VENDOR-AGNOSTIC AI ETHERNET FABRIC TEST DRIVE

Service Overview:

Netnology specializes in the design, deployment and enablement of high-performance Ethernet fabrics for AI. The Test Drive helps customers build a baseline understanding of 400/800GbE AI networking and of the leading platform options (Cisco Nexus, NVIDIA Spectrum-X, HPE Juniper Networking and SONiC) for their AI/ML infrastructure.

This 2–4-hour session demonstrates how an AI fabric is provisioned, how it delivers lossless, high-performance transport for GPU-to-GPU (RoCEv2) traffic, and how it is monitored and operated. The demonstration runs on the platform of the customer's choice, and each platform's approach to automation, congestion control and load balancing is explained side by side, so the customer can judge fit as well as features.

This Test Drive is built and delivered by Netnology subject matter experts (SMEs) with hands-on, multi-vendor experience in data center and AI fabric deployments.

Solution Overview:

AI training and inference clusters place demands on the data center network that general-purpose fabrics were not designed for. GPU-to-GPU communication relies on RDMA over Converged Ethernet (RoCEv2), which generates synchronized, high-bandwidth flows that are highly sensitive to packet loss, latency and congestion. A single congested link or dropped packet can stall collective operations across the entire GPU cluster, extending job completion time and reducing the return on the GPU investment.

A modern AI fabric is built on 51.2 Tbps-class switches with 400GbE and 800GbE ports, deployed in a non-blocking, rail-optimized spine-leaf design with separate back-end (GPU-to-GPU), front-end and storage networks. Whichever platform is chosen, the fabric must deliver the same outcomes: lossless transport for RoCEv2 using Priority Flow Control (PFC) and Explicit Congestion Notification (ECN), load balancing that keeps large RDMA flows from colliding, and telemetry that exposes congestion before it affects job completion time. Netnology delivers these outcomes on three platform tracks:

- Cisco Nexus: Nexus 9000 Series 800GbE switches (such as the Nexus 9364E-SG2 with Cisco Silicon One G200) on NX-OS, with dynamic load balancing, PFC/ECN and DCQCN congestion control, and Cisco Nexus Dashboard for automation and visibility.
- NVIDIA Spectrum-X: Spectrum-4 Ethernet switches (64 × 800GbE) on Cumulus Linux, paired with BlueField-3 or ConnectX-8 SuperNICs for adaptive routing and Spectrum-X congestion control, with NVIDIA Air digital twin and NetQ visibility.
- HPE Juniper Networking: QFX5240 800GbE switches and PTX Series spines on Junos OS Evolved, with dynamic and RDMA-aware load balancing, DCQCN congestion control and Juniper Apstra intent-based automation.
- SONiC: Open network operating system on NVIDIA Spectrum and other 800GbE switch platforms, with FRR BGP, PFC/ECN buffer and QoS profiles, PFC watchdog, ZTP/Ansible automation and gNMI streaming telemetry.
- Open Host Connectivity: GPU servers from leading OEMs with RoCEv2 NICs such as NVIDIA ConnectX/BlueField, Broadcom Thor and AMD Pensando Pollara.
- Common Validation: The same RDMA (perftest) and NCCL/RCCL collective tests on every platform, so results are directly comparable.

Service Benefits:

Netnology has a team of world-class engineers who work across Cisco, NVIDIA, HPE Juniper and SONiC AI fabrics and are passionate about customer success. Netnology will partner with you to provide guidance on:

- Platform Selection: A vendor-neutral view of how each platform fits your GPU cluster size, operating model and team skills.
- AI Fabric Architecture: Back-end, front-end and storage network design options at 400GbE and 800GbE.
- Seamless Integration: Integrating the AI fabric with your existing data center network and operations tooling.
- Optimized Performance: Tuning lossless QoS, congestion control and load balancing for AI traffic.

Scope of Services:

As part of the 2–4-hour Test Drive engagement, Netnology will provide the following services on one selected platform track:

- Solution Overview
- Platform: Cisco Nexus, NVIDIA Spectrum-X, HPE Juniper Networking and SONiC
- Solution Use-case Demo
  - Fabric Provisioning and Automation
    - Fabric build with the platform toolset (Cisco Nexus Dashboard, NVIDIA NVUE and Air, Juniper Apstra, or SONiC with ZTP/Ansible)
    - BGP underlay and, where required, EVPN-VXLAN overlay
    - AI/ML QoS policy deployment
  - Lossless Ethernet for RoCEv2
    - PFC and ECN in action
    - Load balancing: ECMP, dynamic load balancing and adaptive routing
  - GPU Server Fabric Integration
    - Server NIC RoCEv2 settings (DSCP/PFC priority, ECN, MTU)
    - Host-to-fabric connectivity and LLDP verification
  - AI Traffic Validation and Visibility
    - RDMA and NCCL collective test traffic across the fabric
    - Congestion and telemetry visibility (Cisco Nexus Dashboard, NVIDIA NetQ, Juniper Apstra or gNMI streaming telemetry)
    - Link failure and recovery
- Lessons learned and enhancements

Target Audience:

This service is designed for IT leaders, architects and network engineers evaluating Ethernet platforms for AI infrastructure. No prior experience with a specific AI fabric platform is required.

2026 Netnology and Affiliates | Confidential

Service Deliverables:

No.	Deliverable	Service Details
1.	Concepts	• Solution Overview • Platform Comparison • Use-Case Overview
2.	Value-Prop	• Ease of configuration and automation • Ease of manageability • Performance and visibility for AI traffic
3.	Execution	• Fabric provisioning on the selected platform • Lossless Ethernet (RoCEv2) • AI traffic validation and visibility
4.	Wrap-up	• Lessons Learned • Platform fit and recommended next steps • Q&A

2026 Netnology and Affiliates | Confidential