CISCO AI-OPTIMIZED FABRIC MENTORED INSTALL

Service Overview:

Netnology specializes in the design, deployment and enablement of Cisco AI-Optimized Fabrics, the high-performance, lossless Ethernet networks that connect GPU clusters. As part of this Mentored Install service, our subject matter experts (SMEs) work side-by-side with your team to deploy the Cisco AI-Optimized Fabric, validate its performance under AI traffic patterns, and provide the knowledge transfer your staff need to operate it with confidence.

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.

The Cisco AI-Optimized Fabric is built on Cisco Nexus 9000 Series switches at 400G and 800G, deployed in a non-blocking spine-leaf design with separate back-end (GPU-to-GPU), front-end and storage networks. It delivers lossless Ethernet transport using Priority Flow Control (PFC) and Explicit Congestion Notification (ECN), load balancing that spreads RDMA traffic evenly across all available paths, and granular telemetry for congestion visibility. Cisco Nexus Dashboard provides automated fabric provisioning with built-in AI/ML QoS and queuing policies, together with day-2 visibility and analytics, while Cisco UCS GPU servers managed through Cisco Intersight connect the compute nodes into the fabric.

Service Benefits:

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

- Faster time to production: a validated, production-ready AI fabric built to Cisco 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.

Service Scope:

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

- Cisco AI-Optimized Fabric
  - Fabric Design

  - Non-blocking spine-leaf topology for back-end (GPU), front-end and storage networks
  - Rail-optimized GPU server connectivity and oversubscription planning
  - IP addressing, routing and lossless QoS design

- Fabric Deployment with Cisco Nexus Dashboard (up to 10 switches)
  - Onboard Cisco Nexus 9000 switches and create the AI fabric
  - Configure underlay routing and MTU (jumbo frames)
  - Configure VXLAN EVPN overlay (VRFs, VLANs/VTEPs) for front-end and storage networks, as required

- Lossless Ethernet for RoCEv2
  - Configure QoS classification and queuing for RDMA and control traffic
  - Configure Priority Flow Control (PFC) and PFC watchdog
  - Configure ECN/WRED thresholds for congestion management
  - Configure load balancing across fabric links (ECMP / dynamic load balancing)

- GPU Server Fabric Integration (up to 6 nodes)
  - Onboard GPU servers to Cisco Intersight and apply server policies and profiles (BIOS, boot, LAN connectivity)
  - Configure server NICs for RoCEv2 (DSCP/PFC priority, ECN, MTU)
  - Verify host-to-fabric connectivity, LLDP and link health

- Telemetry and Visibility
  - Enable fabric telemetry in Cisco Nexus Dashboard
  - Configure views and alerts for PFC, ECN and buffer congestion
  - Deploy the Splunk OpenTelemetry Collector on the GPU servers and send host and GPU (NVIDIA DCGM) metrics to Splunk Observability Cloud, where licensed.
  - Build a dashboard and detectors for GPU utilization, temperature and host health, viewable alongside the PFC, ECN and buffer congestion views.

- Fabric Test and Validation
  - Connectivity, redundancy and failover testing
  - RDMA performance testing (bandwidth and latency)
  - NCCL collective tests (e.g., all-reduce) 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 Cisco Nexus (NX-OS) and data center networking.
- 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 (Cisco Nexus, Cisco Nexus Dashboard, Cisco Intersight) are procured and available.
- Cisco Nexus Dashboard 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.
- The customer provides a Splunk Observability Cloud organization and ingest access token, and the GPU servers have outbound HTTPS access, a supported OS, NVIDIA drivers and the DCGM exporter installed.

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.

Service Deliverables:

No	Deliverable	Service Details
1.	Project Kickoff	• Project overview and engagement plan • Solution 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) • NX-OS software and license validation • Cisco Nexus Dashboard and Cisco Intersight readiness • Management network and access readiness
3.	Fabric Deployment	• Onboard Nexus 9000 switches (up to 10) • Configure underlay routing and MTU • Configure VXLAN EVPN overlay for front-end and storage networks, as required
4.	Lossless Ethernet (RoCEv2) Configuration	• QoS classification and queuing • PFC and PFC watchdog • ECN/WRED congestion management • Load balancing across fabric links
5.	GPU Server Fabric Integration (up to 6 nodes)	• Cisco Intersight onboarding, server policies and profiles • Server NIC RoCEv2 settings • Host-to-fabric connectivity verification
6.	Telemetry and Visibility	• Enable fabric telemetry in Cisco Nexus Dashboard • Congestion views and alerts (PFC, ECN, buffers) • GPU server host and GPU metrics, dashboard and detectors in Splunk Observability Cloud, where licensed
7.	Test and Validation	• Test plan • Connectivity, redundancy and failover tests • RDMA performance tests

		• NCCL collective tests • Test results report
8.	As-Built Documentation	• As-built configuration document for the deployed fabric
9.	Knowledge Transfer	• Explain how to configure, operate and troubleshoot the deployed fabric in the customer's environment • Day-2 operations best practices