Solution Guide

Solving AI Infrastructure Challenges with Out-of-Band Management

Artificial intelligence is being rapidly adopted and changing the way data centers are designed. ***Per Gartner, 70% of enterprise executives are exploring generative AI for their organizations. The infrastructure that supports AI must meet increasing demands for performance, reliability, and security.***

However, managing this infrastructure brings several challenges that, if not addressed, can lead to operational inefficiencies, increased costs, and compromised performance.

The Challenges of AI Infrastructure Management

AI workloads are resource-intensive and require a complex ecosystem of servers, GPUs, power distribution units (PDUs), and network switches. To effectively manage this infrastructure, data center operators need to overcome three major challenges:

1. Resource Utilization and Power Efficiency

AI workloads require computing and power resources that work at peak efficiency. With modern AI racks consuming 40kW to 200kW of power, organizations can easily incur waste and increased costs when their GPUs, CPUs, and PDUs are underutilized or poorly managed. In large-scale AI environments, every component needs to operate at peak efficiency to maximize performance and ROI.

2. Downtime and Performance Bottlenecks

AI systems are highly sensitive to downtime and latency, with each server relying on the underlying infrastructure of networking and power equipment. When a network switch fails or a PDU experiences a fault, the negative impact has a cascading effect. One server failure affects the entire cluster, interrupting critical workloads for model training or inference. Bottlenecks in latency-sensitive AI applications like real-time analytics or autonomous systems can degrade performance and results.

3. Security and Compliance Risks

AI environments process sensitive data, from proprietary algorithms to regulated personal information. If the management interfaces of network switches or PDUs are exposed to the production network, the overall attack surface increases and security risks grow. Breaches and ransomware easily compromise critical systems and encrypt management access from IT admins, resulting in data theft, system manipulation, or regulatory non-compliance, as seen in many recent cyberattacks.

The Gap in Traditional Management Solutions

Traditional in-band management solutions rely on the primary network to access and control infrastructure components. While this approach works under normal conditions, it has serious limitations when addressing distributed deployments, network outages, or security breaches.

- **Limited Visibility:** In-band solutions often lack comprehensive monitoring capabilities, making it difficult to proactively identify and address issues before they impact performance.
- **Network Dependency:** In-band management tools are useless when the primary network is down. If there's an outage or an admin makes a critical error, teams are cut off.
- **Security Exposure:** Managing switches, PDUs, and compute resources over the same network used for production traffic increases the risk of cyberattacks. This approach is banned by CISA's latest best practice recommendations.

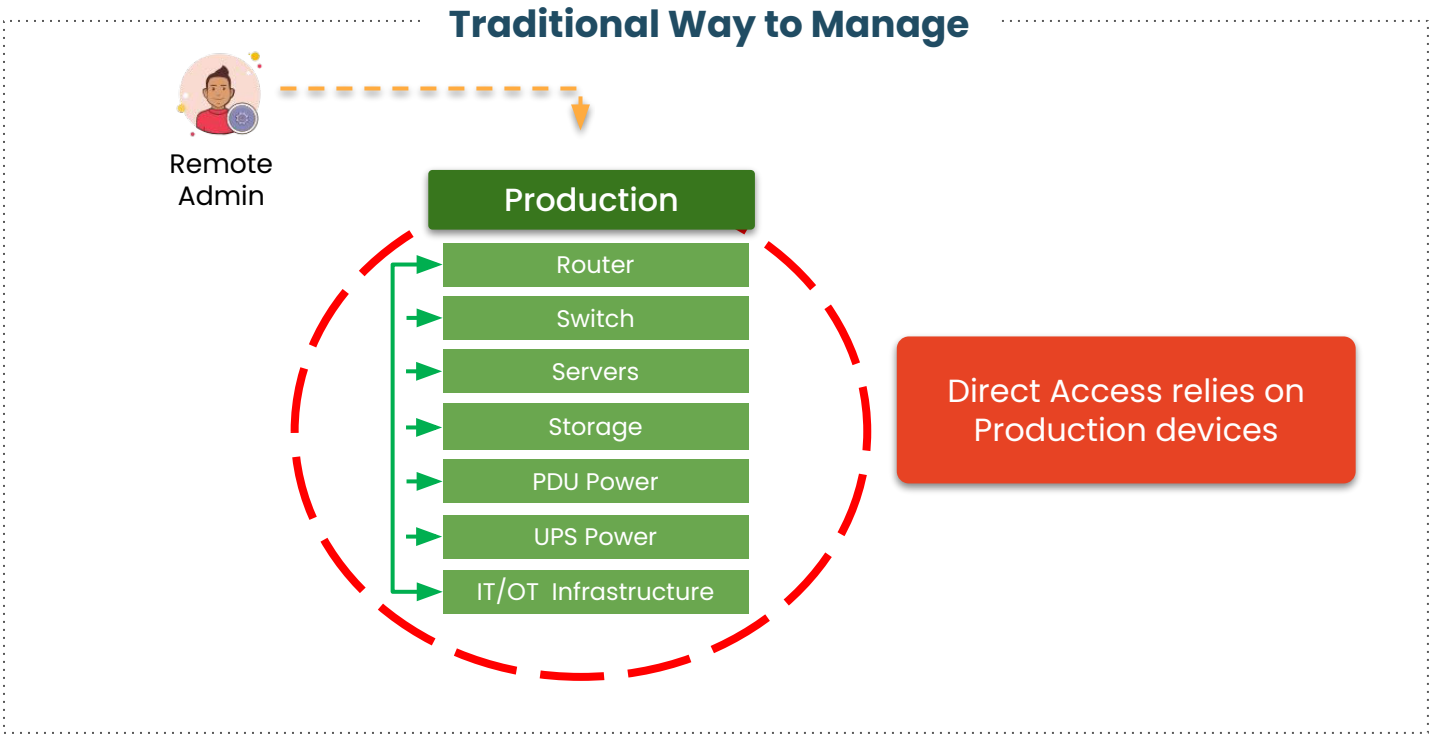

These limitations create operational blind spots and increase the risk of inefficiencies, downtime, and security breaches.

Solution: Out-of-Band Management

Out-of-band (OOB) management provides a dedicated, independent pathway to monitor and control AI infrastructure. It is fully separate from production equipment and provides complete management capabilities even if there is a production failure or outage.

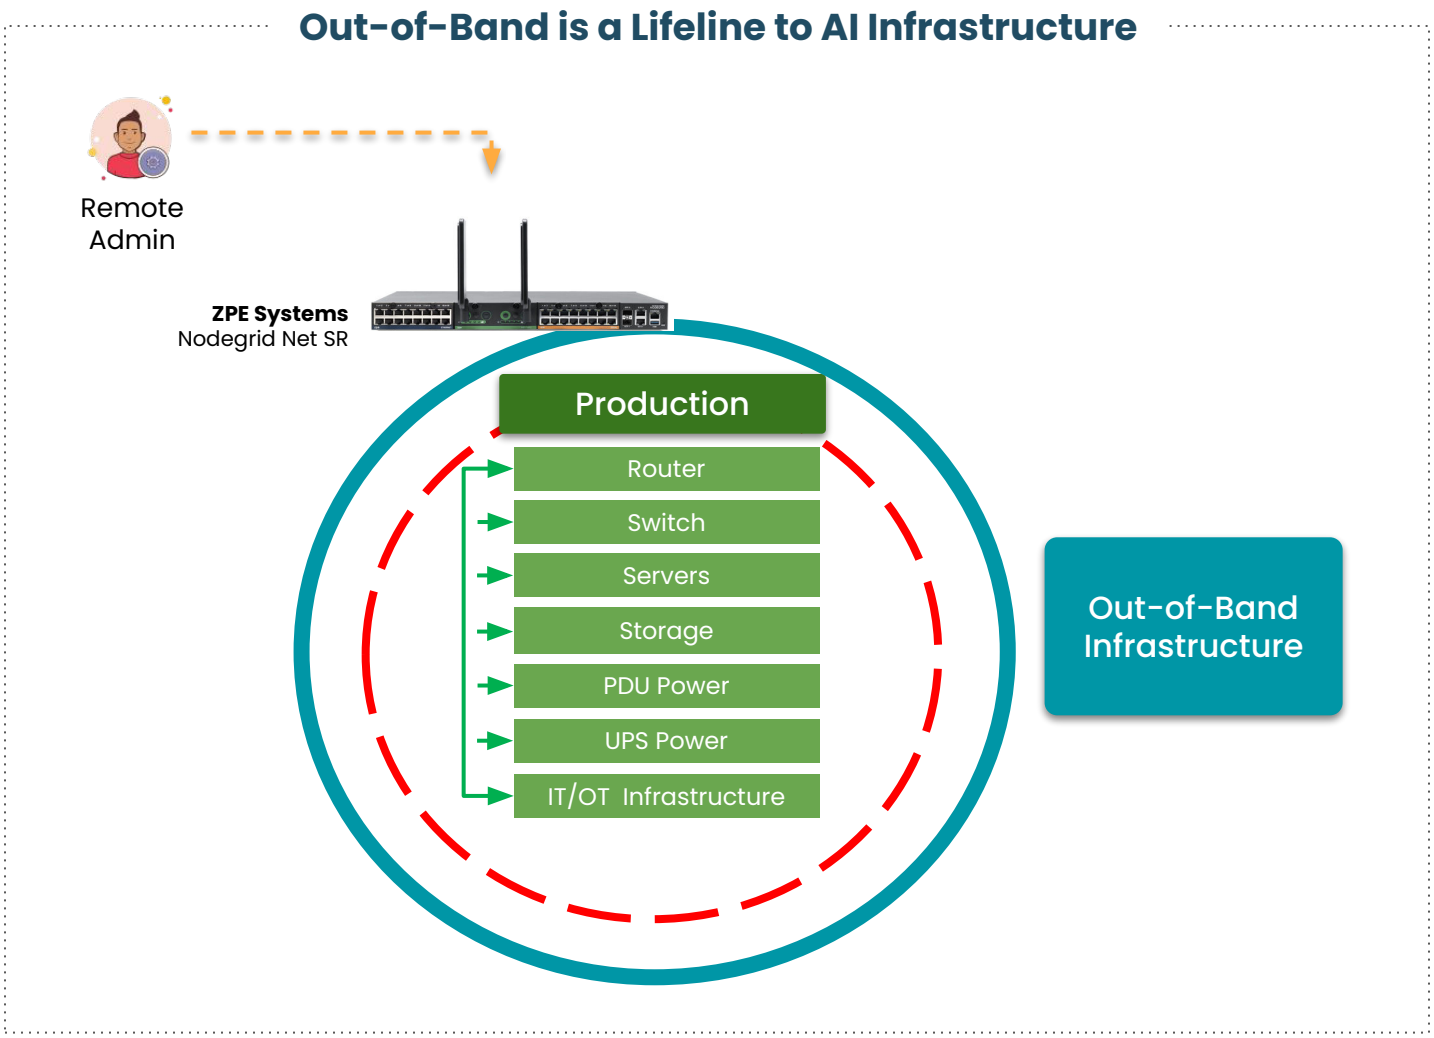

Key Benefits of Out-of-Band Management

1. Optimized Resource and Power Management

With OOB, administrators gain real-time visibility into resource utilization and power consumption. This allows for dynamic resource allocation and proactive power management, ensuring that GPUs, CPUs, and other hardware operate efficiently. Optimizing power distribution through PDUs helps to reduce energy costs and prevent overloads.

2. Uninterrupted Access During Network Failures

OOB management operates independently of the primary network, ensuring continuous access to infrastructure components even during outages. This capability allows administrators to troubleshoot and resolve issues remotely, minimizing downtime and maintaining system availability.

3. Enhanced Security and Compliance

OOB solutions isolate management traffic from production networks, reducing the attack surface and protecting critical systems from unauthorized access. This Isolated Management Infrastructure is called out in CISA's binding operational directive 23-02 as the best practice to secure management interfaces. Advanced security features such as role-based access control and encrypted communication help organizations meet regulatory compliance requirements while safeguarding sensitive data.

4. Streamlined Operations with Remote Monitoring

OOB management enables remote monitoring and troubleshooting for all infrastructure components, including switches, PDUs, and other infrastructure that supports AI. IT teams can diagnose and resolve issues quickly without the need for on-site visits. This is particularly valuable in distributed AI environments, where reducing mean time to recovery (MTTR) by minutes can impact the entire operation.

5. Proactive Issue Resolution

Advanced OOB solutions provide predictive analytics and automated alerts, helping administrators identify potential issues before they escalate. This proactive approach reduces downtime, enhances system reliability, and improves overall infrastructure performance.

Why Choose ZPE Systems for Out-of-Band Management of AI?

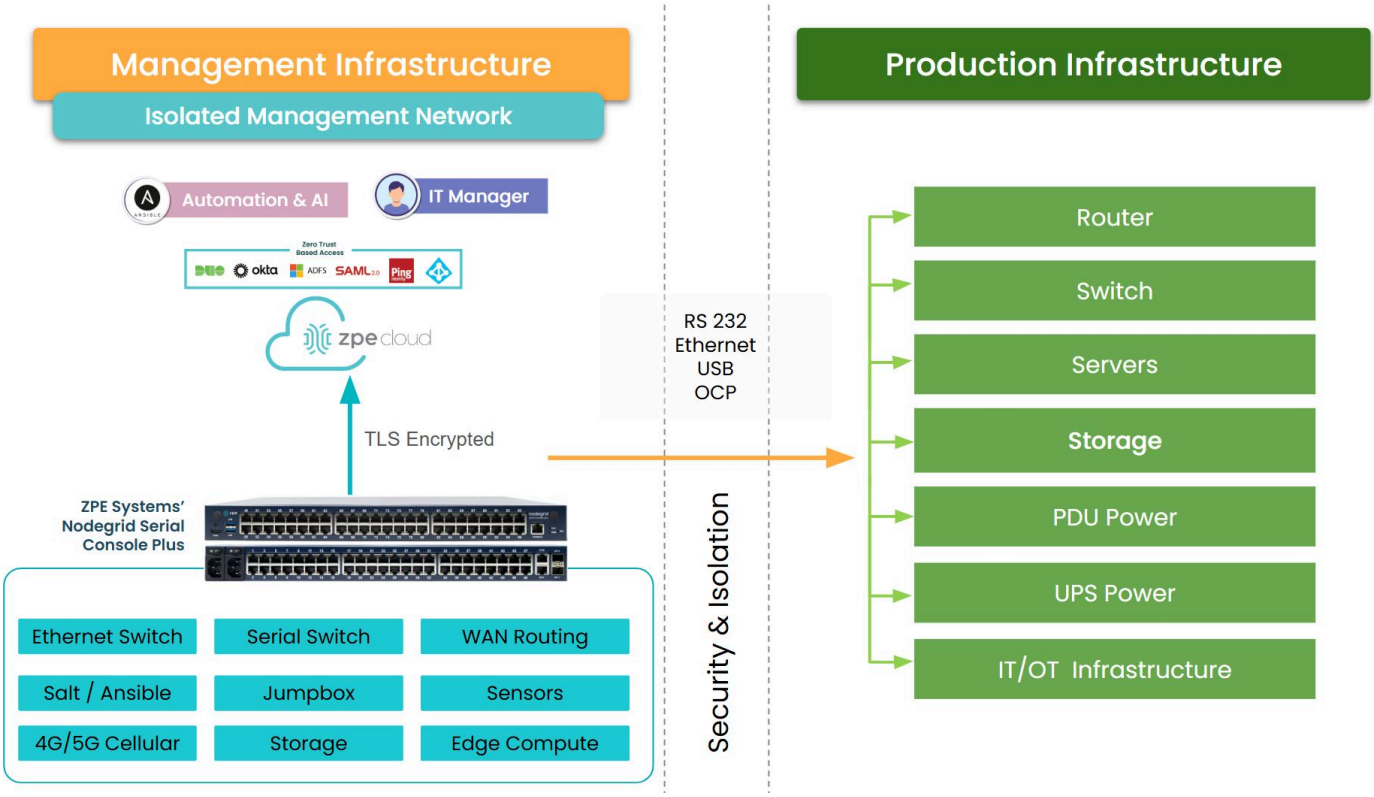

ZPE Systems' Nodegrid combines up to nine functionalities into a 1RU device. This provides a complete out-of-band infrastructure that can drop into any environment, so organizations can manage and optimize the systems that support their AI deployments. Here are a few reasons why Nodegrid is used by hyperscalers and global enterprises:

1. One-Stop Shop

Nodegrid is the only platform to combine complete management infrastructure in a single serial console device, including programmable software and cloud management. This minimizes capex and opex by eliminating the need to deploy separate devices for switching, routing, 5G, storage, automation, and application hosting.

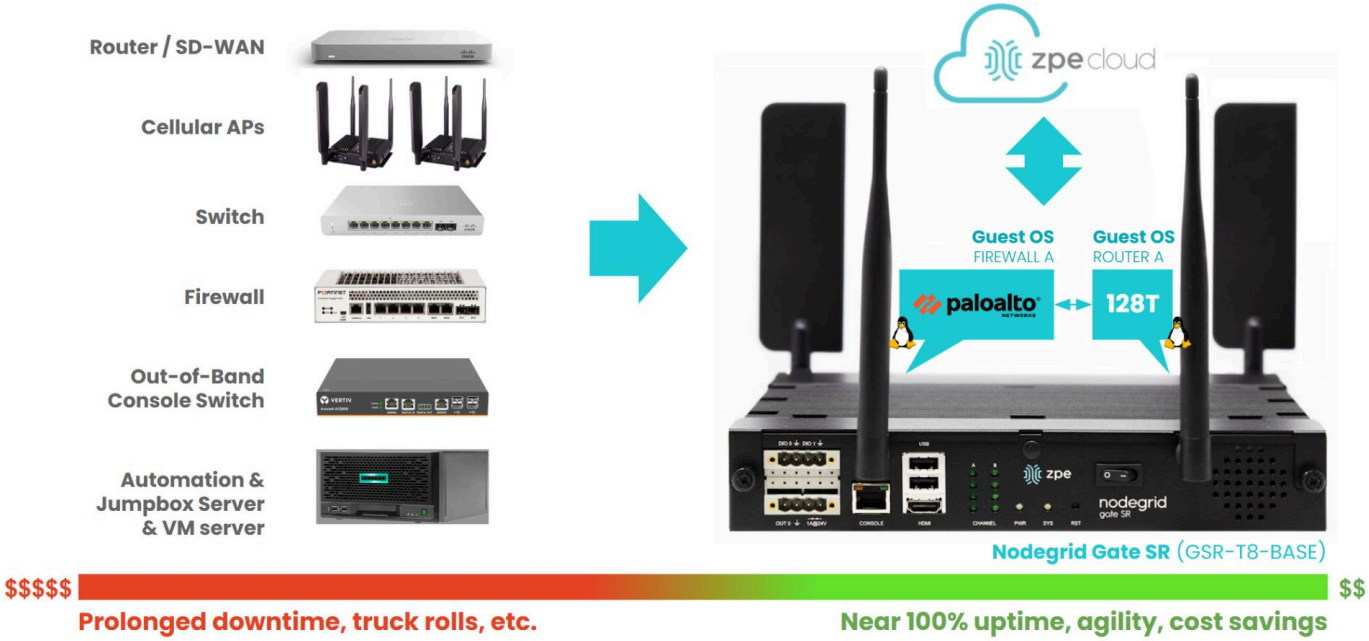

2. Security Best Practices

Nodegrid not only follows security best practices of fully isolating management interfaces; it's the industry's most secure out-of-band platform. Nodegrid is FIPS 140-3, SOC 2 Type 2, and ISO27001 certified, has a Synopsys-validated codebase, and incorporates dozens of hardware and supply-chain security features.

		ZPE Systems	Others
Security Integrations	• CyberArk, Delinia, Horizon3.ai • PaloAlto, Fortinet, Cloudflare	✓ ✓	○ ○
Certification and Processes	• SOC2 Type 2, FIPS 140-3 • PSIRT, Pentesting	✓ ✓	○ ○
Software & Cloud	• Latest Kernel and CVE patches • Zero Trust based access • SAML2 based SSO • MFA • Latest encryption standards	✓ ✓ ✓ ✓ ✓	○ ○ ○ ○ ○
Software Development	• Dynamic Code Analysis • Static Code Analysis • Software BOM analysis including Open Source Software Composition • Continues Security Assessments • Zero CVE Policy • Vulnerability Scan	✓ ✓ ✓ ✓ ✓	○ ○ ○ ○ ○
Hardware	• Secure Signed OS • Password-protected BIOS and Boot Loader • TPM 2.0 • Self Encrypted Disk • Secure Boot • Geo Fencing Protected	✓ ✓ ✓ ✓ ✓ ✓	○ ○ ○ ○ ○ ○

3. Full Integration and Remote Access

Nodegrid is vendor-neutral, meaning it integrates with and enables control of many device types via serial, Ethernet, USB, KVM, or IPMI. Teams use it to remote-into any underlying infrastructure, and it gives them the ability to perform tasks ranging from a simple power cycle, to a full re-image and rebuild of systems. In AI environments, this gives teams secure, isolated management access to troubleshoot and optimize switches, servers, and PDUs, as well as collect and analyze data passed along from temperature sensors inside each rack.

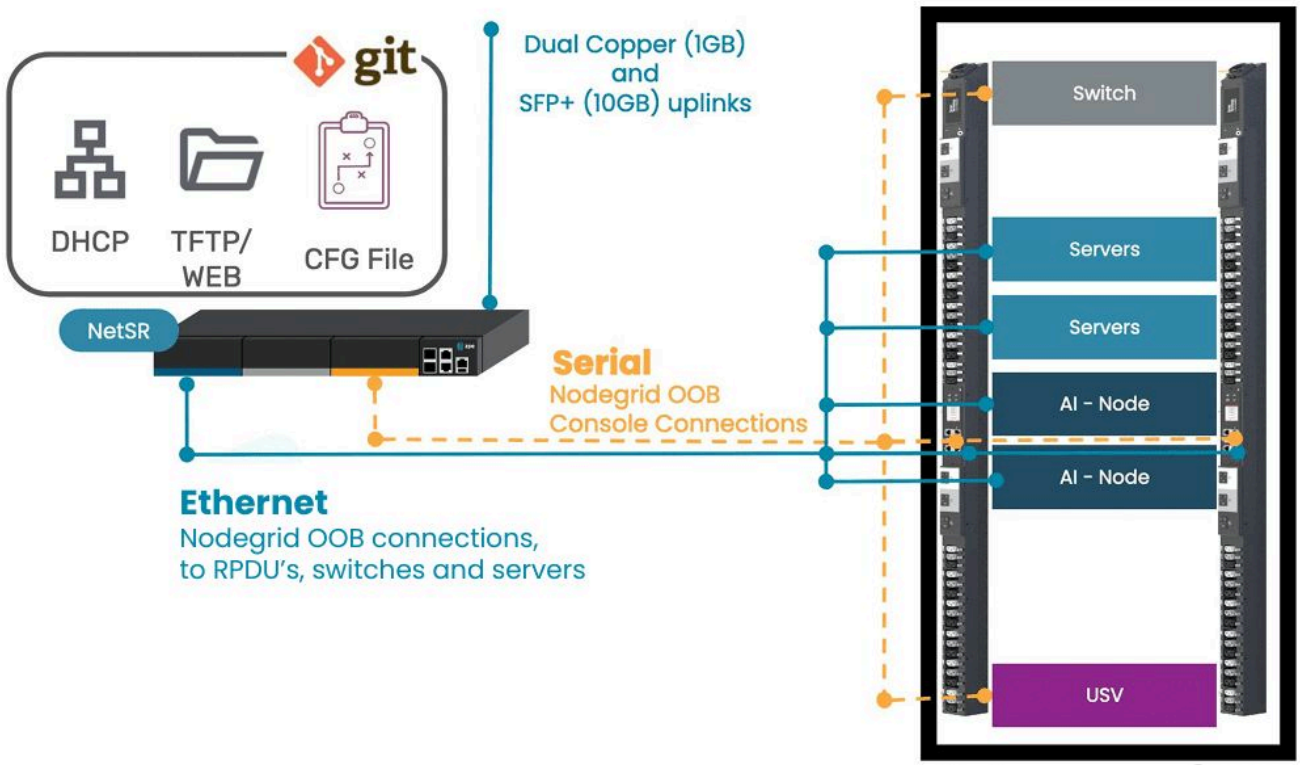

ZPE Systems' Nodegrid is Ideal for NVIDIA's SuperPOD Architecture

Because Nodegrid combines various interface types (serial, Ethernet/fiber, and USB) with up to 96 ports in a 1U appliance, it's the ideal out-of-band management solution for complex, high-density AI environments.

As an example, Nodegrid connects to and provides out-of-band for the following infrastructure listed in NVIDIA's SuperPOD reference design:

- Supermicro HGX H100 GPU SuperServer (via IPMI)
- L2/L3 1/10/25G Ethernet switch (via RS-232)
- NVIDIA QM9700 1U NDR 400Gbps InfiniBand switch (via RS-232)
- NVIDIA Spectrum SN2201 TOR Ethernet switch (via RS-232)
- NVIDIA Spectrum SN4600 Ethernet switch (via RS-232)
- NVIDIA UFM Enterprise Appliance Gen 3.0 (via RJ-45)
- Certified storage: DDN AI400X, Dell PowerScale, IBM Storage Scale, NetApp BeeGFS (E-Series), WEKA, and VAST (via GbE or host CLI)
- Raritan PX3-5008I2R-Q1 Power Distribution Unit (via RJ-45 or USB) and thermal monitoring/alarming passed from required temperature sensors

Sources: NVIDIA DGX SuperPOD Reference Architecture, NVIDIA DGX SuperPOD Reference Guide

Contact our Experts to Discuss Out-of-Band for AI

Organizations implementing OOB management see significant gains in operational efficiency, system reliability, and security. *"Nodegrid paid for itself after just one outage"* is common feedback from ZPE Systems' customers, and our experts are ready to help you implement the leading out-of-band solution. Get in touch to schedule a hands-on demo and discuss how to get the most out of your environment.

Discuss Out-of-Band for AI

Additional Resources for Out-of-Band Management

- Nvidia's Reference Architecture for Out-of-Band Management
- The CrowdStrike Outage: How to Recover Fast and Avoid the Next Outage
- Living Spaces case study: Scaling to 50 sites with only 3 network staff
- Video: How to Make Networks Hard to Fail and Easy to Recover
- ZPE Systems' Supply Chain Security Assurance (pdf)