



Solution Brief

Accelerating AI Networking with NVIDIA Spectrum-X™ and Aviz ONES

NVIDIA Spectrum-X™ Overview

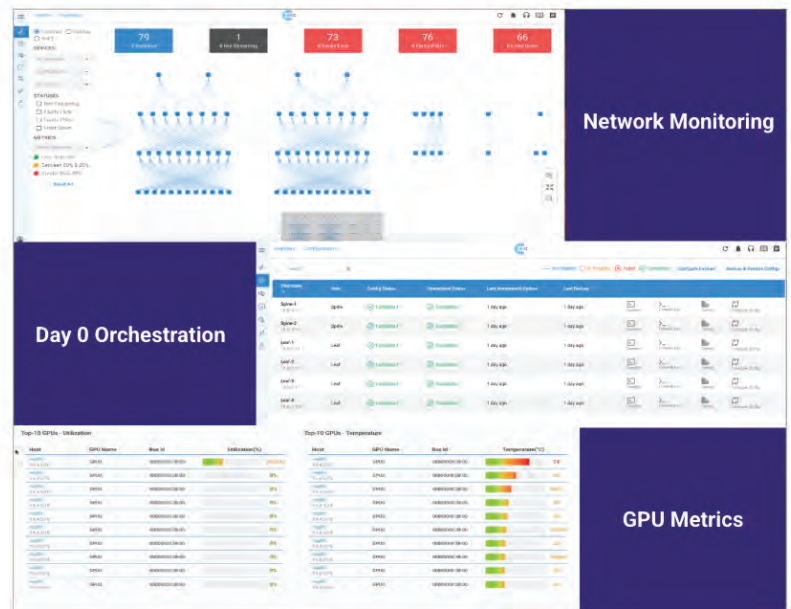
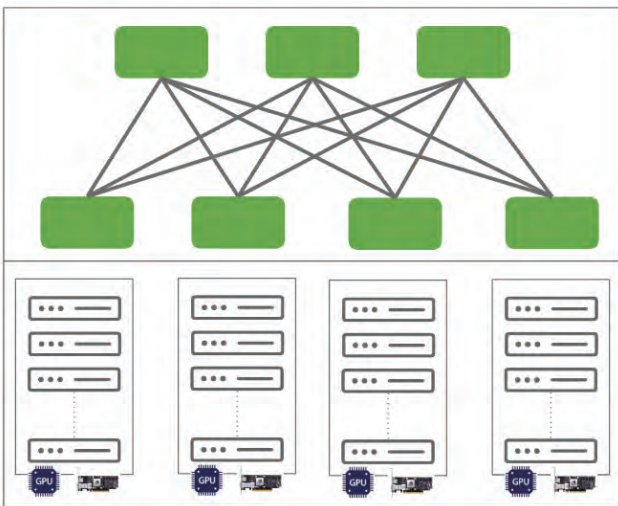
The NVIDIA Spectrum-X™ platform, combined with NVIDIA HGX H100 systems, represents the next-generation architecture for AI cloud infrastructure. Designed to deliver the compute performance needed for solving complex AI challenges in cloud environments, Spectrum-X™ serves as the foundation for AI cloud production clusters when integrated with HGX.

To ensure maximum performance and scalability, the HGX Spectrum-X™ platform leverages several key NVIDIA technologies, including:

1. NVIDIA HGX H100, the most powerful computational building block for AI.
2. Spectrum-4 switch and NVIDIA BlueField-3 SuperNIC, which operate together with SDKs and software to enable full-stack AI cloud network optimization.

NVIDIA Spectrum-X™ and Aviz ONES

The NVIDIA Spectrum-X™ reference architecture defines integration requirements to ensure a straightforward design for customers, enabling end-to-end Day 0 orchestration of network configuration, including IP Fabric and RDMA profiles for both single and multi-tenant environments. Additionally, the integration should support Day 2 operations by handling fabric failures, validating configuration integrity, and facilitating future GPU scaling.



Aviz ONES delivers full-stack observability and automation for NVIDIA Spectrum-X™ Ethernet fabrics, ensuring high-performance AI workload orchestration. Built with an agentless architecture and containerized microservices, ONES simplifies network deployment, monitoring, and scaling for AI-driven infrastructures.

Aviz ONES integrates seamlessly with NVIDIA Spectrum-X™, Spectrum SN5600 switches, and BlueField-3 SuperNICs, enabling a lossless RDMA over Converged Ethernet (RoCEv2) fabric optimized for multi-tenant AI clusters. By reducing network congestion, enhancing visibility, and automating operations, ONES ensures maximum GPU utilization and predictable AI performance—key for enterprises scaling AI workloads.

Aviz ONES Capabilities

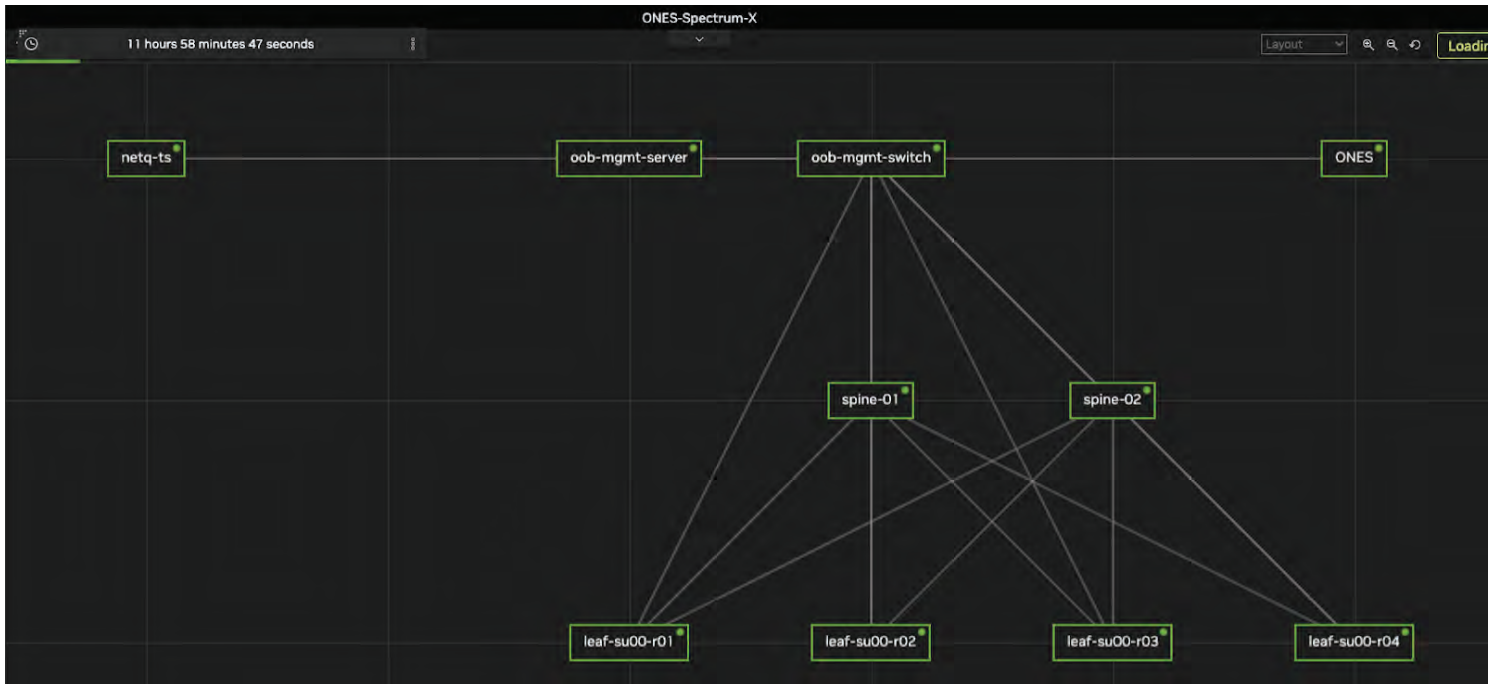
Orchestration	<ul style="list-style-type: none"> • Network Design • YAML Templates • Validate, Apply and Verify Configuration
Network Telemetry	<ul style="list-style-type: none"> • End-2-End Monitoring [Switches NIC GPUs] • Inventory, Health and Traffic • RoCE Telemetry [PFC, ECN, Queue Counters] • GPU Metrics [Performance and Utilization]
Day 2 Operations	<ul style="list-style-type: none"> • Alerting and Integration with Customer tools • Support and Troubleshooting

Integration Benefits

- Seamless AI Fabric Deployment – Automate Day-0/1/2 network provisioning, including switch, NIC, and server configurations.
- Optimized GPU Performance – Ensure high-bandwidth, low-latency AI training with real-time RoCE telemetry, congestion control, and anomaly detection.
- Open-Source – Built on open networking principles, integrating with existing NetOps and customer tools.

- Enterprise-Ready Observability – Track over 250+ unique network, server, and GPU metrics to detect performance bottlenecks before they impact workloads.
- Simplified Management – Automate switch RMAs, config sanity checks, and scaling of GPUs across single- and multi-tenant AI clusters.

Aviz ONES for NVIDIA Spectrum-X™: How It Works



A 3-step workflow for ONES network deployment and operations in an NVIDIA Spectrum-X™ environment:

- **Network Design:** Define GPU requirements, choose between single or multi-tenant deployment, create a ONES blueprint, simulate the setup in NVIDIA AIR, and ensure network services like NTP and ZTP are configured.
- **Configuration (Generate | Apply | Verify):** Generate a ONES template, apply network configurations, validate them for errors, and verify end-to-end operational behavior.
- **Monitoring & NetOps:** Continuously monitor key metrics across Switches, NICs, and GPUs, set up alerts for failures and anomalies, integrate with customer tools, and establish workflows for Day 2 operations.

This structured approach ensures efficient deployment, validation, and monitoring of network infrastructure for GPU workloads.

Ready to transform your AI networking?
Schedule a demo with Aviz Networks today:

**Aviz ONES for
AI-Ready Networking**

