



Nokia and Lenovo Sovereign AI Data Center Blueprint Solution Solution Brief

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 data 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.

Definition of 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. For details, see the NVIDIA blog post, [What Is Sovereign AI?](#).

Key concepts of Sovereign AI include:

- **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. ([Dale, 2025](#))
- **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. ([Alduhishy, Muath, 2024](#))
- **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. ([Sideco, 2025](#))
- **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 ([Alduhishy, Muath, 2024](#))

Sovereign AI is a strategic response to the growing importance of AI in national security, economic competitiveness, and digital sovereignty.

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. The following table shows examples of the industry segments and use cases for sovereign AI.

Table 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/astro-physics	- 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 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:

1. **Data Sovereignty & Path Control.** Sovereign AI must ensure that data used for training and inference does not leave authorized jurisdictions.
2. **High-performance AI Fabric.** AI workloads require massive parallel processing and low-latency communication between GPUs and nodes. ([Ray, 2025](#))
3. **Distributed Workload.** AI training and inference involve large datasets and distributed compute nodes.
4. **Security & Compliance.** Sovereign AI must comply with national regulations and protect against cyber threats.
5. **Modular & Scalable Architecture.** Sovereign AI must support diverse use cases and scale rapidly.
6. **Host Management & Orchestration.** Smooth operation and debugging are essential for AI lifecycle management.
7. **Federated & Hybrid Models.** Full sovereignty may not be feasible for all nations.

Nokia and Lenovo Sovereign AI Data Center Blueprint Solution

Nokia and Lenovo have [partnered](#) 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.

The following figure shows an example of an AI pod unit that supports 64 GPUs.




Figure 1. 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, as shown in the following figure.




Figure 2. Lenovo, Nokia Sovereign AI Validated Design Architecture

The major components in the architecture are as follows:

- The backend “scale-out” network supports GPU to GPU connectivity. It consists of [Lenovo ThinkSystem SR685a V3](#) 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 ThinkSystem SR630 V3](#) servers and [DDN AIX400](#) storage nodes.
- The converged front end and in-band management network is represented by [Lenovo ThinkSystem SR630 V3](#) GPU and management servers connected through [Nokia 7220 IXR-D3L](#).
- [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.

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 packet sizes ranging from 256 to 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 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 test results are included in the [Lenovo-AMD-Nokia AI-DC Validated Design document](#).

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 enterprises include the following:

- **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](#)

References

For more information, see these references:

- Dale, Roberts: Cambridge University Press "Sovereign AI in 2025", August 2025.
<https://www.cambridge.org/core/journals/natural-language-processing/article/sovereign-ai-in-2025/C51560626AF518BDF280891C406E9553>
- Mauth Alduhishy: World Economic Forum "Sovereign AI; What is it? And 6 strategic pillars for achieving it", April 2024.
Muath Alduhishy - Agenda Contributor | World Economic Forum
<https://www.weforum.org/stories/2024/04/sovereign-ai-what-is-ways-states-building/>
- Francis Sideco: Forbes "The AI Era Enters Its Sovereign Phase", June 2025.
<https://www.forbes.com/sites/tiriasresearch/2025/06/03/the-ai-era-enters-its-sovereign-phase/>
- Trisha Ray, "Sovereign remedies: Between AI autonomy and control" Atlantic Council, April 2025.
<https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/sovereign-remedies-between-ai-autonomy-and-control/>

Authors

Mark Wallis is a Senior AI Solutions Architect at Lenovo. He is passionate about building and testing the latest AI solutions to meet customers' needs with a background in managing complex installations, performance testing and software automation. He is originally from the United Kingdom, but has been based in the USA for the last 25 years. When he is not working, he can be found spending time with his family or out and about, hiking and cycling.

David Nowoswiat is a Senior Product and Solutions Marketing Manager at Nokia and has over 25 years of telecom industry experience in both wireless and wireline technologies. He is currently in Nokia's IP and Optical Business Group supporting Cloud Packet Core Marketing. He holds a Bachelor in Industrial Engineering degree from Georgia Institute of Technology and a Masters in Business Administration from Northeastern University.

Miguel Simões is the Alliance Director for the Global Nokia Partnership in the Lenovo Infrastructure Solutions Group. With more than 28 years of industry experience, Miguel is a veteran of the telecommunications world, now leading Lenovo's global alliance with Nokia. In his current role, he drives strategic alignment, joint innovation, and business acceleration across enterprise and telco ecosystems worldwide. Miguel's expertise spans operations, engineering, pre-sales, and business development, with previous leadership roles at Nokia where he led SD-WAN and data center transformation initiatives in Europe. Known for his customer-centric vision and technical depth, Miguel brings a collaborative spirit and global perspective to the stage. He holds a BSC in Electronics and Telecommunications Engineering and is a passionate advocate for agile, sustainable infrastructure

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2339, was created or updated on November 24, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2339>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2339>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

The following terms are trademarks of other companies:

AMD and AMD Instinct™ are trademarks of Advanced Micro Devices, Inc.

Other company, product, or service names may be trademarks or service marks of others.