are essential to your company's success. So too are your IT and Network Operations teams as they struggle to manage and automate the existing network with legacy Operating Systems (OS) that do not provide the observability, programmability, and extensibility to dynamically adapt to a changing environment.

With the adoption of new technologies like AI/ML, new clusters and workloads will require a new generation of high-performance servers with higher density CPU (Central Processing Unit) and GPU (Graphics Processing Unit) processing together with an increase in data center networking bandwidth capacity and improved performance. These workloads can also be distributed across cloud networking infrastructures which require network operations management visibility and automation capabilities that cut across network boundaries.

Nokia and Lenovo have developed a data center networking solution that is built to meet the demanding processing and network performance requirements for your AI workloads. It combines the high-performance Lenovo ThinkSystem line of servers and storage systems together with the Nokia data center network fabric solution (fig 1).

Nokia and Lenovo data center network solutions

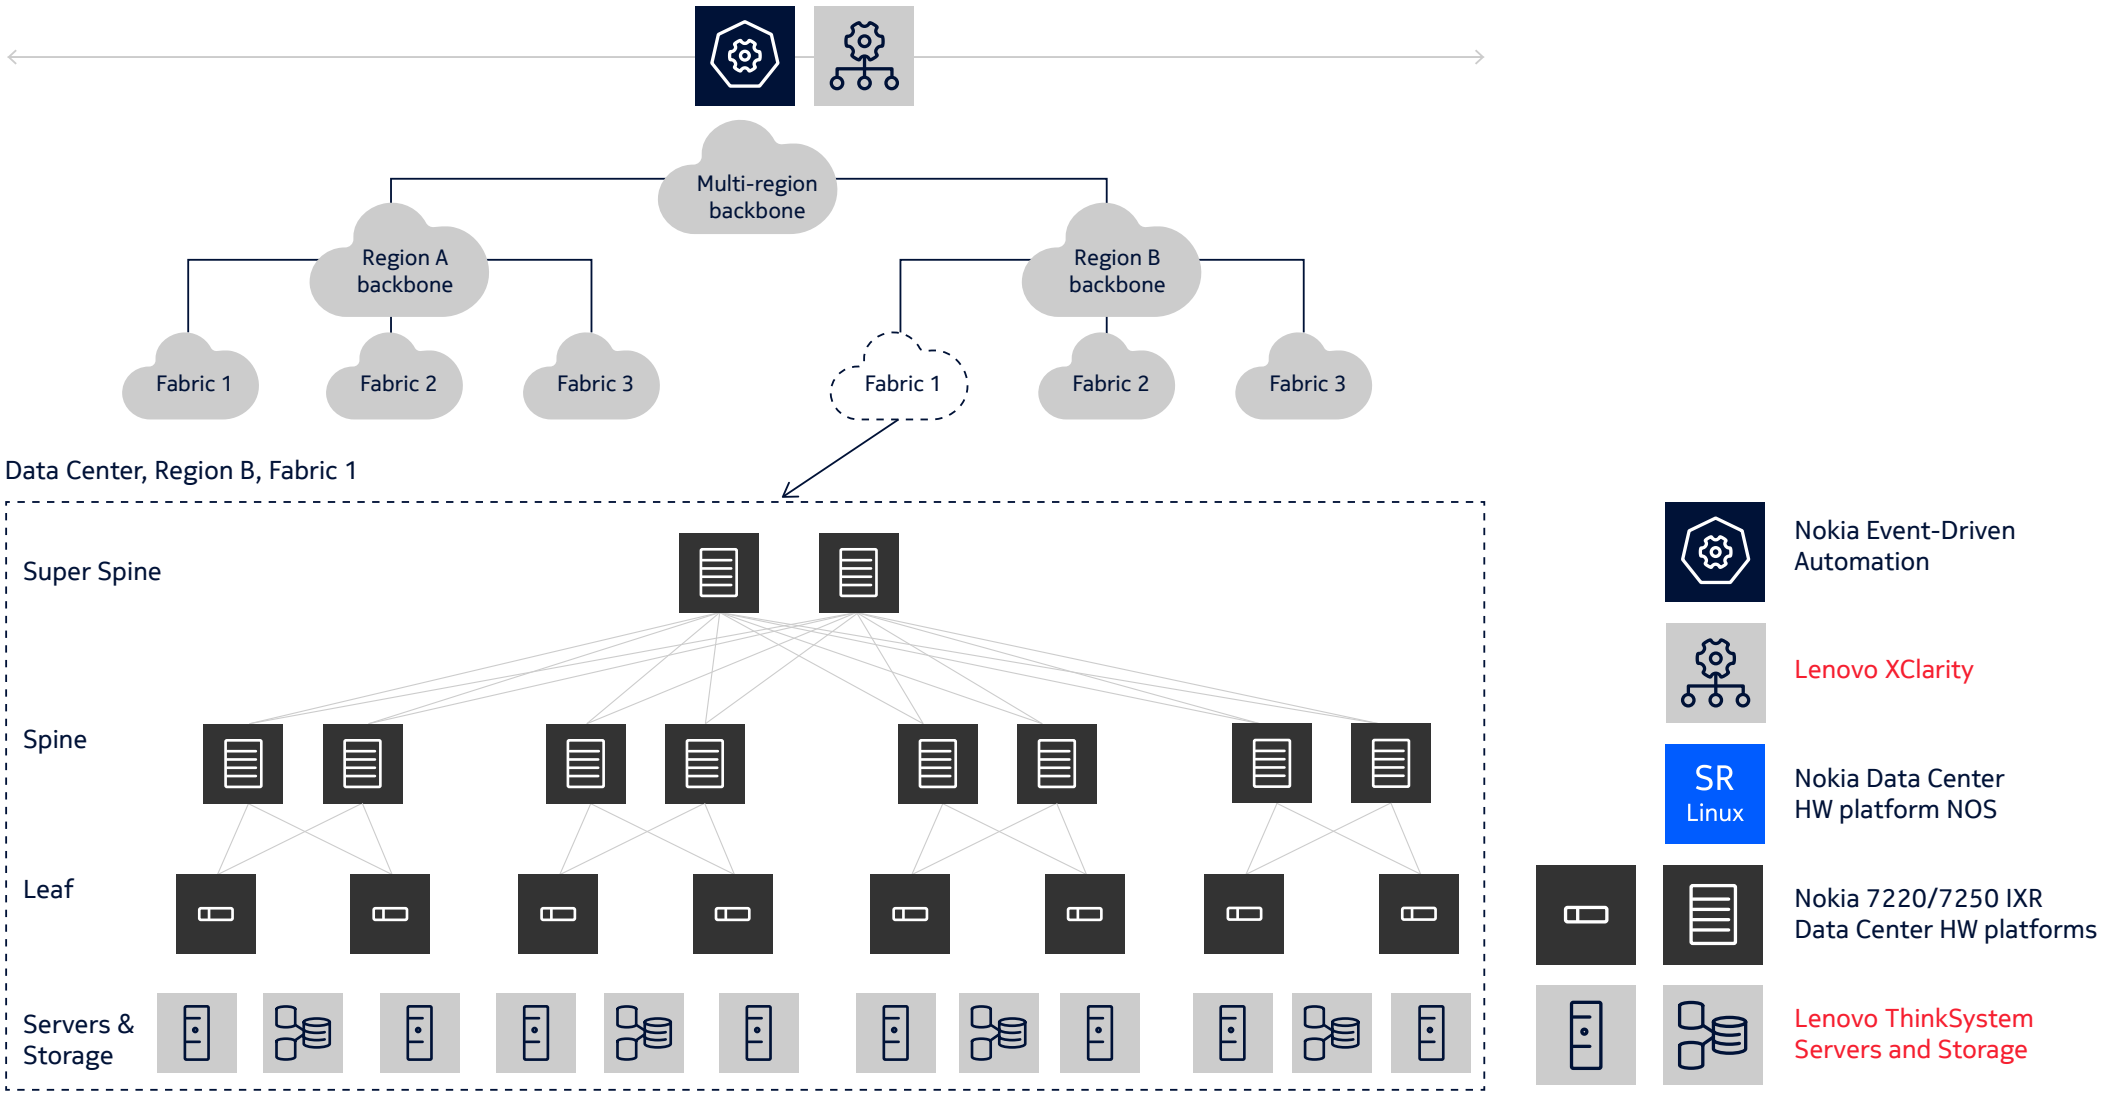

Figure 1: Nokia Data Center Fabric and Lenovo ThinkSystem solution

Nokia Data Center Fabric solution

The Nokia Data Center Fabric solution enables you to scale your data center networks quickly and easily to support high performance workloads, delivering the reliability and performance you need to deliver new applications while maintaining business continuity you require. It is a modern, high-performance solution that brings new levels of openness and automation to data center networking to help your operations teams manage accelerating demand with all the freedom and control they need. Built with merchant silicon, the solution consists of high-capacity hardware that support massively scalable and reliable data center switching architectures; a modern, open, and extensible Network Operating System (NOS), and an intent-based automation platform for all phases of the data center fabric operations life cycle (fig 2).

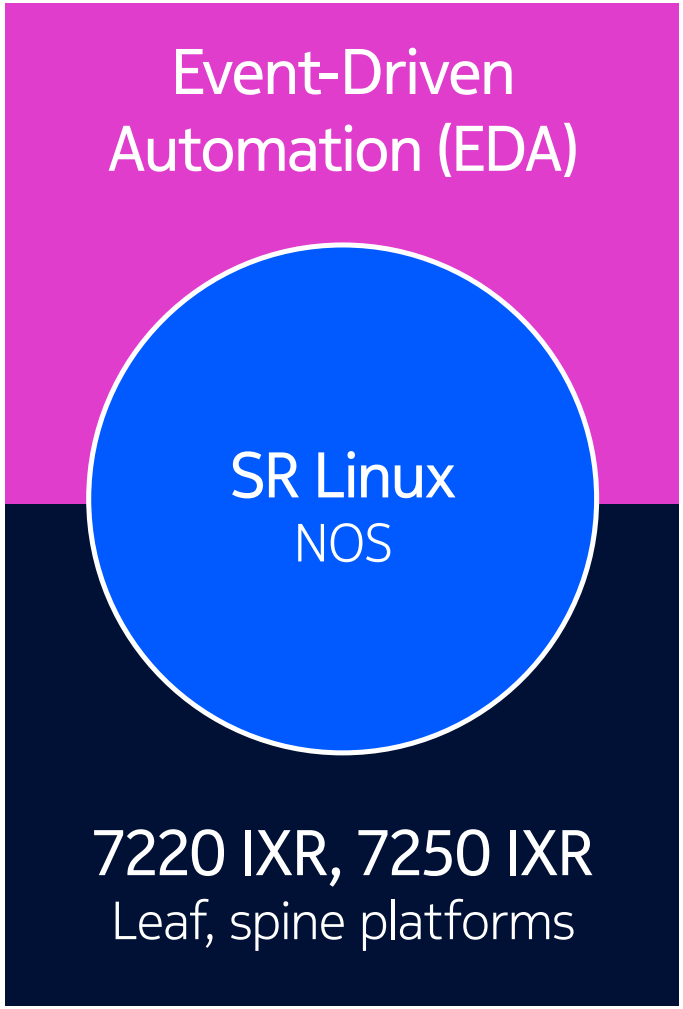

Event-Driven Automation (EDA) platform

Reliable, simplified and adaptable data center operations

SR Linux

Open, extensible and resilient NOS with proven routing capabilities

7220 IXR /7250 IXR

High performance, energy efficient leaf, spine platforms

Figure 2: Nokia Data Center Fabric Solution

Data Center hardware platforms

The Nokia 7220 Interconnect Router (IXR) provides fixed-configuration, high-performance data center platforms that bring unmatched scale, flexibility and operational simplicity to your data center networks and cloud environments. There are multiple 7220 chassis options that scale up to 51.2 Tb/s throughput with support for 400GE, 100GE, 50GE, 40GE, 25GE, 10GE or 1GE port speeds and redundant power and fans that are designed for leaf, and spine deployments.

The Nokia 7250 Interconnect Router (IXR) are modular and fixed configuration, high-capacity platforms supporting spine and super-spine functions. There are multiple chassis options with throughput capacity of up to 460.8 Tb/s and supporting high-density 800GE, 400GE, 100GE, 40GE, 25GE and 10GE interfaces for intra-fabric and server connectivity. These platforms are designed with superior innovations providing optimized power and cooling that provide an upgrade path to future-generation silicon and line cards.

The Nokia 7215 IXS interconnect system provides connectivity for leaf and spine data center fabric management connectivity in data center and cloud environments.

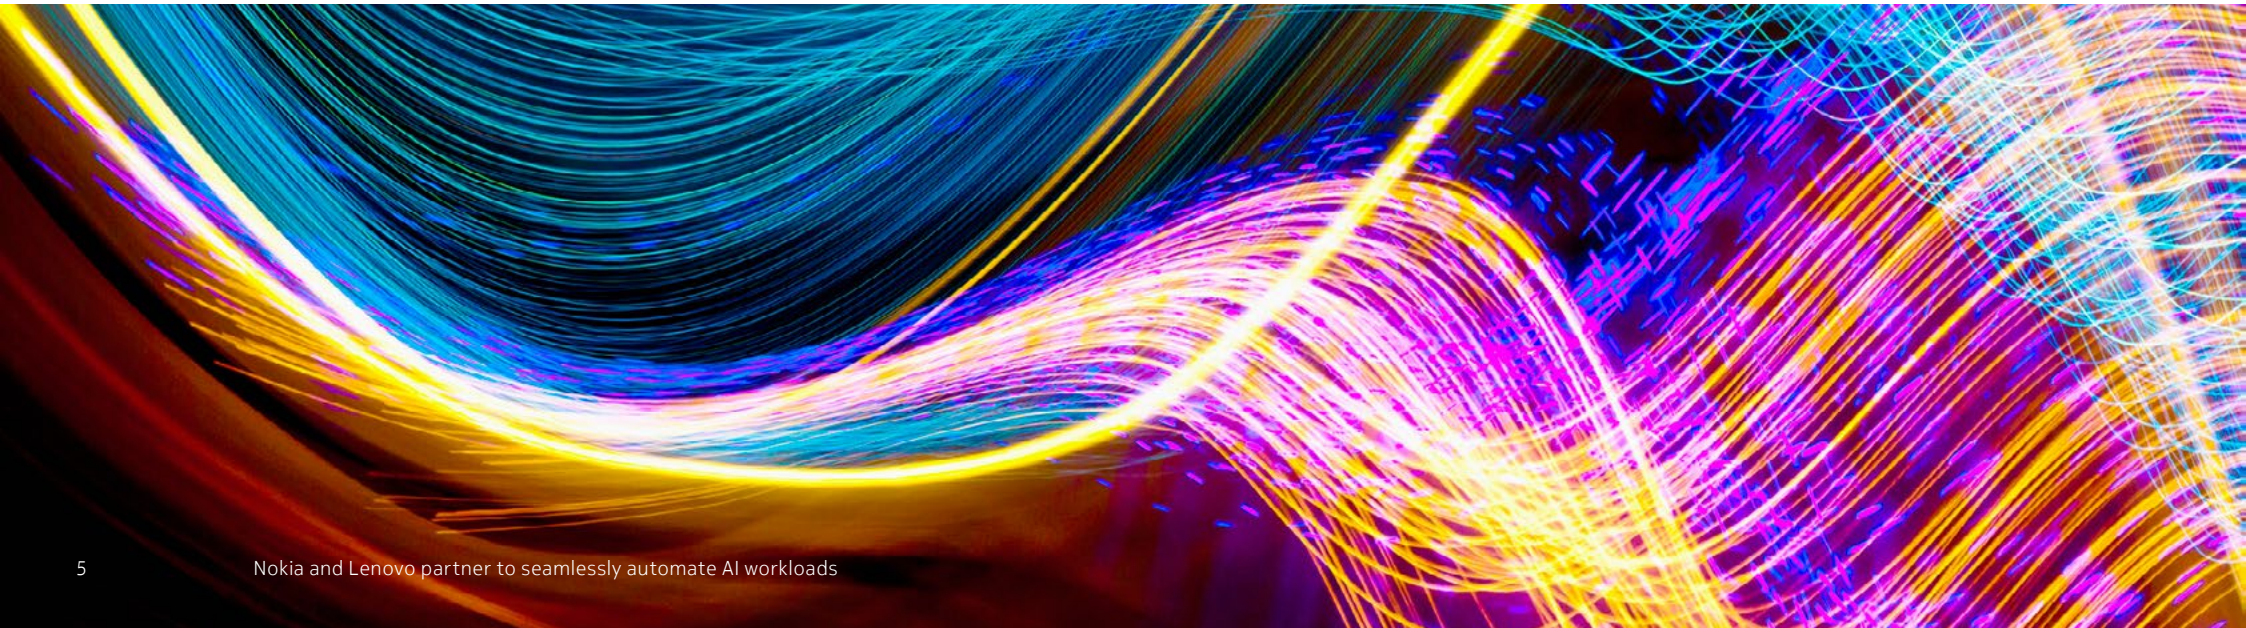

Open, extensible, reliable NOS

Nokia SR Linux is a next-generation NOS that is more scalable, more flexible, and simpler to operate while also being highly reliable that has been built with field-proven routing protocols. Designed for modern IP and data center networks, SR Linux uses an unmodified Linux® kernel to ensure reliability, portability, and straightforward application development.

SR Linux is built with a microservices-based model-driven design that leverages independent YANG data models to provide granular and broad access to network data. Modern standard interfaces such as gNMI and gRPC and others provide efficient and consistent ways to communicate with the network.

It supports an open and scalable telemetry framework that leverages on-change streaming telemetry and gNMI to provide new levels of network visibility.

SR Linux uses tested, hardened and field-proven routing capabilities (e.g., EVPN, MPLS, BGP, VXLAN) inherited from Nokia's SR OS that run the biggest IP routing networks in the world that provide proven performance at scale.

And it includes a NetOps Development Kit (NDK) and open-source CLI plugins that allow data center and cloud builders to create their own customized operational tools.

NetOps Automation

Nokia Event-Driven Automation (EDA) is an intent-based network automation platform that lets your data center operations keep pace with new performance and scalability demands while reducing the operational complexity with simple lifecycle management that adapts to your environment.

Designed to meet day 0/1/2 network operational phases with built-in redundancy, checklists, rollback and network-wide transactions features, EDA gives your operations team the confidence in the reliability of the system to move forward in automation.

It provides a single dashboard for all data center operations enabling true operational lifecycle management of your data center fabric with “on-change” telemetry for instant observability and an AI assistant to help with training and troubleshooting.

EDA adopts a cloud-native microservice design that is built on top of an open Kubernetes foundation along with easy integration with an eco-system of many open-source projects. This approach embraces the latest innovations in cloud computing to run a scalable, distributed and flexible platform in modern, dynamic IT environments.

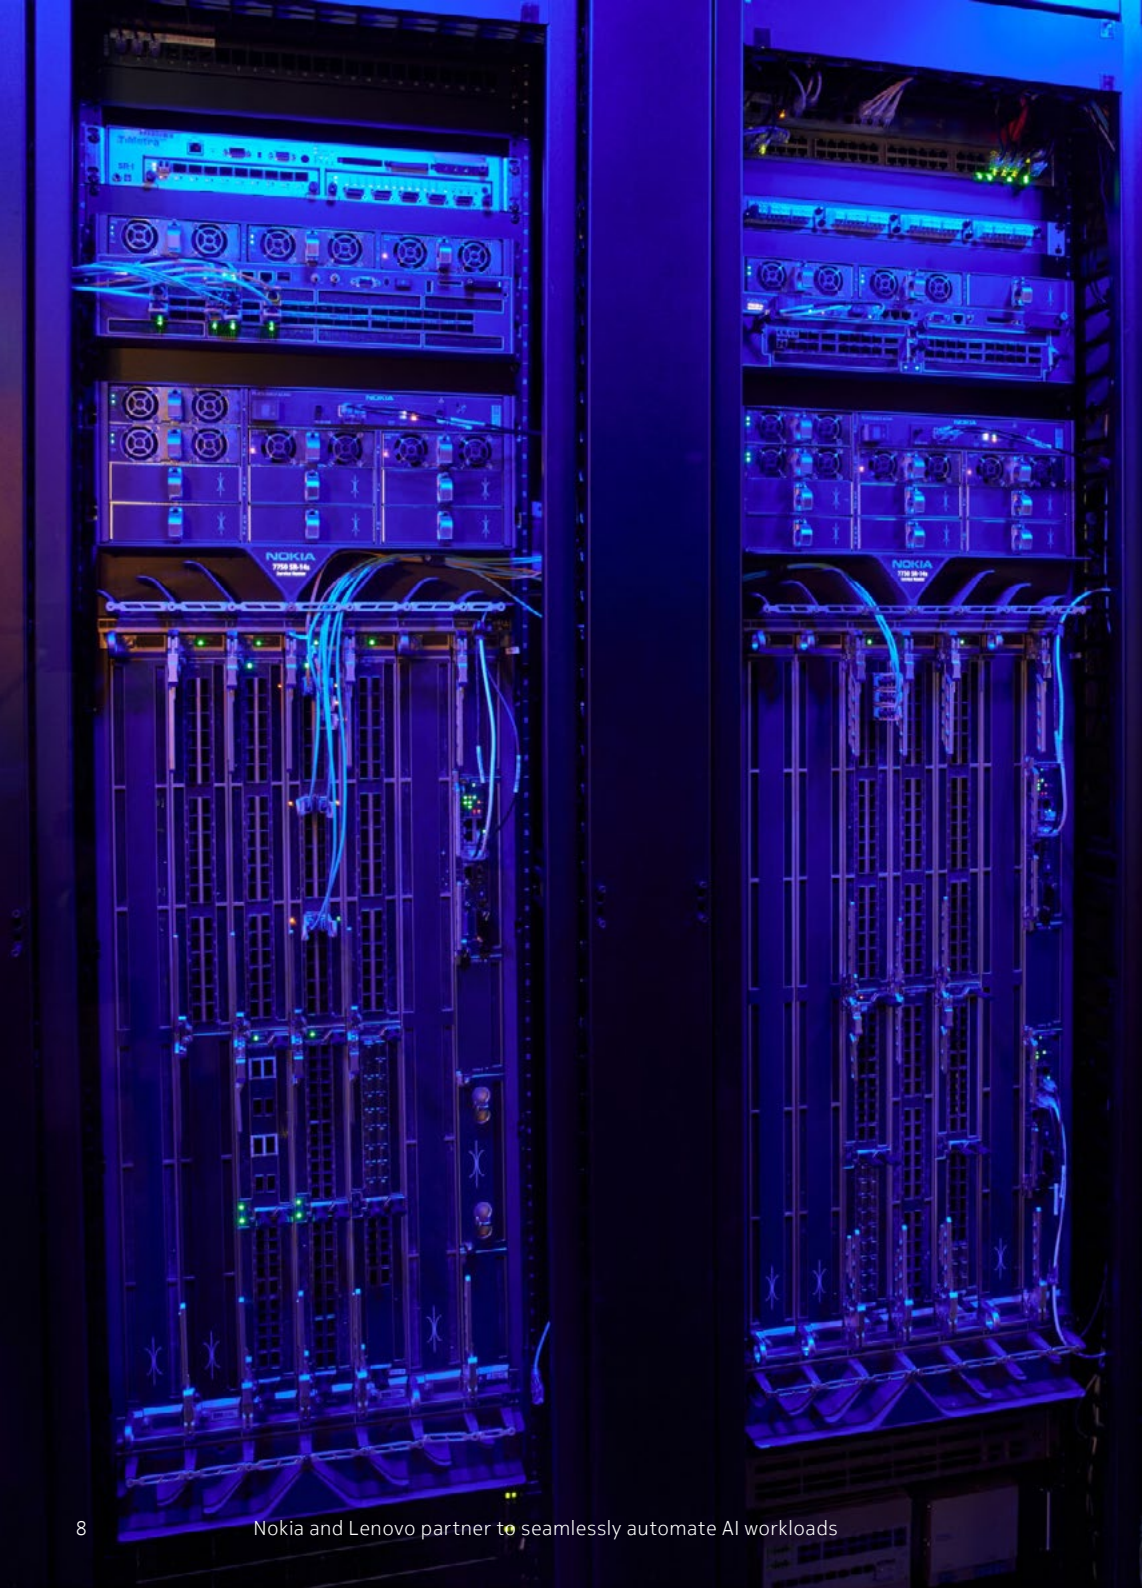

The Power of Five 9's Uptime – Lenovo Ranks #1 in Reliability

According to ITIC's Reliability Study, Lenovo has ranked #1 in server reliability for 10 consecutive years, establishing Lenovo as a best-in-class partner for modern AI and telecom networking infrastructure.

This amounts to an annual per server, per minute potential cost of $52.60 assuming hourly downtime losses of $100,000. In the ITIC 2022 Reliability study, Lenovo ThinkSystem servers averaged 1.10 minutes per server, per minute downtime with a potential associated cost of $1,837 per server, per minute based on an estimated Hourly Downtime cost of $100,000.

The result proved that Lenovo ThinkSystem meet these carrier-grade high-availability reliability requirements for telecommunications network infrastructure.

Read the ITIC Report

AI and HPC Lenovo ThinkSystem Servers Provide Trusted Infrastructure

The New ThinkSystem 8-GPU Servers

Providing massive computational power to the data center, ideal for AI and HPC compute-intensive workloads.

The new Lenovo ThinkSystem SR680a V3, SR685a V3, and SR780a V3 GPU systems deliver massive computational performance for Artificial Intelligence (AI), High-Performance Computing (HPC), and graphical and simulation workloads across various industries. The family of servers supports eight high-performance GPUs, either from AMD or NVIDIA, with planned support for GPUs from Intel.

To meet today's HPC requirements your servers should have a modern design. With Lenovo supercomputing servers, even the most technical simulations can be performed with gusto. These dense systems pack a punch for all your artificial intelligence, cloud, grid, inferencing and analytics needs. Designed to be scalable, unique water-cooled systems maximize density while putting security first.

The SR680a V3 and SR685a V3 servers feature two Intel® Xeon® and two AMD EPYC™ processors, respectively, and are 100% air-cooled servers suitable for most data centers. The SR780a V3 server features two Intel Xeon processors and uses a hybrid water-air cooling system to take advantage of the efficiencies of cooling the CPUs and GPUs using a direct water cooled infrastructure.

New 8-GPU AI Servers from Lenovo > Lenovo Press

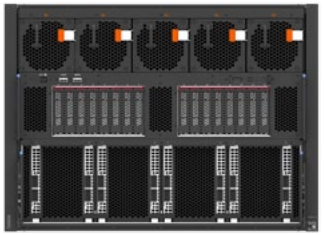

From Pocket to Cloud - AI for ALL

AI-optimized Portfolio from Model Development to Inferencing

From Enterprise, Cloud Service Providers and Telco Service Providers to Hyperscalers, Lenovo has a broad range of Servers from one to for sockets, from Edge to High Performance Computing, and from air to hybrid and liquid cooling systems.

To meet today's HPC requirements your servers should have a modern design. With Lenovo supercomputing servers, even the most technical simulations can be performed with gusto. These dense systems pack a punch for all your artificial intelligence, cloud, grid, inferencing and analytics needs. Designed to be scalable, unique water-cooled systems maximize density while putting security first.

For a broad view of the extensive server portfolio of Lenovo, please visit our portfolio guide page.

Lenovo Servers and Storage Portfolio Guide > Lenovo Press

Lenovo Storage Systems Built for Growth

Lenovo has a complete line of enterprise-grade storage solutions that adapt to your growing virtual environments, fit into your existing budget, and ensure data is ready when you need it.

- Unified storage
- Storage Area Network (SAN)
- Direct-attached storage
- Tape storage

New 6th-Gen Lenovo Neptune™ Liquid Cooling Powers the Era of AI

Boost AI and HPC performance with Lenovo Neptune™ Liquid Cooling. With over a decade of liquid-cooling expertise and more than 40 patents, we leverage our experience in large-scale supercomputing and AI to help organizations deploy high-performance AI at any scale.

Neptune™ | Data Center & GPU Liquid Cooling Technologies | Lenovo US

Adapt to Dynamic IT Needs with AI-based Smart Support with Lenovo XClarity

XClarity One is the next milestone in Lenovo's portfolio of systems management products. Now you can leverage the benefits of a true AI-powered, hybrid cloud-based solution for the deployment, management, and maintenance of your infrastructure through a single, intuitive interface that delivers a consistent user experience across all Lenovo products.

XClarity One is designed to know what you need before you need it and improve operational efficiency while keeping your data and devices secure from external threats.

XClarity One delivers:

- An Industry-First Hosted Hybrid Cloud Architecture
- Predictive Failure Analytics
- Risk Mitigation & Resiliency
- Cross-Platform Compatibility
- Authentication and Authorization:
- AI Customizable Insights and Reporting

www.lenovo.com/us/en/servers-storage/software/xclarityone

Lenovo Intelligent Computing Orchestration

Lenovo Intelligent Computing Orchestration (LiCO) is a software solution that simplifies the use of clustered computing resources for Artificial Intelligence (AI) model development and training, and HPC workloads. LiCO interfaces with an open-source software orchestration stack, enabling the convergence of AI onto an HPC cluster.

Lenovo Open Cloud Automation for AI

Lenovo Open Cloud Automation (LOC-A) is a software solution that rapidly deploys and manages the lifecycle of on-prem cloud infrastructure, regardless of its location: inside of a datacenter or distributed in your remote edge locations. LOC-A uses Infrastructure as Code and GitOps best practices for consistent and error-free results.

www.lenovo.com/us/en/servers-storage/software/open-cloud-automation/

Tried and Trusted Lenovo and Nokia for Data Center Solutions

Nokia and Lenovo's data center solution is built to support the capacity, scale, performance, and reliability for including the addition of AI workloads into your network. Combining data center networking and IP routing expertise from Nokia together with Lenovo's deep knowledge and experience in the high-performance data center server and storage systems provides the trusted infrastructure and the flexible foundation to grow with the increasing demands of AI data.

Benefits

- Orderable, standard server racks bundles validated with blueprint testing are readily available at a discounted package price
- A range of DC network management options – from DIY with open-source tools and open APIs to a complete life-cycle management and automation solution
- Field-proven networking stack for high-reliability, robustness
- High-availability and network uptime with built in component redundancy, active-active links
- Network pre-validation and restoration capabilities with digital twin and software version roll-back features
- Openness for configuration and observability
- Comprehensive open configuration API
- Streaming telemetry push based with more granular data
- Peace of mind with a trusted DC network (your own assets)

Lenovo's $1 Billion Dollar Investment in AI Innovation

Lenovo has made significant investments in AI infrastructure and innovation. They recently announced reaching a record annual AI infrastructure revenue of over **US$2 billion** and unveiled their next growth phase, committing an additional **US$1 billion** over the next three years to accelerate artificial intelligence (AI) deployment for businesses worldwide[1]. This investment aims to simplify the often-complex implementation of new AI capabilities by delivering AI directly to the source of data. Lenovo's industry-leading AI-ready portfolio includes smart devices, infrastructure solutions, and services that enable generative AI and cognitive decisions at scale across various sectors such as finance, manufacturing, healthcare, retail, and smart cities.

Lenovo's $100 Million Allocated to Grow AI Innovators Program

As part of this commitment, Lenovo is also allocating **US$100 million** to grow **the Lenovo AI Innovators program**, which has already delivered more than 150 cutting-edge AI solutions created in collaboration with 45 leading ISV partners in its first year1. These solutions cover areas like generative AI, computer vision, voice AI, and virtual assistants, helping businesses rapidly deploy and leverage advanced AI capabilities[1].

Lenovo's investment demonstrates their dedication to advancing AI infrastructure and supporting innovators in the field.

lenovo.com/nextlevelAI

Resources

Nokia Data Center Networking for AI workloads

Nokia Event-Driven-Automation

Nokia supercharges AI with Ultra Ethernet

lenovo.com/nextlevelAI

Nokia OYJ
Karakaari 7
02610 Espoo
Finland

Tel. +358 (0) 10 44 88 000

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.

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.