Nokia and Lenovo Sovereign AI Data Center Blueprint Solution

Empowering Sovereign AI with scalable, secure and high-performance infrastructure

White paper

Contents

Design-optimized data center network architecture for Sovereign AI

As AI becomes more pervasive in our lives, its data sources, data rights/permissions, where it is stored and where it can be used has become a critical determinant of where training and inference models are built and operated. Nation states and regulators are continuing to define new policies and laws that protect the rights of the data owners and ensure the security and privacy of that data. Gartner predicts that by 2028, **65% of governments worldwide will introduce technological sovereignty requirements** to improve independence and protect from extraterritorial regulatory interference, highlighting the urgency for infrastructure that supports Sovereign AI.

New Sovereign AI clusters and workloads will require a new generation of high-performance servers with higher density CPU and GPU processing together with bandwidth capacity, scalability and improved performance of the date center network. These workloads can also be distributed across cloud networking infrastructures which requires network operations management visibility and automation capabilities that cut across network boundaries.

Nokia and Lenovo have architected, tested and validated a Sovereign AI data center networking blueprint solution that's built to meet the demanding processing, network performance and security requirements for sovereign AI. It combines the high-performance Lenovo ThinkSystem line of servers and storage systems together with the Nokia Data Center Fabric solution.

Sovereign AI

Sovereign AI can be defined as AI systems that are designed, developed and deployed under the control and jurisdiction of a specific nation, organization or enterprise that removes any dependency of a foreign nation, technology, platform or outside influence.[1]

It's a strategic response to the growing importance of AI in national security, economic competitiveness, and digital sovereignty.

Key concepts behind Sovereign AI

1. **National control and governance:** Sovereign AI means having full control over the entire AI lifecycle— from data collection, data storage and model training to deployment and governance—using domestic infrastructure, internal resources and data.[2]

2. **Strategic resilience:** In order to reduce reliance on foreign AI systems, nations need to prepare and adapt their own AI System to protect against geopolitical risks, potential security or system vulnerabilities in their critical infrastructure.[3]

3. **Ethical and regulatory alignment:** Sovereign AI systems are designed to comply with local laws, ethical standards, strategic priorities and core values, ensuring that AI technologies serve national interests.[4]

4. **Economic and innovation boost:** By investing in domestic AI capabilities, nations can accelerate innovation, create high paying jobs that attracts new talent in order to remain competitive in the global digital economy.[5]

---

1 What Is Sovereign AI? | NVIDIA Blog
2 (Dale, 2025)
3 (Alduhishy, Muath 2024)
4 (Sideco, 2025)
5 (Alduhishy, Muath 2024)

Sovereign AI use cases

Lenovo and Nokia have designed and optimized a high-performance, scalable AI pod blueprint solution that allows Government, Service Providers and Enterprises to quickly deploy and turn-up new AI services and use cases. Figure 1 provides examples of the industry segments and use cases for sovereign AI:

Figure 1. Sovereign AI use cases across vertical segments

Financial Services	Life Sciences and Healthcare	Research and Education	Media and Entertainment	OT Enterprise	Government/ Public Safety
• Real-time fraud detection • Conversational chatbots • Automated loan processing • Robo-advisors • Compliance monitoring	• Clinical trial optimization • Drug discovery • Enhanced patient risk management • Medical imaging • Genomic data	• Personalized education programs • HPC data analysis • Quantum physics • Astronomy/ astrophysics	• Personalized content recommendations • Copyright infringement • Real-time translation/ dubbing	• Oil & gas exploration • Transportation/ logistics optimization • Enhanced air traffic control • Nuclear-based power plant	• Seismology • Climatology • Crime prediction • Disaster simulation and planning • Planetary exploration • Defense R&D • Military strategic planning

Key Sovereign AI DC network design considerations

Designing networks for **Sovereign AI** involves building infrastructure that ensures **data control**, **security**, **performance** and **compliance** within national borders. According to IDC research commissioned by Lenovo, data sovereignty, compliance and availability are the top success factors for AI implementation among CIOs.

Here are the **key network design considerations** and why they matter:

Data sovereignty and path control. Sovereign AI must ensure that data used for training and inference does not leave authorized jurisdictions.

High-performance AI. Fabric AI workloads require massive parallel processing and low-latency communication between GPUs and nodes.[6]

Distributed workload. AI training and inference involve large datasets and distributed compute nodes.

Security and compliance. Sovereign AI must comply with national regulations and protect against cyber threats.

Modular and scalable architecture. Sovereign AI must support diverse use cases and scale rapidly.

Host management and orchestration. Smooth operation and debugging are essential for AI lifecycle management.

Federated and hybrid models. Full sovereignty may not be feasible for all nations.

---

6   (Ray, 2025).

# Nokia and Lenovo Sovereign AI data center blueprint solution

Nokia and Lenovo are partnering to deliver high-performance, scalable, and secure blueprint solution for AI/ML data center combining Nokia's Data Center Fabric and Event-Driven Automation with Lenovo's ThinkSystem AI-ready portfolio to meet the demands of modern workloads.

This blueprint solution is built in modular AI pod units that support a wide range of AI applications or workload processing, footprint requirements or budget.

The basic AI pod unit consists of high-performance AI GPU servers, storage systems and management equipment together with a rail-optimized switching network. The AI pods units can be scaled out to support hundreds of servers and petabytes of storage.

Shown in figure 2 is an example of an AI pod unit that supports 64 GPUs.

Figure 2. AI-DC Blueprint Solution Equipment Rack Layout example for 64 GPUs
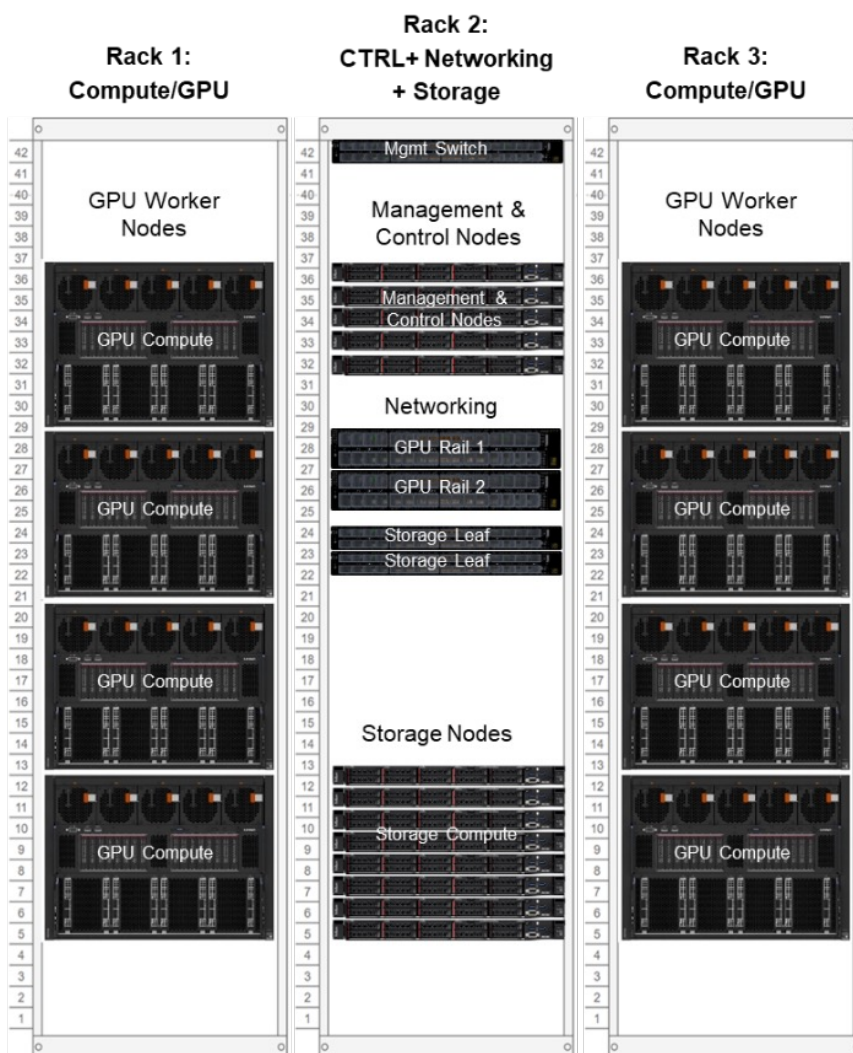
Nokia-Lenovo Sovereign AI Validated Design

The Lenovo-Nokia Validated Design reference architecture is composed of three main networking segments. The backend network, the backend storage network, and the combined front-end and management network (figure 3).

Figure 3. Lenovo, Nokia Sovereign AI Validated Design Architecture
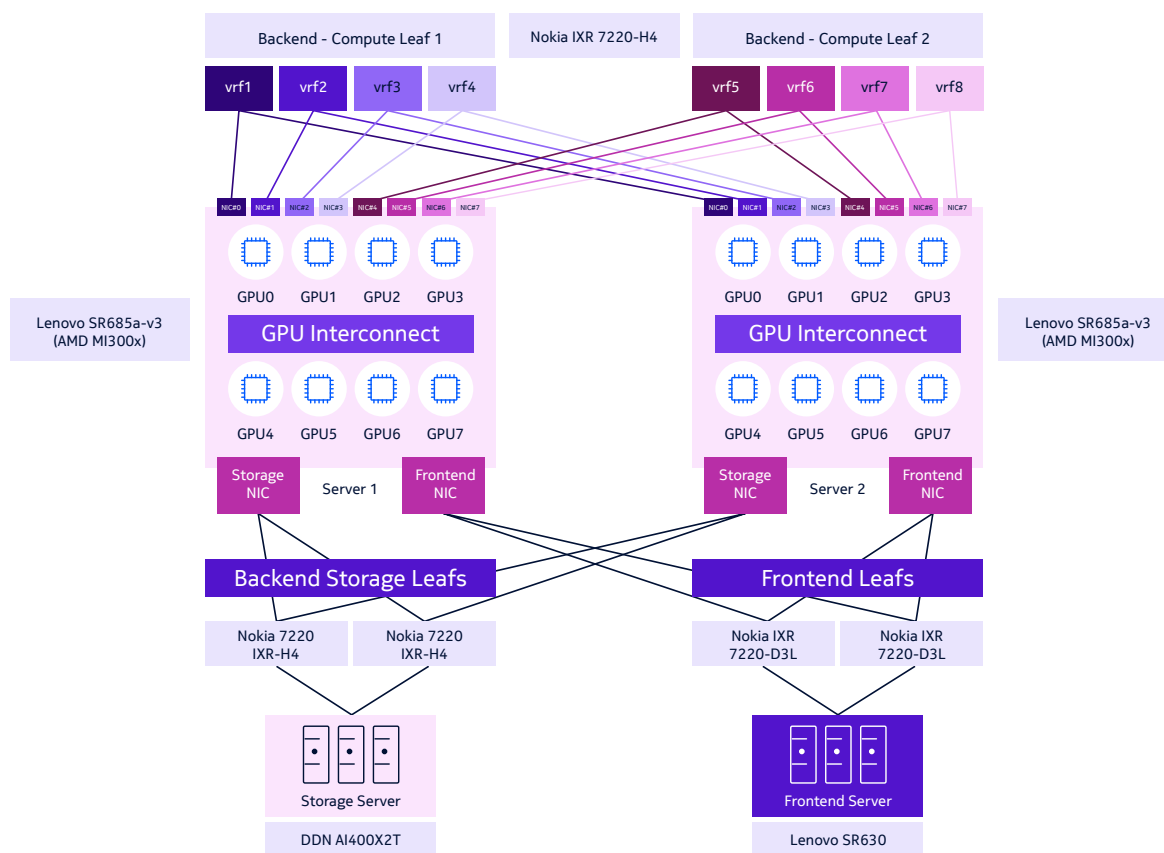
The backend "scale-out" network supports GPU to GPU connectivity. It consists of Lenovo SR635 air-cooled AI GPU servers that host AMD Instinct GPUs and 400G (Gig) NICs that are interconnected with Nokia 7220 IXR-H4 GPU leaf switches. This design ensures that only a single hop is needed for GPU to GPU inter server communication.

The backend storage network Nokia 7220 IXR-D5 high-performance switches that interconnects to Lenovo SR630 servers and DDN AIX400 storage nodes.

The converged front end and in-band management network is represented by Lenovo SR630 GPU and management servers connected through Nokia 7220 IXR-D3L switches.

Nokia's Event Driven Automation (EDA) platform was used to simplify the configuration, validation and operational aspects of managing an AI cluster. This included the use of its back-end configuration application to design, configuration and deployment of the back-end fabric including multi-tiered architectures, multi-tenancy, IP addressing and network parameter settings such as for congestion avoidance and traffic prioritization.

Lenovo-Nokia validated design results

Lenovo and Nokia ran an extensive suite of MLCommons.org performance tests on the referenced Sovereign AI design. These include:

- MLPerf tests for standard RoCEv2(Rocky v2) packet sizes ranging from 256–4096 bytes. The testing results met or exceeded industry standards for the average bi-directional bandwidth achieved across 400Gb (gigabits) interfaces.

- MLPerf tests were also run for maximum RoCEv2 (Rocky v2) packet size of 4096 bytes. The testing results met or exceeded industry standards for the average bi-directional bandwidth achieved across all 400Gb Interfaces.

- MLCommons benchmark testing was run on the performance of a 70 billion Llama 2 training and inference models that measured both throughput and latency metrics.

The Nokia Data Center Validated Design Hub is a repository of rigorously tested, multi-vendor network designs to ensure performance and reliability standards are met as well as providing a template for rapid deployment and ease of operations. The Lenovo-Nokia AI Data Center Design and test results are included here.

Blueprint solution business benefits

The Nokia and Lenovo data center networking solution has been rigorously tested and validated to support the capacity, scale, performance, and reliability required for integrating AI workloads into enterprise environments. It combines Nokia's expertise in data center networking and IP routing with Lenovo's deep experience in high-performance servers and storage systems.

Key business benefits for enterprise:

- **Pre-validated, orderable server rack bundles**
  Ready for deployment and available at discounted package pricing to accelerate time-to-value.

- **Flexible network management options**
  From DIY setups using open-source tools and APIs to full lifecycle automation via integrated orchestration platforms.

- **Carrier-grade networking stack**
  Proven reliability and robustness for mission-critical AI workloads.

- **High availability and uptime**
  Enabled by built-in component redundancy and active-active link configurations.

- **Digital twin and rollback capabilities**
  Support pre-validation, version control and rapid restoration in complex environments.

- **Open configuration and observability**
  Designed for transparency, customization and operational control.

- **Comprehensive API support**
  Facilitates seamless integration across compute, storage and networking layers.

- **Streaming telemetry with granular insights**
  Enables real-time monitoring and proactive performance optimization.

- **Infrastructure sovereignty and control**
  Peace of mind with a trusted data center network built on your own assets.

Why Nokia and Lenovo for your Sovereign AI data center network?

Lenovo and Nokia together bring together a Sovereign AI solution that is unmatched in best-of-breed technology, services and support.

- **Open multi-vendor eco-system solution**
  Providing flexibility of choice to build a solution tailored to your own needs.

- **GPU/CPU chipset agnostic**
  Multi-GPU support: AMD/NVIDIA/Other GPUS.

- **Built with industry standard specifications**
  No proprietary protocols: based on standard RoCEv2/future UEC protocols.

- **Certified full AI DC solution**
  Blueprints that are rigorously tested and jointly validated in the labs.

- **Scalable AI pod configurations for simplified ordering**
  Pre-defined full-stack AI blueprint solutions that take the guess-work out of ordering.

- **Network customization and operation the way you want it**
  Open-APIs to interface both with the server and networking layers.

- **Rapid trouble-resolution and support**
  Global tech support to answer your call, wherever you are located.

- **End-to-end data center solution across your network**
  Seamless integration between the data center and the Wide Area Network (and across DCs).

- **Industry leaders in AI, HPC Data Centers and IP/Optical**
  Lenovo is deployed in over 25% of the top 500 Supercomputers in the world and Nokia is #1 market-share leader in IP edge routing worldwide.

- **Advanced automation with AI for data center operations**
  Network Automation to gain insights, simplify operations and reduce human errors.

- **Power-efficient platforms**
  Optimized designs for lower power consumption.

- **Best-of-breed data center technologies**
  Deep Networking Expertise from Nokia combined with advanced AI Server technologies from Lenovo.

- **Carrier-grade reliability**
  Track record of reliability from Lenovo and Nokia with field-proven, carrier-grade systems that combined to deliver a powerful, reliable and resilient solution.

- **Global presence and expertise**
  Expert teams present in the field to engage with customers and partners to identify the best solutions and support with any technical topics.

Learn more

Visit our website or talk to one of our experts.

Nokia Data Center Networks →

Talk to a Nokia expert →

Lenovo AI Solutions →

Lenovo Data Center Support →