NVIDIA INFINIBAND AI FABRIC MENTORED INSTALL

Service Overview:

Netnology specializes in the design, deployment and enablement of NVIDIA Quantum InfiniBand fabrics, the high-performance, natively lossless networks that connect large GPU clusters. As part of this Mentored Install service, our subject matter experts (SMEs) work side-by-side with your team to deploy an NDR (400 Gb/s) or XDR (800 Gb/s) InfiniBand fabric, validate its performance under AI traffic patterns, and provide knowledge transfer.

Solution Overview:

Large-scale AI training and inference clusters place demands on the network that general-purpose fabrics were not designed for. GPU-to-GPU communication generates synchronized, high-bandwidth collective operations (all-reduce, all-gather, all-to-all) that are highly sensitive to latency, jitter and congestion. A single congested link or degraded cable can slow collectives across the entire GPU cluster, extending job completion time and reducing the return on the GPU investment.

NVIDIA Quantum InfiniBand is purpose-built for this traffic. It is a natively lossless, credit-based fabric, centrally managed by a subnet manager, available at 400 Gb/s (Quantum-2 NDR, 64 × 400 Gb/s switches) and 800 Gb/s (Quantum-X800 XDR, Q3200 and Q3400 switches with up to 144 × 800 Gb/s ports). Deployed as a non-blocking, rail-optimized fat tree, it combines adaptive routing and congestion control to keep large flows off congested paths, NVIDIA SHARP in-network computing to offload collective operations into the switches, and ConnectX-7 and ConnectX-8 adapters at the GPU servers. NVIDIA Unified Fabric Manager (UFM) provides fabric provisioning, subnet management, telemetry and day-2 monitoring.

Service Benefits:

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

- Faster time to production: a validated, production-ready InfiniBand fabric built to NVIDIA reference architecture best practices.
- Predictable performance: adaptive routing, congestion control and SHARP tuned and verified end-to-end, from switch to ConnectX adapter.
- 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 with UFM.

Service Scope:

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

- NVIDIA Quantum InfiniBand Fabric (NDR or XDR)
  - Fabric Design
    - Non-blocking fat-tree topology (two- or three-tier) sized for the GPU cluster
    - Rail-optimized GPU server connectivity, scalable-unit layout, optics and cabling plan
    - Subnet manager, routing engine, partition (PKey) and service level design
  - Fabric Deployment with NVIDIA UFM (up to 10 switches)
    - Deploy UFM (Enterprise software or appliance) with high availability
    - Validate switch firmware and software, and upgrade as required
    - Discover the fabric and validate topology and cabling with ibdiagnet
    - Configure the subnet manager, routing engine and partitions
  - Performance Features
    - Configure adaptive routing
    - Configure congestion control
    - Enable NVIDIA SHARP in-network computing (aggregation manager and host libraries)
  - GPU Server Fabric Integration (up to 6 nodes)
    - Configure ConnectX adapters (firmware, InfiniBand link type, DOCA-Host drivers)
    - Verify host-to-fabric connectivity, link speed/width and error counters
  - Telemetry and Visibility
    - Enable UFM telemetry and fabric health monitoring
    - Configure views and alerts for link errors, congestion and port health
  - Fabric Test and Validation
    - Connectivity, redundancy and failover testing (including subnet manager failover)
    - RDMA performance testing (bandwidth and latency)
    - NCCL collective tests (e.g., all-reduce, all-to-all), with and without SHARP
  - Knowledge Transfer

Target Audience:

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

Prerequisites:

- Working knowledge of data center networking and Linux. Familiarity with InfiniBand concepts (subnet manager, LIDs, partitions) is helpful.
- All switches, GPU servers, optics and cables are racked, cabled to the agreed cabling plan and powered up prior to the kick-off.
- Required NVIDIA licenses and support entitlements (UFM and switch software support) are procured and available.
- A UFM appliance, or a server that meets UFM requirements, is available with management network access to the fabric.
- GPU servers are installed with a supported operating system, GPU drivers, DOCA-Host (or MLNX_OFED), NCCL, HPC-X/SHARP libraries 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.
- Ethernet front-end and storage network configuration beyond its connectivity to the GPU servers.
- NVLink and NVLink Switch (scale-up) domain configuration.

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) • Switch and adapter firmware and license validation • UFM readiness • Management network and access readiness
3.	Fabric Design	• High-level fabric design • Cabling, optics and rail connectivity plan
4.	Fabric Deployment	• UFM deployment with high availability • Switch firmware validation (up to 10 switches) • Topology and cabling validation (ibdiagnet) • Subnet manager, routing engine and partitions
5.	Performance Features	• Adaptive routing • Congestion control • NVIDIA SHARP
6.	GPU Server Fabric Integration (up to 6 nodes)	• ConnectX adapter configuration • Host-to-fabric connectivity verification
7.	Telemetry and Visibility	• Enable UFM telemetry • Alerts for link errors, congestion and port health
8.	Test and Validation	• Test plan • Connectivity, redundancy and subnet manager failover tests • RDMA performance tests • NCCL collective tests (with and without SHARP) • Test results report
9.	As-Built Documentation	• As-built configuration document for the deployed fabric

No	Deliverable	Service Details
10.	Knowledge Transfer	• Explain how to configure, operate and troubleshoot the deployed fabric in the customer's environment • Day-2 operations best practices