Solution Guide

zpe

Best Practices Guide for Deploying Nvidia DGX and Other AI Pods

Security and resilience lessons learned from customers who operate AI Pods

A brand of legrand

AI Infrastructure Lacks Best Practices

AI is causing a strategic shift for modern enterprises, enabling financial modeling, anomaly detection, and language model processing.

But as organizations invest hundreds of millions in AI infrastructure, leadership must ask: What's the most efficient way to operate AI equipment and secure our multi-million-dollar investment?

AI infrastructure poses three major challenges:

1. **Delays**
   Deploying AI equipment is time-consuming and prone to error

2. **Revenue Loss**
   Recovering from failures takes too long and can cost millions of dollars

3. **Increased Risks**
   Manually operating equipment increases the risk of error-induced outages and security breaches

As Legrand (ZPE Systems' parent company) supports AI data center build outs, we see hyperscalers refining their best practices to address these challenges. These experienced AI hyperscalers are now adopting **Isolated Management Infrastructure (IMI).** IMI creates a dedicated system for automating operations and recovering from outages, while isolating infrastructure from the vulnerabilities of traditional management approaches.

ZPE Systems has developed IMI over the past 10 years by working with hyperscalers and global enterprises who have stringent requirements for uptime and security. Although Nvidia's DGX SuperPOD reference design is comprehensive, it omits IMI.

This guide shows the best practices for enabling automated deployments, rapid recovery, and improved security of Nvidia DGX and other AI Pods, by using an IMI approach.

The Problem: Deploying and Recovering Multi-Node AI Pod Infrastructure

Too Time Consuming: Once a Pod is physically installed and the cables are connected, engineers must manually set up each component. Errors or a lack of staff can cause delays and missed deadlines. It often takes weeks or months to configure AI equipment and bring the Pod online.

Too Far Away: Traditional management tools are fragmented solutions that require some manual, on-site intervention. To manage distributed infrastructure, teams find themselves in operational firefighting mode – fixing config errors, troubleshooting hanging servers, and reconfiguring networks.

Image: Traditional management tools require on-site intervention by engineering staff.

Too Many Parts: Managing AI Pods at scale is extremely complex due to the number of components required. Here are some components that teams must operate as part of one Nvidia DGX SuperPOD:

- Compute/GPU devices (HGX H100 servers)
- Storage appliances (Dell PowerScale)
- Storage fabric (InfiniBand high-speed)
- Network fabric (InfiniBand high-speed)
- Management fabric (via jumpbox entry point)
- Power and cooling systems (Raritan PX3 PDU)
- Environmental/device health monitoring (temperature, humidity, fan speed)

Adding to this complexity, the SuperPOD requires teams to build four distinct networks:

1. **High-speed InfiniBand network:** Connects GPU to GPU directly
2. **Storage InfiniBand network:** Connects compute node to storage array
3. **In-band Ethernet network:** Connects compute nodes to all other compute nodes, and to customer corporate network, Internet, etc.
4. **Out-of-band Ethernet network:** Used for management/IPMI, telemetry/observability, automation, etc.

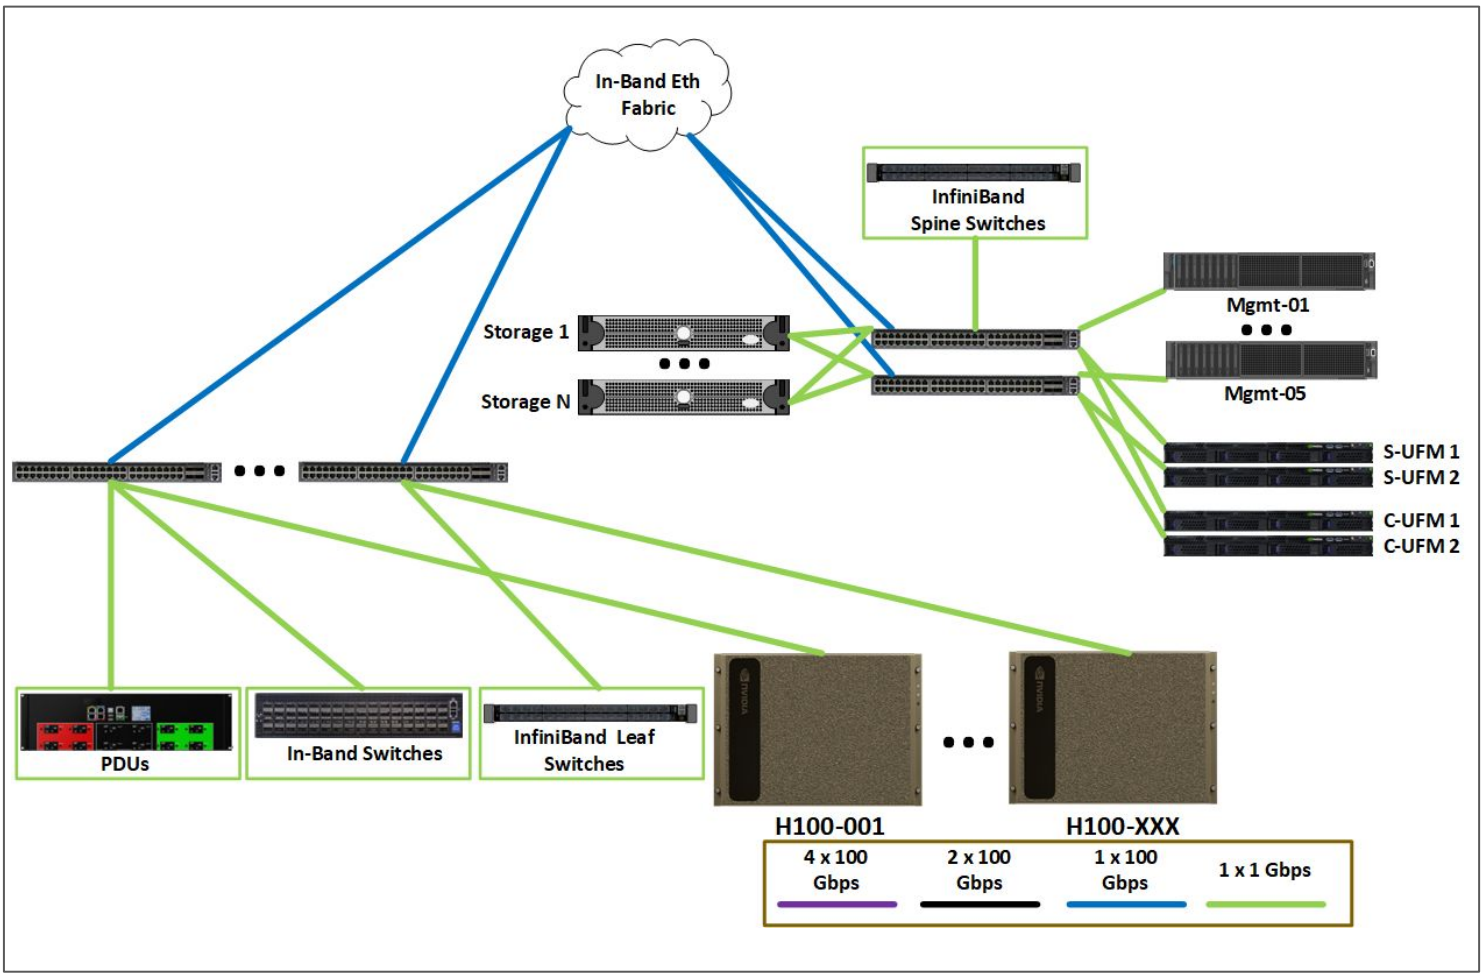

Diagram: Nvidia DGX SuperPOD management topology.

Solving the overarching problem requires addressing three areas of the infrastructure lifecycle:

1. **How can AI Pods be deployed without manual network configuration?**
2. **How can AI infrastructure be recovered if the production network fails?**
3. **How can teams reduce the risk of causing outages and security incidents?**

The Solution: Isolated Management Infrastructure

1. How Can AI Pods Be Deployed Without Manual Network Configuration?

Problem:

Setting up AI pods requires configuring networking components, such as VLANs, IP addresses, and PXE boot servers. This is often done manually, leading to deployment delays, errors, and inconsistencies that increase costs and operational overhead.

Gap:

Traditional infrastructure management solutions lack automated setup capabilities, also known as zero touch provisioning (ZTP). This forces teams to rely on manual, error-prone processes.

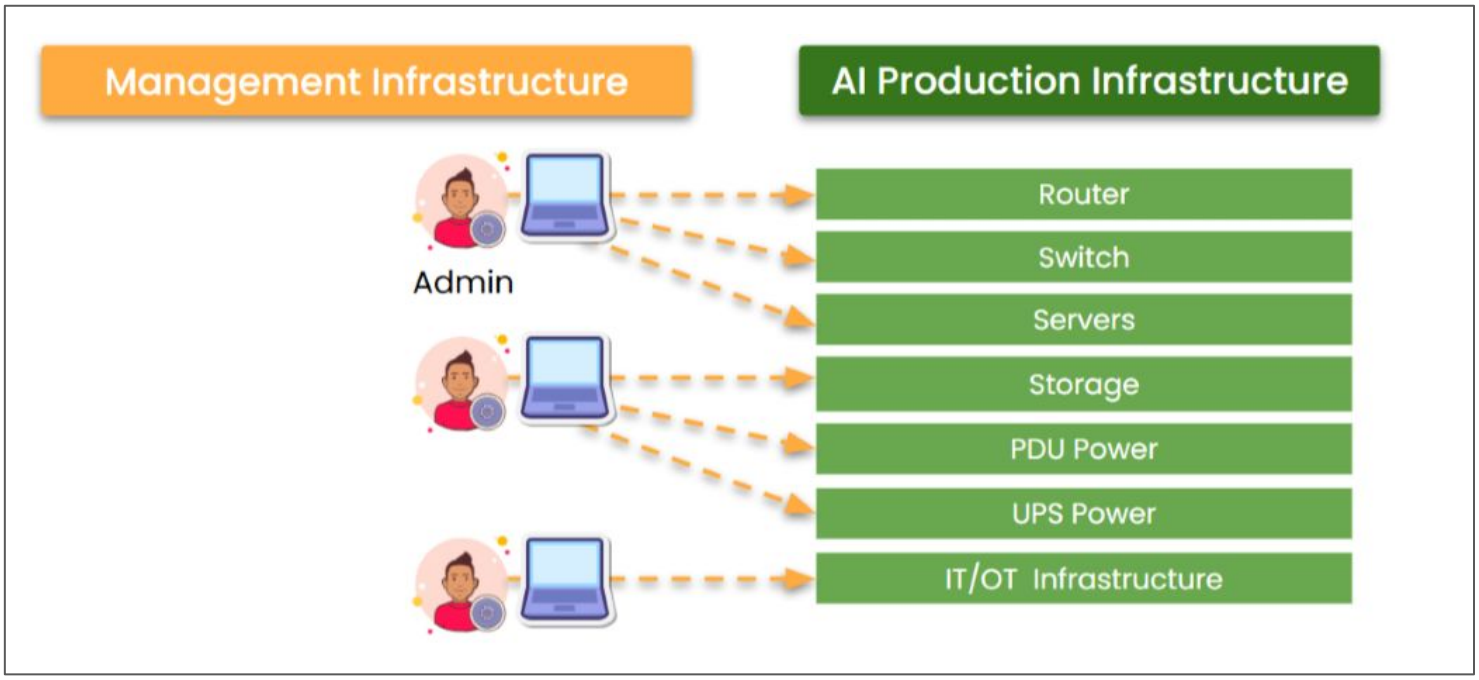

Diagram: Admins must manually configure individual components of the AI Pod.

Best Practice:

Isolated Management Infrastructure provides the dedicated environment where teams can build, test, and deploy automation and orchestration. This allows them to automatically configure the networking components necessary for setting up the rest of the AI Pod. However, this requires additional servers for running VMs/Docker and automation.

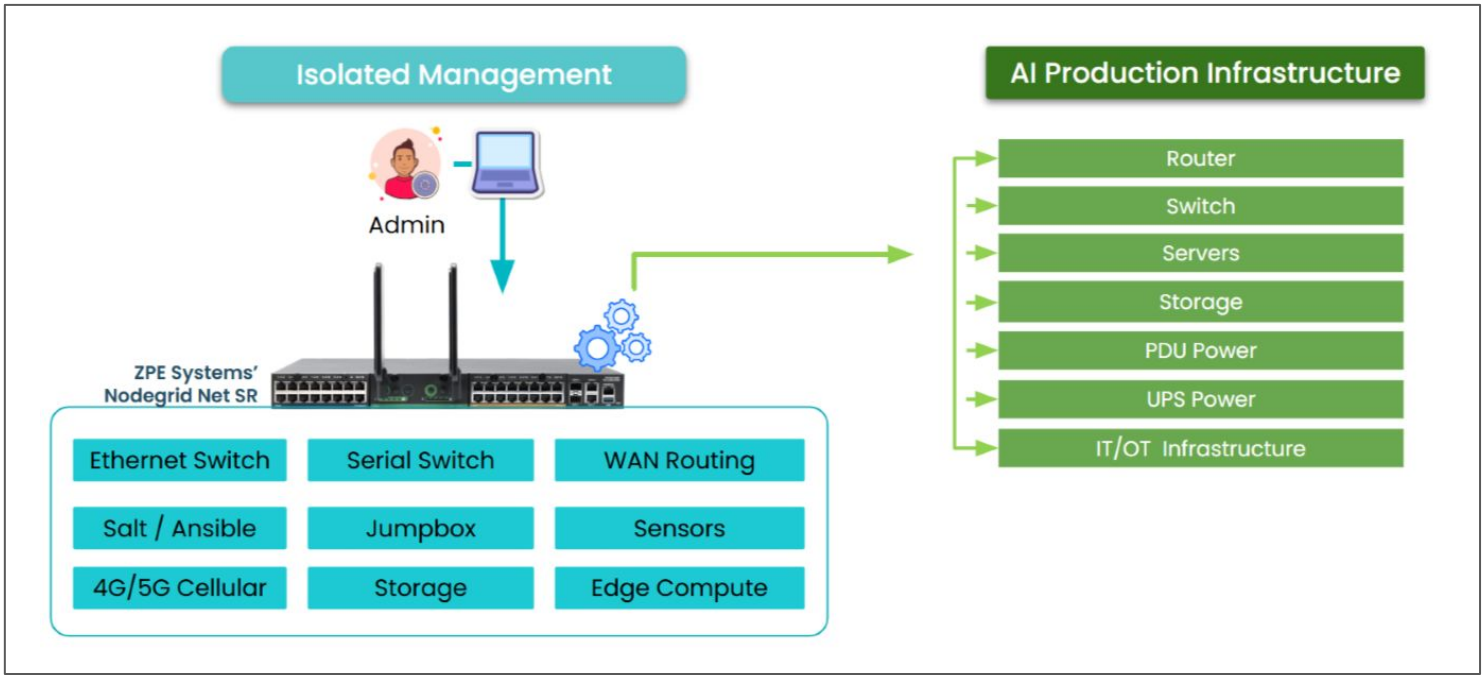

Diagram: IMI is a dedicated environment for running automation to set up the complete DGX Pod.

How ZPE Improves Best Practices:

- **Complete IMI in One Device:** Consolidates 9+ management functions into one device, serving as a jumpbox entry point as well as a VM and automation server.
- **Zero-Touch Provisioning (ZTP):** Automates network deployment as well as the rest of the stack, including compute, storage, IoT, power, and sensor devices.
- **Faster Time-to-Market:** Reduces deployment timelines and minimizes human intervention, lowering the risk of configuration errors.
- **Automated Environmental Monitoring:** Tracks airflow, heat, and power usage for efficiency in ongoing operations.
- **Integrated Security & Access Control:** Enables monitoring of access control for enhanced security.
- **Reliable Connectivity for Administrators:** Ensures remote access and control via backup links including 5G, Starlink, broadband, and other connection types.

2. How Can AI Infrastructure Be Recovered If the Production Network Fails?

Problem:

If the primary network fails, restoring AI clusters becomes impossible without an independent recovery path, leading to extended downtime, lost productivity, and revenue impact.

Gap:

The default SuperPOD design focuses on out-of-band (OOB) via Ethernet. This ignores the fact that many AI Pod components — like DGX nodes, networking gear, storage, and power systems — require serial (RS-232), Ethernet, USB, and virtualized interfaces for full management access.

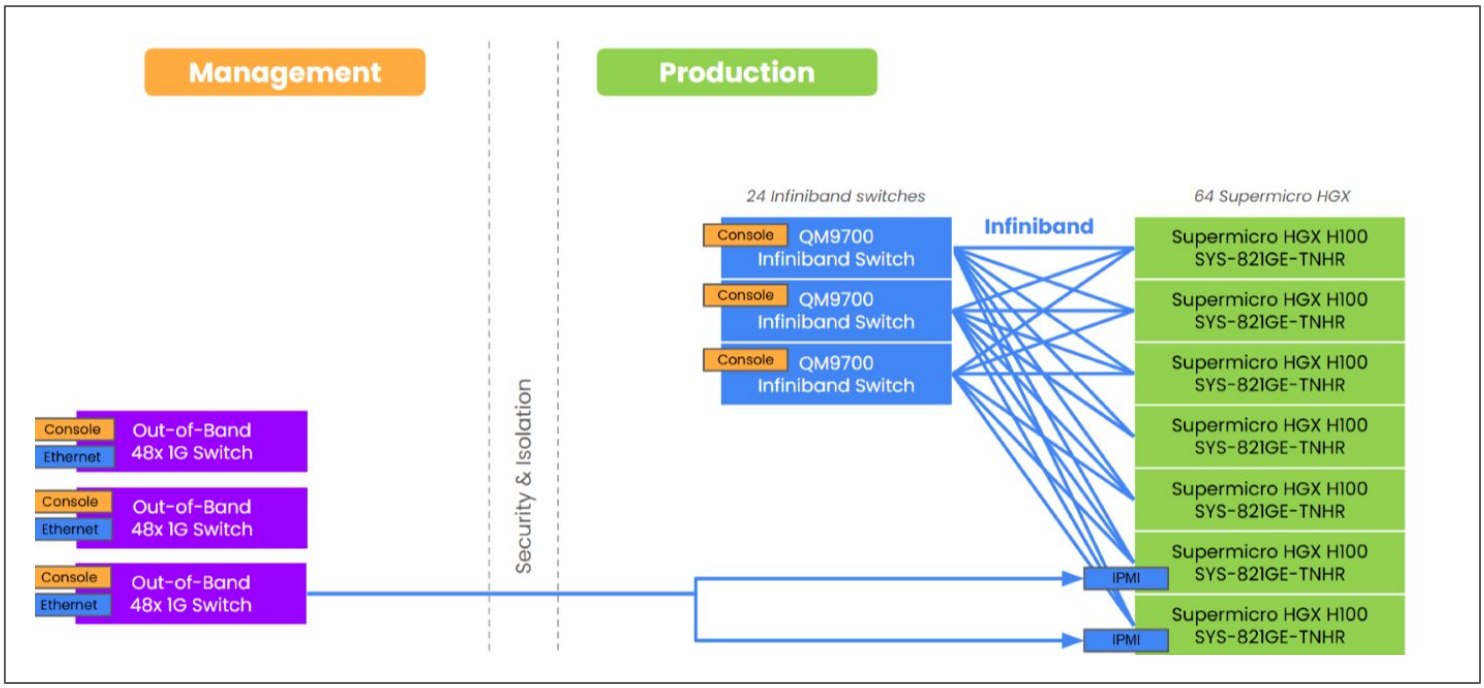

Diagram: The default SuperPOD design limits out-of-band access to Ethernet interfaces only.

Best Practice:

As part of IMI, proper OOB provides admin access to all production equipment, but does not rely on any of this production equipment. However, implementing a comprehensive OOB solution requires dedicated devices. This adds to the overall management stack, and because AI environments are so packed full, there's no rack space or power to accommodate this stack.

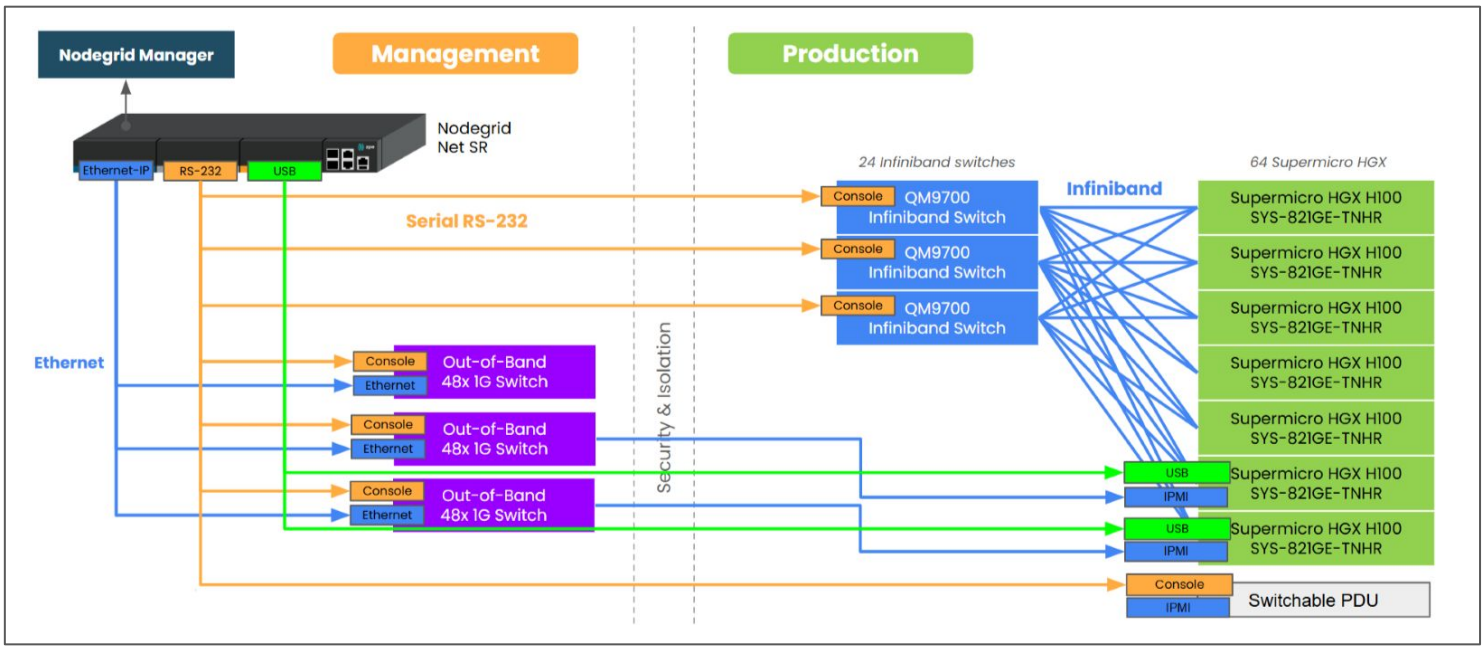

Diagram: IMI via ZPE Systems' Nodegrid provides out-of-band remote access to all components in the DGX Pod via Ethernet, serial, USB, and virtualized interfaces.

How ZPE Improves Best Practices:

- **Comprehensive OOB**: One ZPE device provides OOB for the entire DGX Pod via RS-232 serial, USB, Ethernet, and virtualized interfaces (REST API, IPMI, Redfish).
- **Network Isolation & Failover:** Provides a fully isolated management network with LTE/5G, fiber, or Wi-Fi failover, ensuring access even if the primary network is down.
- **Granular Remote Control:** Enables full remote BIOS/UEFI access, power cycling, and firmware updates via console servers, reducing the need for on-site interventions.
- **Ransomware Recovery:** Allows teams to deploy an Isolated Recovery Environment where infrastructure can be isolated, cleansed, and restored even during active attacks.
- **Automated Rollback:** Allows teams to recover the entire Pod quickly using a golden image and zero touch provisioning.

3. How Can Teams Reduce the Risk of Causing Outages and Security Incidents?

Problem:

Normal system maintenance and administration can lead to error-induced outages because traditional management gives direct access to production infrastructure. Traditional tools grant broad access privileges that also increase security risks, such as through stolen credentials.

Gap:

Directly managing infrastructure leaves no room for error when making changes or deploying automations. This approach is banned by CISA. Traditional tools also lack proper segmentation and Role-Based Access Control (RBAC), giving users access to environments that are outside of their expertise.

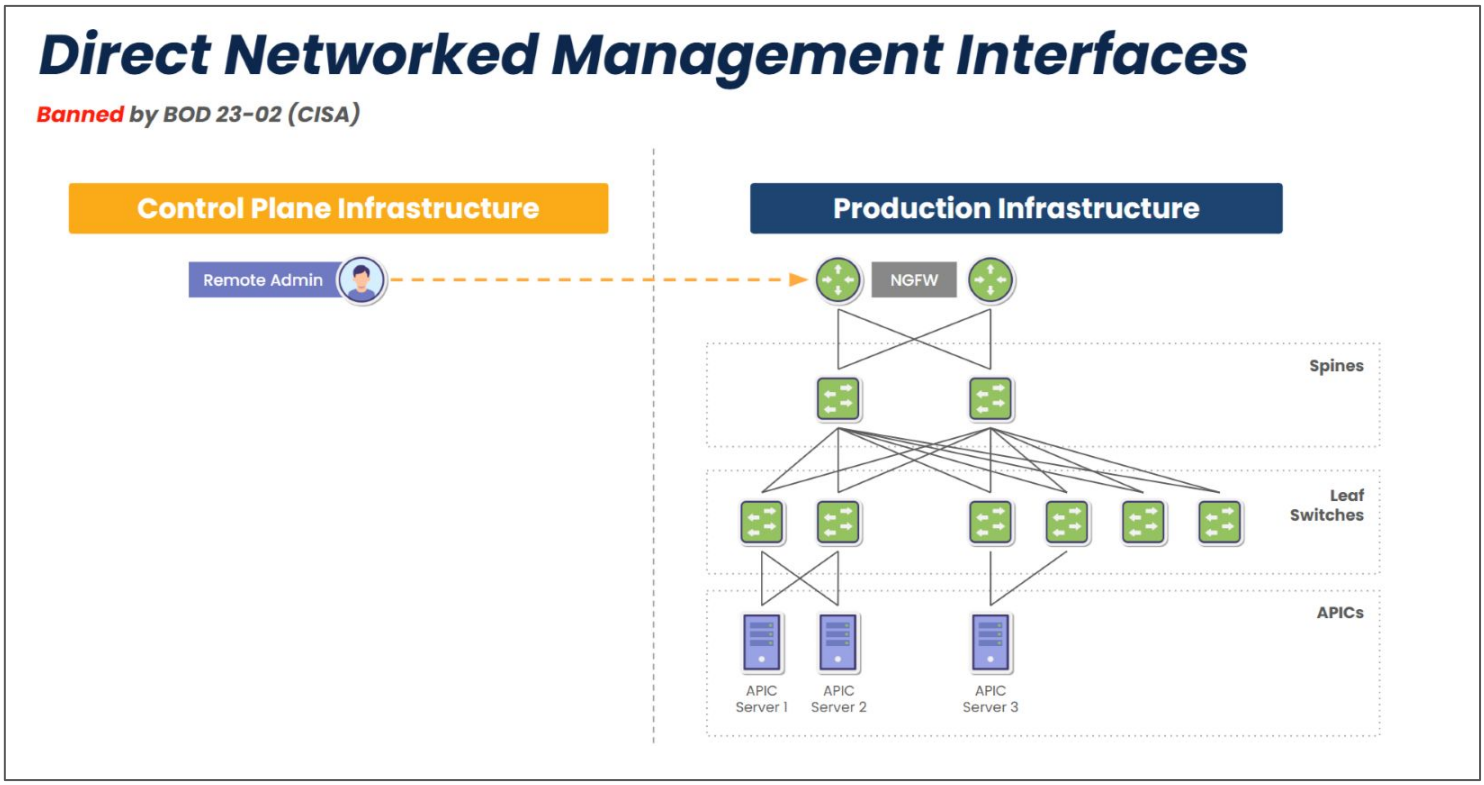

Diagram: Traditional network management relies on direct access to production infrastructure.

Best Practice:

Per CISA best practice, IMI separates management from production traffic but also provides a safe environment for admins to test their changes and automations. IMI also enables separation of duties so that every change/automation must go through multiple levels of approval before being pushed to production. RBAC enforces access, however some solutions may lack authentication controls or integration with SSO providers.

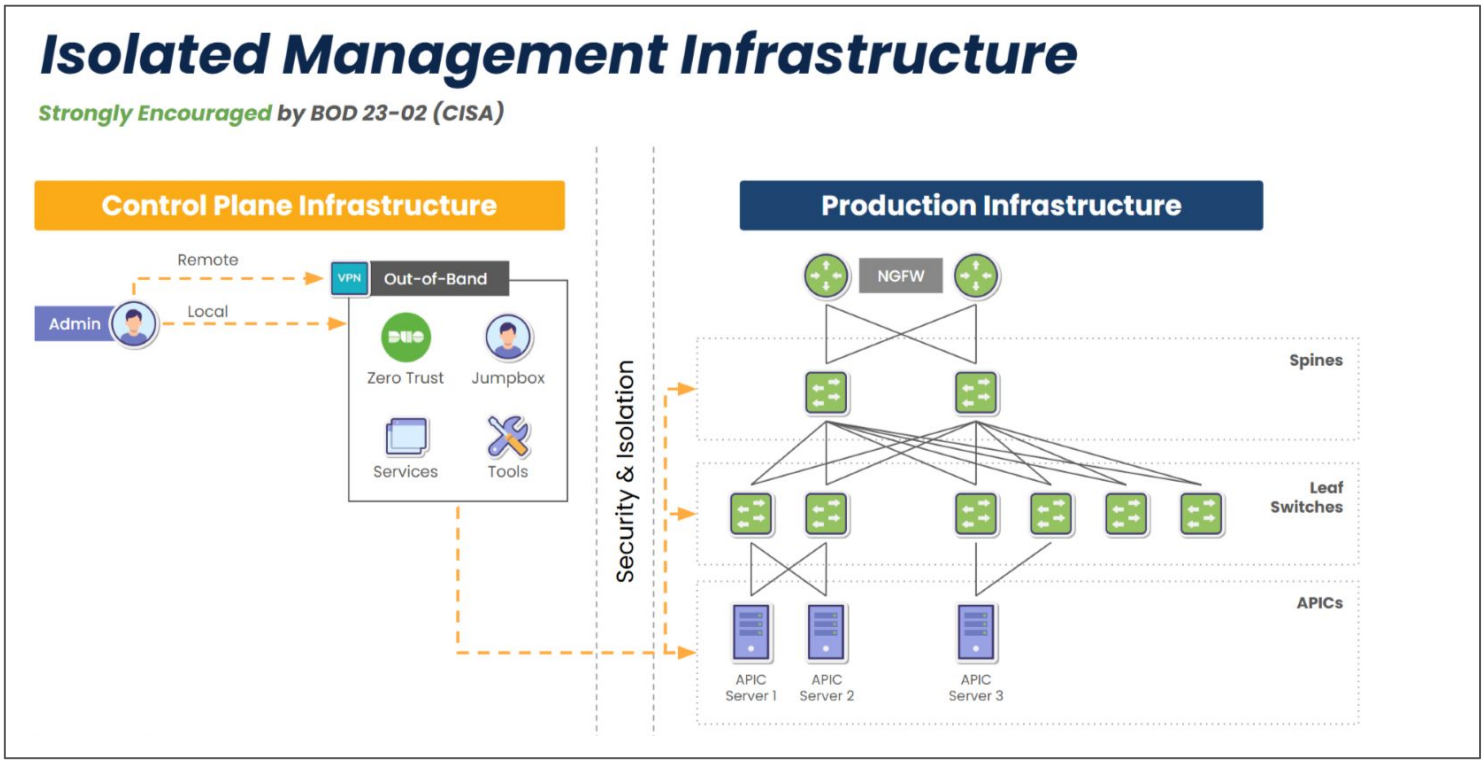

Diagram: IMI separates management so teams can test changes and automations before implementing.

How ZPE Improves Best Practices:

- **Full Isolation:** Fully isolates management from production traffic to ensure changes and automations can be executed without impacting production systems.
- **Zero-Trust Security:** RBAC and policy enforcement allow for proper segmentation and separation of duties, aligning with zero-trust security principles.
- **Granular RBAC with MFA & SSO**: Enforces least-privilege access, with built-in controls for multi-factor authentication (MFA) and integration with major SSO providers.
- **Encrypted Remote Access**: Guarantees out-of-band connectivity via encrypted tunnels to prevent unauthorized entry.
- **Centralized Visibility and Control**: Provides a unified management platform to track user access, enforce policies, and ensure AI operations follow strict security guidelines.

Blueprint for Isolated Management Infrastructure

To deploy AI per hyperscaler best practices, organizations should build IMI as follows:

1. **Connect everything out-of-band:**
   Connect IP Ethernet switches, serial consoles, servers, PDUs, routing, and the entire networking stack to the out-of-band network.

2. **Build a dedicated automation framework:**
   Ensure automation does not go directly to production equipment and instead must pass through the management network.

3. **Follow zero trust security principles:**
   Use RBAC and other controls to segment management access, separate duties, and minimize risks.

Contact ZPE Systems for Drop-In IMI

ZPE Systems developed these best practices working with hyperscalers. The **modular Nodegrid Net SR** is ideal for the confined racks of AI environments, and fully supports the out-of-band, automation, and security demands of IMI.

ZPE also offers **professional services for installation and configuration,** ensuring seamless deployment and integration into existing IT workflows.

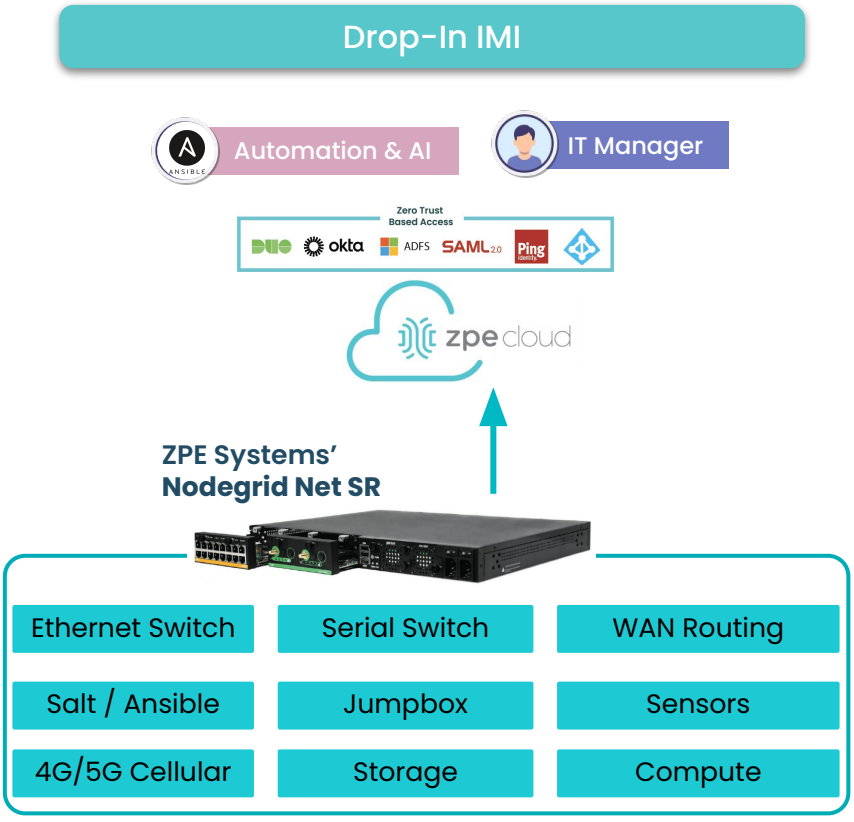

Get in touch for a demo!

sales@zpesystems.com
844-4ZPE-SYS