VENDOR-AGNOSTIC AI ETHERNET FABRIC MENTORED INSTALL

Service Overview:

Netnology specializes in the design, deployment and enablement of high-performance, lossless Ethernet fabrics that connect GPU clusters, across Cisco Nexus, NVIDIA Spectrum-X, HPE Juniper Networking and SONiC. As part of this Mentored Install service, our subject matter experts (SMEs) work side-by-side with your team to deploy a 400/800GbE AI fabric on your selected platform, validate its performance under AI traffic patterns, and provide knowledge transfer.

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 the service on one of four platform tracks: Cisco Nexus (Nexus 9000 Series 800GbE switches on NX-OS with dynamic load balancing, DCQCN and Cisco Nexus Dashboard), NVIDIA Spectrum-X (Spectrum-4 switches on Cumulus Linux with BlueField-3/ConnectX-8 SuperNICs, adaptive routing, NVIDIA Air and NetQ), HPE Juniper Networking (QFX5240 and PTX Series on Junos OS Evolved with dynamic load balancing, DCQCN and Juniper Apstra), or SONiC (on NVIDIA Spectrum or other 800GbE switch platforms, automated with ZTP/Ansible and monitored with gNMI streaming telemetry).

Service Benefits:

Netnology has a team of world-class engineers who specialize in high-performance AI data center fabrics across multiple vendors and are passionate about customer success. This Mentored Install engagement delivers:

- Faster time to production: a validated, production-ready AI fabric built to the selected vendor's reference architecture and best practices.
- Lossless, predictable performance: RoCEv2, PFC and ECN tuned and verified end-to-end, from switch to GPU server NIC.
- Proven results: fabric performance validated with RDMA and NCCL collective tests before workloads go live.
- Operational readiness: hands-on mentoring so your team can operate, monitor and troubleshoot the fabric.

- Platform choice: the same delivery method and validation standard whichever platform you select.

Service Scope:

As part of the 10-day (up to 80 hours) engagement, Netnology will provide the following services on the selected platform track:

- AI Ethernet Fabric (Cisco Nexus, NVIDIA Spectrum-X, HPE Juniper Networking or SONiC)
  - Fabric Design
    - Non-blocking spine-leaf topology for back-end (GPU), front-end and storage networks
    - Rail-optimized GPU server connectivity, 400/800GbE port and breakout plan, optics and oversubscription planning
    - IP addressing, BGP routing and lossless QoS design
  - Fabric Deployment and Automation (up to 10 switches)
    - Onboard switches and build the fabric with the platform toolset (Cisco Nexus Dashboard, NVIDIA NVUE with Air digital twin, Juniper Apstra blueprints, or SONiC with ZTP/Ansible)
    - Configure BGP underlay and MTU (jumbo frames)
    - Configure EVPN-VXLAN overlay (VRFs, VNIs) for front-end and storage networks, as required
  - Lossless Ethernet for RoCEv2
    - Configure QoS classification and queuing for RDMA and CNP traffic
    - Configure Priority Flow Control (PFC) and PFC watchdog
    - Configure ECN thresholds and congestion control (DCQCN, or Spectrum-X congestion control on NVIDIA)
    - Configure load balancing across fabric links (ECMP, dynamic load balancing on Cisco, dynamic or RDMA-aware load balancing on Juniper, adaptive routing on Spectrum-X)
  - GPU Server Fabric Integration (up to 6 nodes)
    - Configure server NICs/SuperNICs for RoCEv2 (DSCP/PFC priority, ECN, MTU) using the NIC vendor's host tools
    - Verify host-to-fabric connectivity, LLDP and link health
  - Telemetry and Visibility
    - Enable fabric telemetry (Cisco Nexus Dashboard, NVIDIA NetQ, Juniper Apstra, or gNMI/OpenConfig streaming telemetry to the customer's collector)
    - Configure views and alerts for PFC, ECN, buffer congestion and link errors
  - Fabric Test and Validation
    - Connectivity, redundancy and failover testing
    - RDMA performance testing (bandwidth and latency)
    - NCCL/RCCL collective tests (e.g., all-reduce, all-to-all) to validate fabric behavior under AI traffic
  - Knowledge Transfer

Target Audience:

This service is designed for Network Architects, Network Engineers and Data Center Administrators responsible for designing, deploying and operating the AI fabric.

Prerequisites:

- Working knowledge of data center networking (BGP, EVPN-VXLAN, QoS) and of the selected platform's operating system (Cisco NX-OS, Cumulus Linux/NVUE, Junos OS Evolved or SONiC).
- All switches, GPU servers, optics and cables are racked, cabled to the agreed cabling plan and powered up prior to the kick-off.
- Required software licenses and subscriptions for the selected platform (for example Cisco NX-OS and Nexus Dashboard, Cumulus Linux and NetQ, Junos OS Evolved and Juniper Apstra, or the chosen SONiC distribution) are procured and available.
- Fabric management and telemetry tooling (Cisco Nexus Dashboard, NVIDIA NetQ, Juniper Apstra or a gNMI telemetry collector) is deployed and reachable, or compute resources are available for its deployment.
- GPU servers are installed with a supported operating system, GPU and NIC drivers/firmware, and RDMA/NCCL test tools to support fabric validation.
- Remote or on-site access to the environment is provided to Netnology engineers.

Out of Scope:

- AI/ML model deployment, LLM tuning, data pipelines and application development.
- Operating system installation and GPU software stack build-out on the servers.
- Storage system configuration beyond its connectivity to the fabric.
- InfiniBand fabric design and deployment.

Service Deliverables:

No	Deliverable	Service Details
1.	Project Kickoff	• Project overview and engagement plan • Solution and platform overview • Confirm GPU cluster size, AI workload profile and success criteria • Gather customer requirements
2.	Pre-Requisite Validation	• Review/confirm hardware, optics and cabling readiness (switches and GPU servers) • Switch OS software and license validation • Management and telemetry tooling readiness • Management network and access readiness
3.	Fabric Design	• High-level fabric design • Port, breakout, optics and IP addressing plan
4.	Fabric Deployment	• Onboard switches (up to 10) • Configure BGP underlay and MTU • Configure EVPN-VXLAN overlay for front-end and storage networks, as required

No	Deliverable	Service Details
5.	Lossless Ethernet (RoCEv2) Configuration	• QoS classification and queuing • PFC and PFC watchdog • ECN and congestion control • Load balancing across fabric links
6.	GPU Server Fabric Integration (up to 6 nodes)	• Server NIC RoCEv2 settings • Host-to-fabric connectivity verification
7.	Telemetry and Visibility	• Enable fabric telemetry • Congestion views and alerts (PFC, ECN, buffers)
8.	Test and Validation	• Test plan • Connectivity, redundancy and failover tests • RDMA performance tests • NCCL collective tests • Test results report
9.	As-Built Documentation	• As-built configuration document for the deployed fabric
10.	Knowledge Transfer	• Explain how to configure, operate and troubleshoot the deployed fabric in the customer's environment • Day-2 operations best practices