NVIDIA SuperPOD DGX B300 Systems, Spectrum-4 Ethernet and DC Busbar Power Reference Architecture

Release latest

NVIDIA Corporation

Aug 01, 2025

Reference Architecture

Document Number	RA-11337-001 V01
Publication Date	2025-05-29

Tip

The NVIDIA SuperPOD DGX B300 Systems, Spectrum-4 Ethernet and DC Busbar Power Reference Architecture Reference Architecture is also available as a PDF.

Chapter 1. Abstract

The NVIDIA DGX SuperPOD architecture has been designed to power the next-generation AI factories with unparalleled performance, scalability, and innovation that supports all customers in the enterprise, higher education, research, and the public sector. It is a physical twin of the main NVIDIA research and development system, meaning the company's infrastructure software, applications, and support are first tested and vetted on the same architecture.

This DGX SuperPOD Reference Architecture (RA) is based on DGX B300 systems powered by Blackwell GPUs. The RA discusses the components that define the scalable and modular architecture of DGX SuperPOD. DGX SuperPOD is built on the concept of scalable units (SU); each SU contains 64 DGX B300 systems, which enables rapid deployment of DGX SuperPOD of any size. The RA also includes details regarding the SU design and specifics of InfiniBand, Ethernet fabric topologies, storage system specifications, recommended rack layouts, and wiring guidelines.

This RA combines the latest NVIDIA technologies to help companies and industries develop their own AI factories. To achieve the most scalability, DGX SuperPOD is powered by several key NVIDIA technologies and solutions, including:

- NVIDIA DGX B300 system provides the most powerful computational building block for AI and HPC.
- NVIDIA XDR (800 Gbps) InfiniBand: High performance, low latency, and scalable network interconnect.

- NVIDIA Spectrum-X (800 Gbps) Ethernet: High performant, low latency, and scalable ethernet connectivity for compute interconnect.

- NVIDIA NVLink® technology—networking technologies that connect GPUs at the NVLink layer to provide unprecedented performance for the most demanding communication patterns.
- NVIDIA Mission Control: a unified operations and orchestration software stack for managing AI factories.

The DGX SuperPOD architecture integrates NVIDIA software solutions including NVIDIA Mission Control, NVIDIA AI Enterprise, CUDA, and NVIDIA Magnum IO™. These technologies help keep the system running at the highest levels of availability, performance, and with NVIDIA Enterprise Support (NVEX), keeps all components and applications running smoothly.

This reference architecture (RA) discusses the components that define the scalable and modular architecture of DGX SuperPOD. The system is built on the concept of scalable units (SU), each containing 64 DGX B300 systems, which provides for rapid deployment of systems of multiple sizes. This RA includes details regarding the SU design and specifics of InfiniBand, NVLink network, Ethernet fabric topologies, storage system specifications, recommended rack layouts, and wiring guides.

The NVIDIA DGX B300 is offered in a DC Busbar version as well as the more traditional AC Power supply version as do many of the fabric and other components. This RA specifically focuses on the DC Busbar version for modern and power-efficient datacenters. The centralized power shelves used for DC power supply provide better power efficiency while keeping the required performance and redundancy for AI Factories.

More information on the NVIDIA DGX B300 power supply version and other AC powered components can be found at: https://www.nvidia.com/en-us/data-center/dgx-superpod/.

1.1. Scope of DGX SuperPOD

NVIDIA DGX SuperPOD is a product that is defined by this RA with specific bill of materials for both hardware and software. It is designed for single-tenant, multi-user enterprise environment.

The designed mode of operation for DGX SuperPOD is as follows:

- NVIS led installation and white-glove services for initial bring up and commissioning
- Customer-owned infrastructure-as-a-product, where the customer is responsible for day-to-day operation with premium support from NVIDIA Technical Account Manager and continuous bring-up service.
- Functional and performance updates during the product lifecycle.
- Administration of the DGX SuperPOD with NVIDIA Mission Control software.
- User and application access using SLURM or NVIDIA Run:AI.

For customers who wish to modify their SuperPOD after delivery for additional capabilities such as general Kubernetes cluster access, or reconfiguration into a multi-tenant environment for enterprise private cloud or public cloud services, modifications are performed at customers' own discretion and risk. NVIDIA does not support these capabilities and performance with the DGX SuperPOD as a product. However, enterprise-level support will still be available for individual components (for both software and hardware) that are available in the SuperPOD.

Chapter 2. Key Components of the DGX SuperPOD

The DGX SuperPOD architecture has been designed to maximize performance for state-of-the-art model training, scale to exaflops of performance, provide the highest performance to storage and support all customers in the enterprise, higher education, research, and the public sector. It is a digital twin of the main NVIDIA research and development system, meaning the company's software, applications, and support structure are first tested and vetted on the same architecture. By using SUs, system deployment times are reduced from months to weeks. Leveraging the DGX SuperPOD design reduces time-to-solution and time-to-market of next generation models and applications.

DGX SuperPOD is the integration of key NVIDIA components, as well as storage solutions from partners certified to work in the DGX SuperPOD environment.

2.1. NVIDIA DGX B300 System

The NVIDIA DGX B300 system (Figure 2.1) is an AI powerhouse that enables enterprises to expand the frontiers of business innovation and optimization. The DGX B300 system delivers breakthrough AI performance with the most powerful chips ever built, in an eight GPU configuration. The NVIDIA Blackwell GPU architecture provides the latest technologies that brings months of computational effort down to days and hours, on some of the largest AI/ML workloads.

Figure 2.1: DGX B300 system

Some of the key highlights of the DGX B300 system when compared to the DGX B200 system include:

- DC Busbar powered, MGX-rack capable design for high density deployment in modern datacenters
- Alternative UshAC PSU powered appliance design
- 72 petaFLOPS FP8 training and 144 petaFLOPS FP4 inference
- Fifth generation of NVIDIA NVLink.
- 1,440 GB of aggregated HBM3 memory

2.1.1. MGX Racks and DGX Power Shelves

DGX B300 system features a DC busbar powered MGX v1.1 design that is similar in the DGX GB200/300 system. This design enables higher rack density, a better power efficiency, and data center level compatibility to DGX GB200/300 system.

The power shelf used for DGX B300 SuperPOD has six 5.5kW Power Shelves configured as N redundancy and can deliver up to 33kW of power. There are four total power shelves in a single DGX B300 rack. At the rear of the power shelf is a set of RJ45 ports used for power brake and current sharing feature. The power shelves are daisy chained to each other using these RJ45 ports. At the front of the power shelf is the BMC port. Figure 2.2 shows the front of the power shelf.

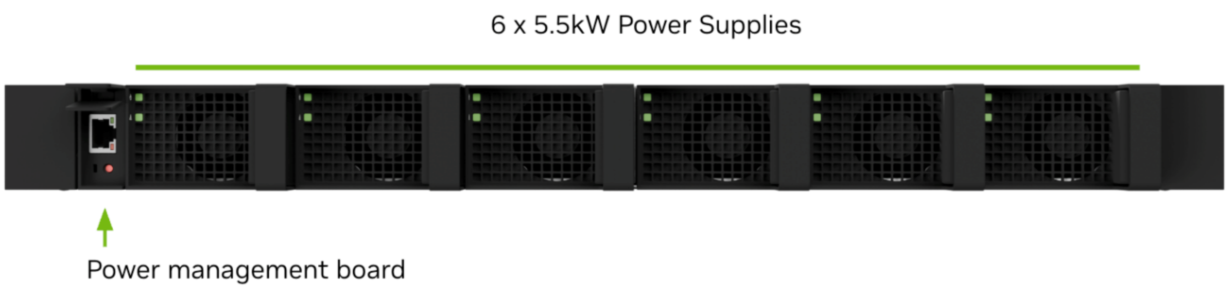

Figure 2.2: Power Shelf

2.2. NVIDIA Spectrum-X Technology

The NVIDIA Spectrum™-X Ethernet platform is designed specifically to improve the performance and efficiency of Ethernet-based AI clouds. This breakthrough technology achieves 1.6X better AI networking performance, along with consistent, and predictable performance. Spectrum-X is built on network innovations powered by the tight coupling of the NVIDIA Spectrum-4 Ethernet switch and NVIDIA® ConnectX-8 Smart NIC. Spectrum-X network optimizations reduce runtimes of massive transformer-based generative AI models and deliver faster time to insight.

Spectrum-X 2.0 provides the same 800Gbps connectivity with comparable latency characteristics when compared with XDR InfiniBand based on more affordable switches and optics.

2.3. NVIDIA InfiniBand Technology

InfiniBand is a high-performance, low latency, RDMA capable networking technology, proven over 20 years in the harshest compute environments to provide the best inter-node network performance. InfiniBand continues to evolve and lead data center network performance.

NVIDIA InfiniBand XDR has a peak speed of 800 Gbps per direction with an extremely low port-to-port latency and is backwards compatible with the previous generations of InfiniBand specifications. InfiniBand is more than just peak bandwidth and low latency. InfiniBand provides additional features to optimize performance including Adaptive Routing (AR), collective communication with SHARP™, dynamic network healing with SHIELD™, and supports several network topologies.

2.4. NVIDIA Mission Control

The DGX B300 SuperPOD Reference Architecture represents the best practices for building high-performance AI factories. There is flexibility in how these systems can be presented to customers and users. NVIDIA Mission Control software is used to manage all DGX B300 SuperPOD deployments.

NVIDIA Mission Control is a sophisticated full-stack software solution. As an essential part of the DGX SuperPOD experience, it optimizes developer workload performance and resiliency, ensures unmatched uptime with automated failure handling, and provides unified cluster-scale telemetry and manageability. Key features include full-stack resiliency, predictive maintenance, unified error reporting, data center optimizations, cluster health checks, and automated node management.

NVIDIA Mission Control software incorporates the same technology that NVIDIA uses to manage thousands of systems for our award-winning data scientists and provides an immediate path to AI Factory for organizations that need the best of the best.

DGX SuperPOD is to be deployed on-premises, meaning the customer owns and manages the hardware. This can be within a customer's data center or co-located at a commercial data center. In each case the customer owns the hardware, the service it provides, and is responsible for their cluster infrastructure as well as providing the building management system for integration.

2.5. Components

The hardware components of DGX SuperPOD are described in Table 2.1. The software components are shown in Table 2.2.

Table 2.1: DGX SuperPOD hardware components by NVIDIA

Component	NVIDIA Technology	Description
Compute nodes	NVIDIA DGX B300 system with eight Blackwell Ultra GPUs	The world's premier purpose-built AI systems featuring NVIDIA Blackwell Ultra GPUs, fifth-generation NVIDIA NVLink, and fourth-generation NVIDIA NVSwitch™ technologies.
Compute node transceiver and cable	NVIDIA OSFP twin port flat top transceiver, MMF passive fiber cable	Transceiver and cables for DGX nodes
Compute fabric	NVIDIA Quantum-3 Q3400-RA 800 Gbps InfiniBand	Rail-optimized, non-blocking, twin-plane, fat tree topology for next-generation extreme-scale AI factory.
InfiniBand Storage fabric	NVIDIA Quantum QM9700 NDR 400 Gbps InfiniBand Switch	The fabric is optimized to match peak performance of the configured storage array
InfiniBand storage fabric switch transceiver and cable	NVIDIA QSFP single port flat top transceiver, MMF passive fiber cable	Transceiver and cables for InfiniBand switches
Ethernet Storage fabric	NVIDIA Spectrum-4 SN5610/SN5600D 800 Gbps Ethernet	Optional storage fabric for ethernet based storage solutions
Ethernet Storage fabric switch transceiver and cable	NVIDIA QSFP single port flat top transceiver 800GB (NIC side) and NVIDIA OSFP twin port finned (switch side) transceiver 1600GB MMF passive fiber cable	Transceiver and cables for Spectrum-4 switches
Storage InfiniBand fabric management	NVIDIA Unified Fabric Manager 3.5 Appliance, Enterprise Edition	NVIDIA UFM combines enhanced, real-time network telemetry with AI powered cyber intelligence and analytics to manage scale-out InfiniBand data centers
In-band management network	NVIDIA SN5610/SN5600D switch	64 port 800 Gbps and up to 256 ports of 200 Gbps Ethernet switch providing high port density with high performance
Ethernet In-band fabric switch transceiver and cable	NVIDIA QSFP single port flat top transceiver 400GB or NVIDIA OSFP twin port finned transceiver 800GB MMF passive fiber cable	Transceiver and cables for Spectrum-4 switches
Out-of-band (OOB) management network	NVIDIA SN2201(M) switch	48 port 1 Gbps Ethernet and 4 x 100 Gbps switch leveraging copper ports to minimize complexity

Table 2.2: DGX SuperPOD software components

Component	Description
NVIDIA Mission Control[1]	Simplified AI data center operations, cluster management, and workload orchestration with agility, resilience, and hyperscale efficiency for enterprises.
NVIDIA Run:ai	Cloud-native AI workload and GPU orchestration platform enabling fractional, full, and multi-node support for the entire enterprise AI lifecycle including interactive development environments, training and inference
NVIDIA AI Enterprise	NVIDIA AI Enterprise is an end-to-end, cloud-native software platform that accelerates data science pipelines and streamlines development and deployment of production-grade co-pilots and other generative AI applications.
Magnum IO	Enables increased performance for AI and HPC
NVIDIA NGC	The NGC catalog provides a collection of GPU-optimized containers for AI and HPC
Slurm	A classic workload manager used to manage complex workloads in a multi-node, batch-style, compute environment

Note

NVIDIA Mission Control now includes Base Command Manager and Run:ai functionality. No separate purchase is needed. SuperPOD only supports multiteam environments through Base Command Manager; multitenancy is not supported with SuperPOD currently.

2.6. Design Requirements

DGX SuperPOD is designed to minimize system bottlenecks throughout the tightly coupled configuration to provide the best performance and application scalability. Each subsystem has been thoughtfully designed to meet this goal. In addition, the overall design remains flexible so that data center requirements can be tailored to better integrate into existing data centers.

2.6.1. System Design

DGX SuperPOD is optimized for a customers' particular workload of multi-node AI and HPC applications:

- A modular architecture based on SUs of 64 DGX B300 systems each.
- MGX-rack based, DC busbar powered, integrated data-center scale design.
- A fully tested system scales to four SUs, but larger deployments can be built based on customer requirements.
- Single rack that can support up to four DGX B300 systems per rack, enabling modification to accommodate different data center requirements.
- Storage partner equipment that has been certified to work in DGX SuperPOD environments.

[1] NVIDIA Mission Control includes Base Command Manager and Run:ai functionality. No separate purchase needed.

- Full system support, including compute, storage, network, and software is provided by NVIDIA Enterprise Experience (NVEX).

2.6.2. Compute Fabric

- The compute fabric is rail-optimized, twin-planar, full-fat tree topology
- Managed Quantum-3 and Spectrum-X switches are used throughout the design to provide better management of the fabric.

2.6.3. Storage Fabric (High Speed Storage)

The storage fabric provides high bandwidth to shared storage. It also has the following characteristics:

- It is independent of the compute fabric to maximize performance of both storage and application performance.
- Provides single-node bandwidth of at least 40 Gbps to each DGX B300 system.
- Storage is provided over InfiniBand or RDMA over Converged Ethernet to provide maximum performance and minimize CPU overhead.
- It is flexible and can scale to meet specific capacity and bandwidth requirements.
- Connectivity to management nodes required to provide storage access independent of compute nodes.

2.6.4. In-Band Management Network

- The in-band management network fabric is Ethernet-based and is used for node provisioning, data movement, Internet access, and other services that must be accessible by the users.
- The in-band management network connections for compute and management nodes operate at 200 Gbps and are bonded for resiliency.

2.6.5. Out-of-Band Management Network

The OOB management network connects all the base management controller (BMC) ports, as well as other devices that should be physically isolated from users. The Switch Management Network is a subset of the Out-Of-Band Network that provides additional security and resiliency.

2.6.6. Storage Requirements

The DGX SuperPOD compute architecture must be paired with a high-performance, balanced, storage system to maximize overall system performance. DGX SuperPOD is designed to use two separate storage systems, high-performance storage (HPS) and user storage, optimized for key operations of throughput, parallel I/O, as well as higher IOPS and metadata workloads.

2.6.7. High-Performance Storage

High-Performance Storage is provided via InfiniBand connected storage from a DGX SuperPOD certified storage partner, and is engineered and tested with the following attributes in mind:

- High-performance, resilient, POSIX-style file system optimized for multi-threaded read and write operations across multiple nodes.
- RDMA on InfiniBand or Ethernet support
- Local system RAM for transparent caching of data.
- Leverage local flash device transparently for read and write caching.

The specific storage fabric topology, capacity, and components are determined by the DGX SuperPOD certified storage partner as part of the DGX SuperPOD design process.

2.6.8. User Storage

User Storage differs from High-Performance storage in that it exposes an NFS share on the in-band management fabric for multiple uses. It is typically used for "home directory" type usage (especially with clusters deployed with Slurm), administrative scratch space, and shared storage as needed by DGX SuperPOD components in a High Availability configuration (e.g., Base Command Manager), and log files.

With that in mind, User Storage has the following requirements:

- 100 Gb/s DR1 connectivity is required.
- Designed for high metadata performance, IOPS, and key enterprise features such as checkpointing. This is different than the HPS, which is optimized for parallel I/O and large capacity.
- Communicate over Ethernet, using NFS.

User storage in a DGX SuperPOD is often satisfied with existing NFS servers already deployed, such that a new export is created and made accessible to the DGX SuperPOD's in-band management network. For the best performance, we require 100 Gb/s minimum bandwidth for the user storage.

Chapter 3. DGX SuperPOD Architecture

The DGX SuperPOD architecture is a combination of DGX systems, Ethernet networking, InfiniBand Networking, management nodes, and storage. Figure 3 the rack layout of a single SU. With DGX SuperPOD B300, we introduce the MGX-based, DC busbar powered design for the best datacenter density and efficiency.

In our reference design, for DGX B300s are within a single rack. The rack-level power consumption per rack exceeds 50 kW. The rack layout can be adjusted to meet local data center requirements, such as maximum power per rack and rack layout between DGX systems and supporting equipment to meet local needs for power and cooling distribution.

For legacy datacenter without the possibility for DC busbar, it is still possible to build SuperPOD with traditional PDU and AC powered EIA racks.

Figure 3.1 shows 64 x NVIDIA DGX B300 Busbar systems in MGX racks each with DC Power shelfs and MGX rack stiffeners.

Figure 3.2 shows an example management rack configuration with networking switches, management servers, storage arrays, and UFM appliances. Sizes and quantities will vary depending upon models used.

This reference architecture is focused on 4 SU units with 256 DGX nodes. DGX SuperPOD can scale to much larger configurations up to and beyond 64 SU with 2000+ DGX B300 nodes. See Table 3.1 for more information.

Table 3.1: Larger DGX SuperPOD component counts

SU Count	**Node Count**	**GPU Count**	**Cable Count**			
			Leaf	**Spine**	**Node-Leaf**	**Leaf-Spine**
1	64	512	16	8	512	512
2	128	1024	32	16	1024	1024
4	256	2048	64	32	2048	2048
8	512	4096	128	64	4096	4096
16	1024	8192	256	128	8192	8192

Contact NVIDIA for information regarding DGX SuperPOD solutions of four scalable units or more.

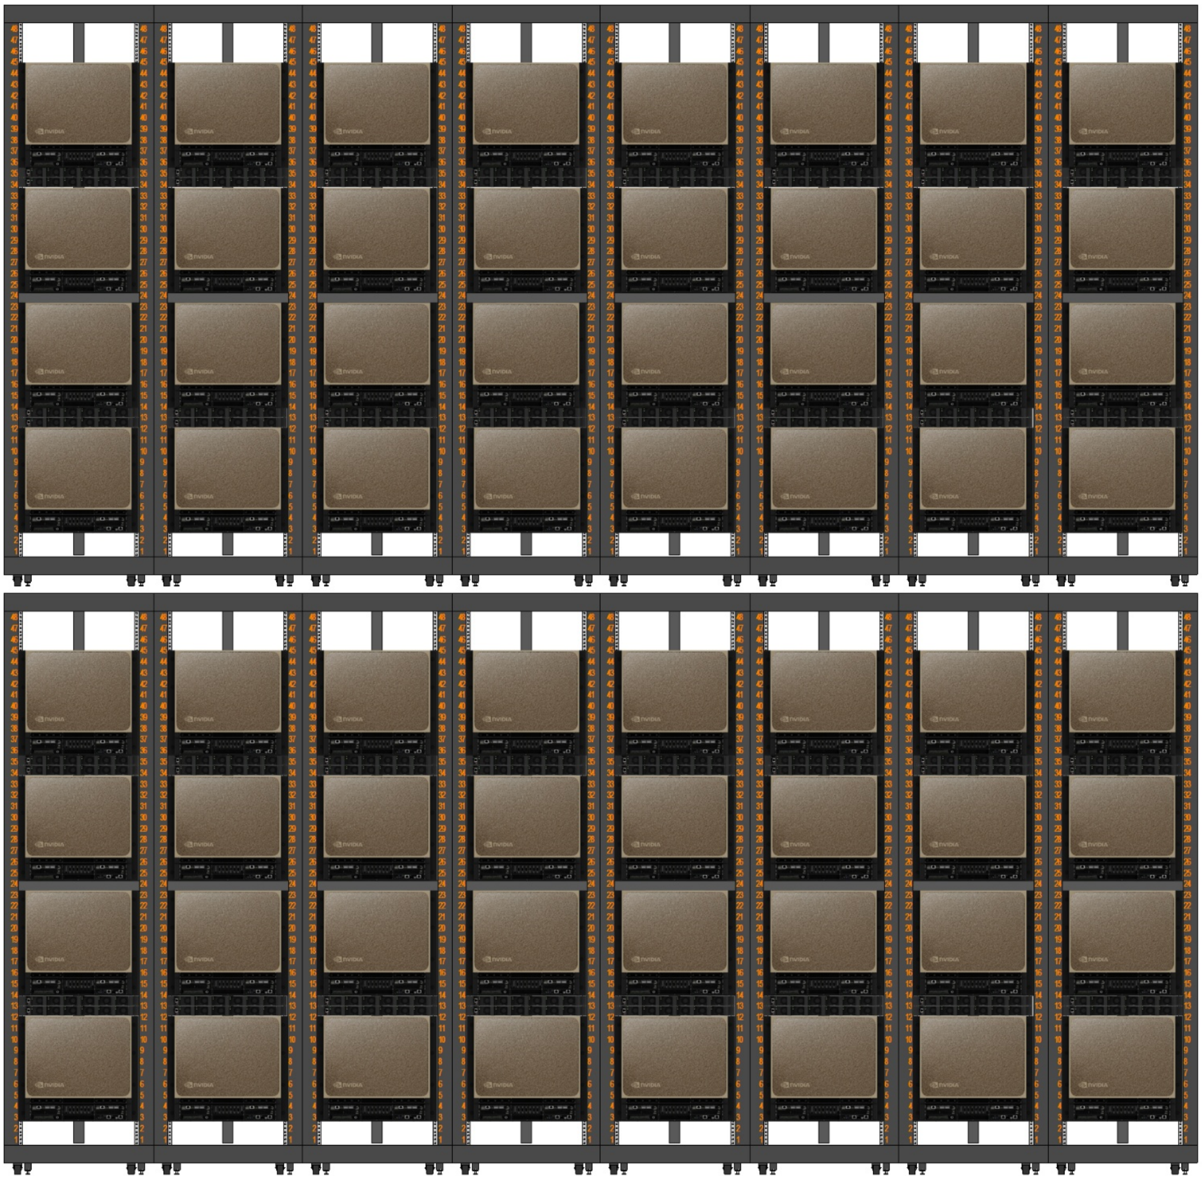

Figure 3.1: DGX B300 in Racks

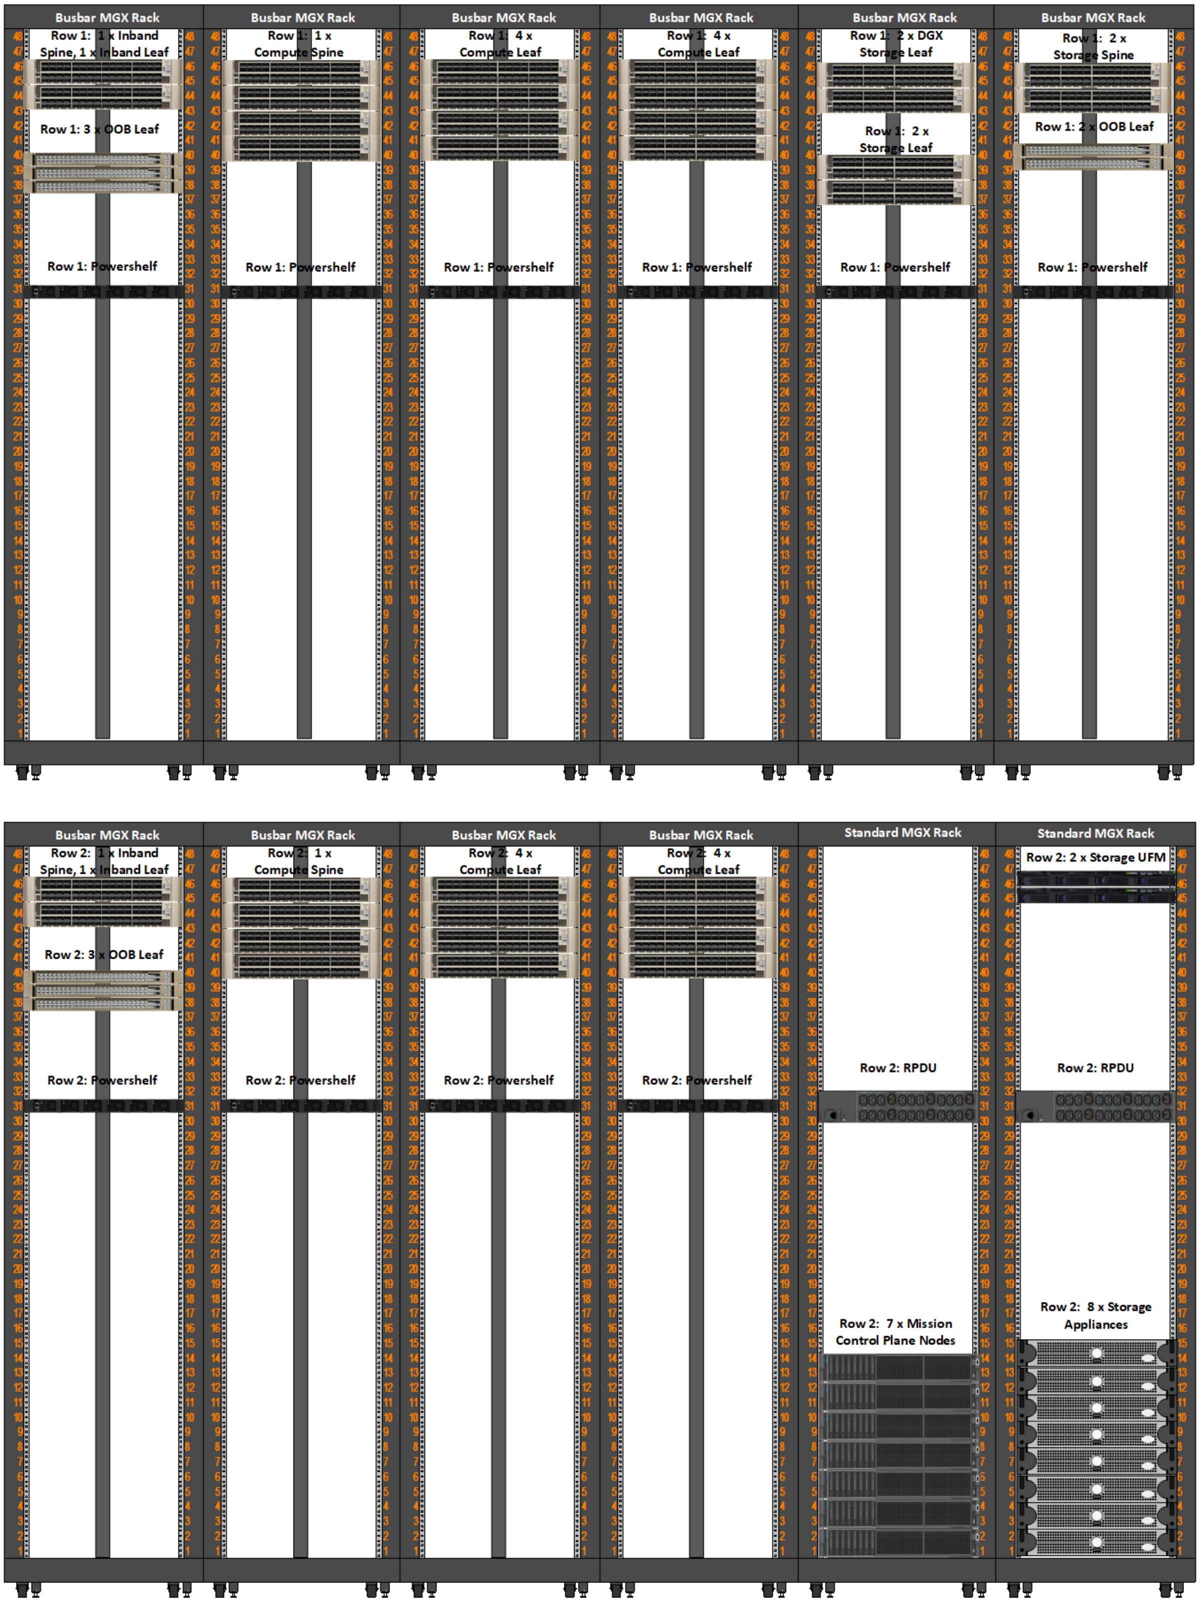

Figure 3.2: Management rack configuration

Chapter 4. Network Fabrics

Building systems by SU provide the most efficient designs. However, if a different node count is required due to budgetary constraints, data center constraints, or other needs, the fabric should be designed to support the full SU, including leaf switches and leaf-spine cables, and leave the portion of the fabric unused where these nodes would be located. This will ensure optimal traffic routing and ensure that performance is consistent across all portions of the fabric.

DGX SuperPOD configurations utilize four network fabrics:

- Compute Fabric
- Storage Fabric
- Ethernet Fabric with two network segments
  - Inband Network
  - Out-of-band Network

Figure 4.1 shows the ports on the back of the DGX B300 CPU tray and the connectivity provided. The compute fabric ports are on the edges of the compute tray and provide access to all eight GPUs. Each pair of in-band management and storage ports provide parallel pathways into the DGX B300 system for increased performance. The BMC port is used for BMC access.

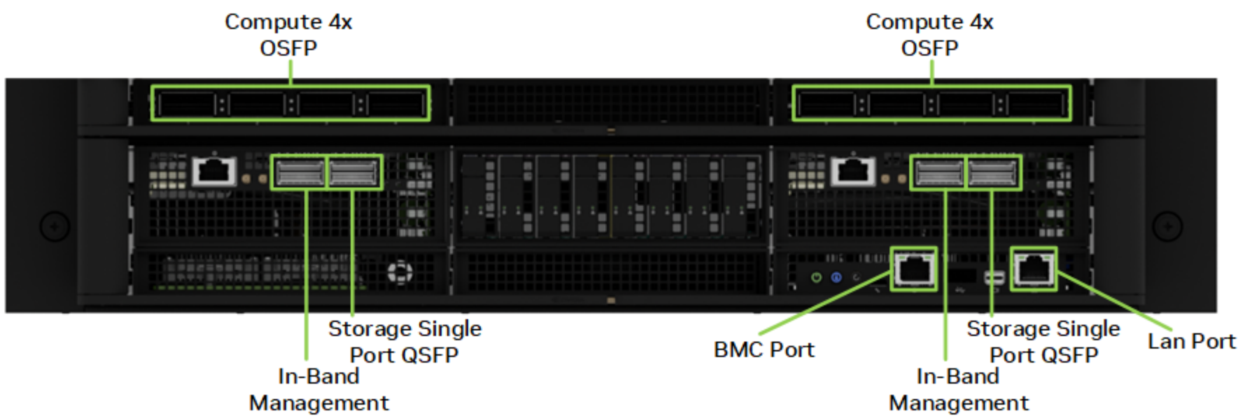

Figure 4.1: DGX B300 network ports

4.1. Compute Fabric

Figure 4.2 shows the compute fabric layout for the full 512-node DGX SuperPOD. Each group of 64 nodes is rail-aligned. Traffic per rail of the DGX B300 systems is always one hop away from the other 64 nodes in a SU. Traffic between nodes, or between rails, traverses the spine layer.

The Spectrum-X based compute fabric features a twin-planar design (denoted in blue and green). Each GPU has 2x 400GbE connectivity through two different planes. The multi-planar design provides not only the high performance and low latency required for AI training, but it also enhances the fault tolerance as two independent data paths are available. A single switch failure, transceiver failure or cable failure will not lead to a catastrophic job abortion. Instead, the application can continue to operate on ½ of the original bandwidth.

Another benefit of the multi-planar design is a significantly increased SuperPOD dimension with a two-layer fabric using spine-leaf switches only. As these two independent planes are not connected at a core layer level, twice as many nodes can be connected to the same fabric on 2-layer fabric, reducing the total cost for construction while keeping the high standard for performance.

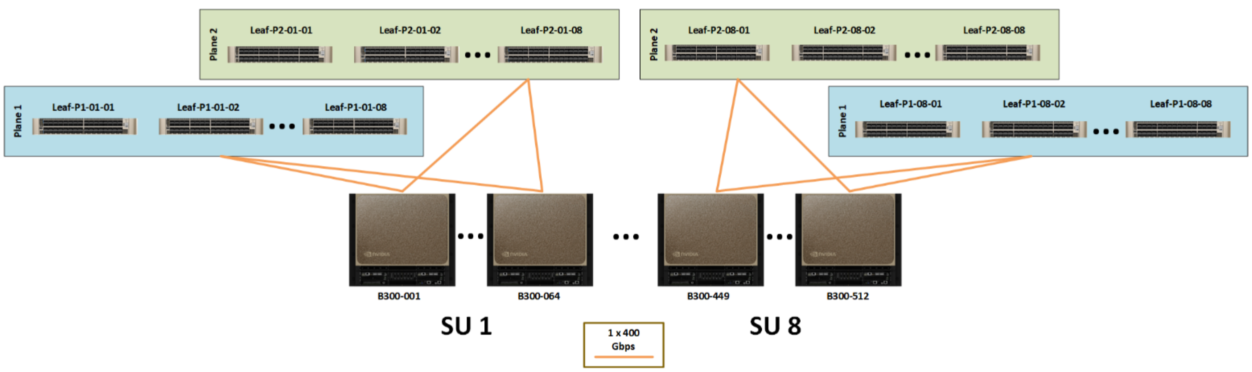

Figure 4.2: Compute fabric for full 512-node DGX SuperPOD

The multi-planar (in this case – 2 planes) fabric provides auto-load balancing based on both hardware and software load balancer. With NVIDIA Route software and the onboard plane load balancer unit on the ConnectX-8 NIC, traffic can be automatically routed and balanced over the best path. In case of a single link failure, GPU application will continue to operate as it will not receive an aborted connection error through the RDMA Driver.

NVIDIA NetQ delivers real-time network telemetry and monitoring for Spectrum-X, enabling operators to gain visibility into the health and performance of the compute fabric. By providing tools for rapid troubleshooting and network validation, NetQ helps ensure reliable operation and efficient management of large-scale AI workloads in DGX SuperPOD environments.

4.2. InfiniBand Storage Fabric

The storage fabric employs an InfiniBand network fabric that is essential to maximum bandwidth (Figure 4.3). This is because the I/O per-node for the DGX SuperPOD must exceed 40 GBps. High bandwidth requirements with advanced fabric management features, such as congestion control and AR, provide significant benefits for the storage fabric.

There are two storage fabric options for NVIDIA DGX SuperPOD. The InfiniBand storage fabric uses MQM9700-NS2R (AC power) or *MQM9701-NS2R* (DC power) NDR switches (Figure 4.4). The high-speed storage devices are connected at a 1:1 port to uplink ratio. The DGX B300 system connections are

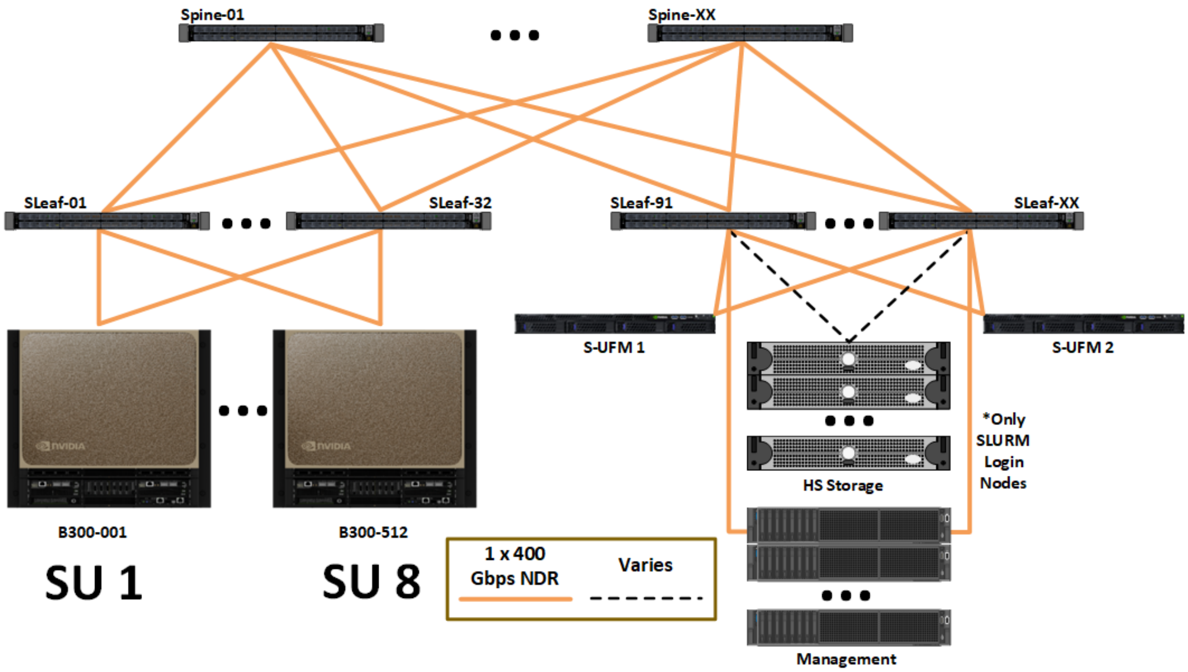

Figure 4.3: Storage fabric logical design

slightly oversubscribed with a ratio near 4:3 with adjustments as needed to enable more storage flexibility regarding cost and performance.

Figure 4.4: MQM9701-NS2R switch

4.3. Ethernet Storage Fabric

There are two storage fabric options for NVIDIA DGX SuperPOD. The Ethernet storage fabric employs a high-speed Ethernet network fabric that is essential to maximum bandwidth (Figure 4.5). This is because the I/O per-node for the DGX SuperPOD must exceed 40 GBps. High bandwidth requirements with advanced fabric management features, provide significant benefits for the storage fabric. Supported ethernet storage appliance leverages RoCE to provide best performance and minimizes CPU usage.

The storage fabric uses SN5600 (AC power) or SN5600D (DC power) switches (Figure 4.6). The high-speed storage devices are connected at a 1:1 port to uplink ratio. The DGX B300 system connections are slightly oversubscribed with a ratio near 4:3 with adjustments as needed to enable more storage flexibility regarding cost and performance.

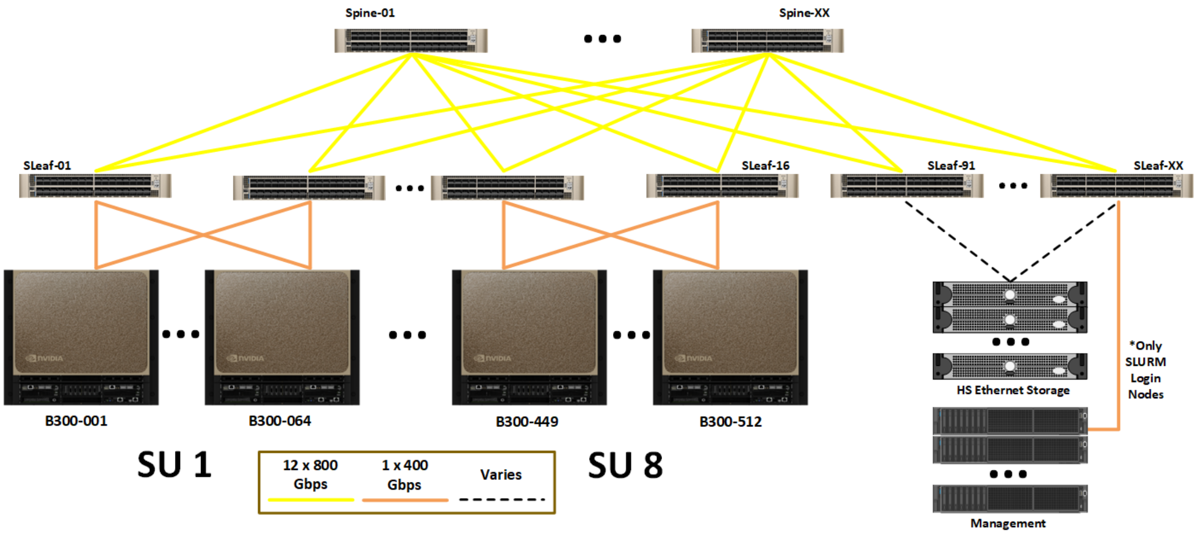

Figure 4.5: Ethernet Storage fabric logical design

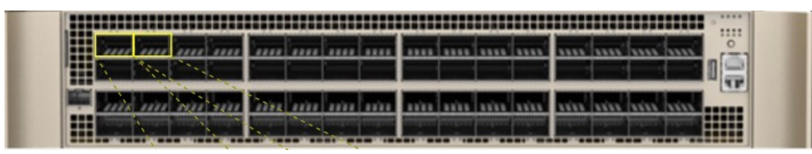

Figure 4.6: NVIDIA Spectrum SN5600D Ethernet Switch

4.4. Network Segmentation of the Ethernet Fabric

The ethernet fabric is segmented into these segments on the SuperPOD:

- *In-band Network*
- *Out-of-band Management Network*

In this reference design, the entire ethernet fabric (except for potential dedicated storage and compute fabric), is built on a common physical network, and segregated with VXLAN and EVPN to achieve network isolation between control traffic and admin traffic. VTEPs for different network segments terminate on the leaf switches (either SN2201 or SN5600D) for the access to different networks.

In the following, we introduce these networks in detail.

4.4.1. In-band Management Network

The in-band management network provides several key functions:

- Connects all the services that manage the cluster.
- Enables access to the lower-speed, NFS tier of storage.
- Provides uplink (border) connectivity for the in-cluster services such as Mission Control, Base Command Manager, Slurm, and Kubernetes to other services outside of the cluster such as the

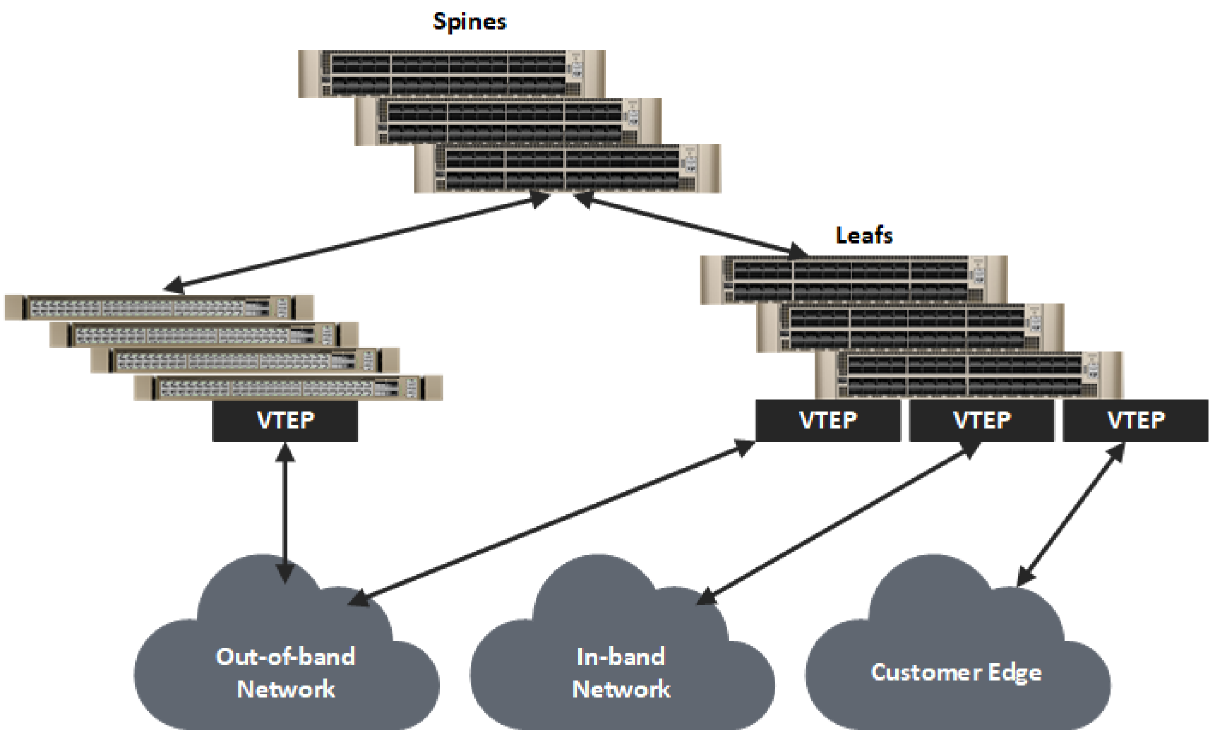

Figure 4.7: Network Segmentation diagram

NGC registry, code repositories, and data sources.

- Provides end user access to the Slurm head nodes and Kubernetes services.

The in-band management network uses SN5600D switches (Figure 4.6 and Figure 4.8)

4.4.2. Out-of-band Management Network

Figure 4.9 shows the OOB Ethernet fabric. It connects the management ports of all devices including DGX B300 compute trays, switches, and management servers, storage, networking gear, rack PDUs, and all other devices. These are separated onto their own network. There is no use-case where a non-privileged user needs direct access to these ports and are secured using logical network separation.

The OOB network carries all IPMI related control traffic and serves as the network for fabric management of the compute InfiniBand fabric and compute fabric.

The OOB network is physically rolled up into the aggregation layer (spine layer) of each SU as a dedicated VXLAN. The OOB management network use SN2201 switches, shown in Figure 4.10.

The OOB management network uses SN2201 switches (Figure 4.10).

4.4.2.1 Customer Edge Connectivity

For connecting DGX SuperPOD to customer edge for uplink and customer corporate network access, we recommend at least 2x 100GbE links with DR1 single-mode connectivity to cope with the growing demand on high-speed data transfer into and from DGX SuperPOD.

For route handover, we prepare BGP protocol to peer with customer's network. Routes to/from in-band and out-of-band are announced.

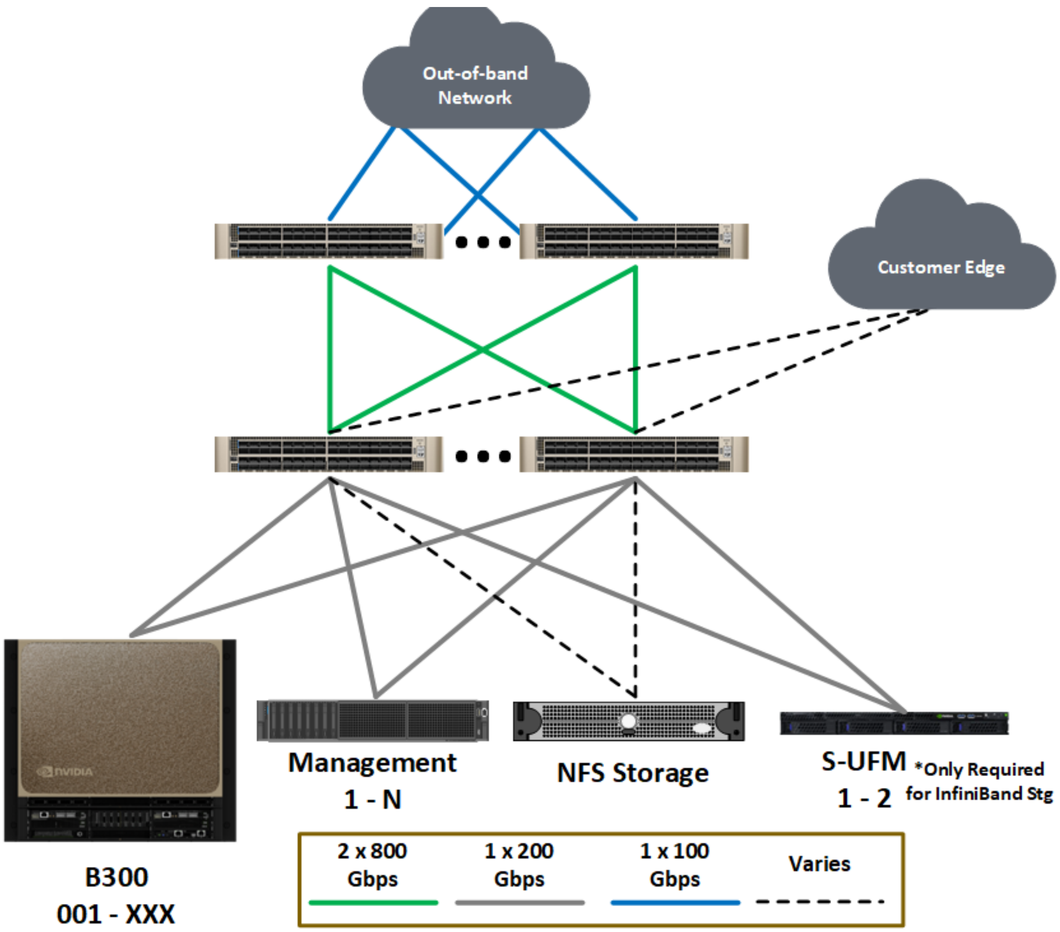

Figure 4.8: In-band Ethernet network

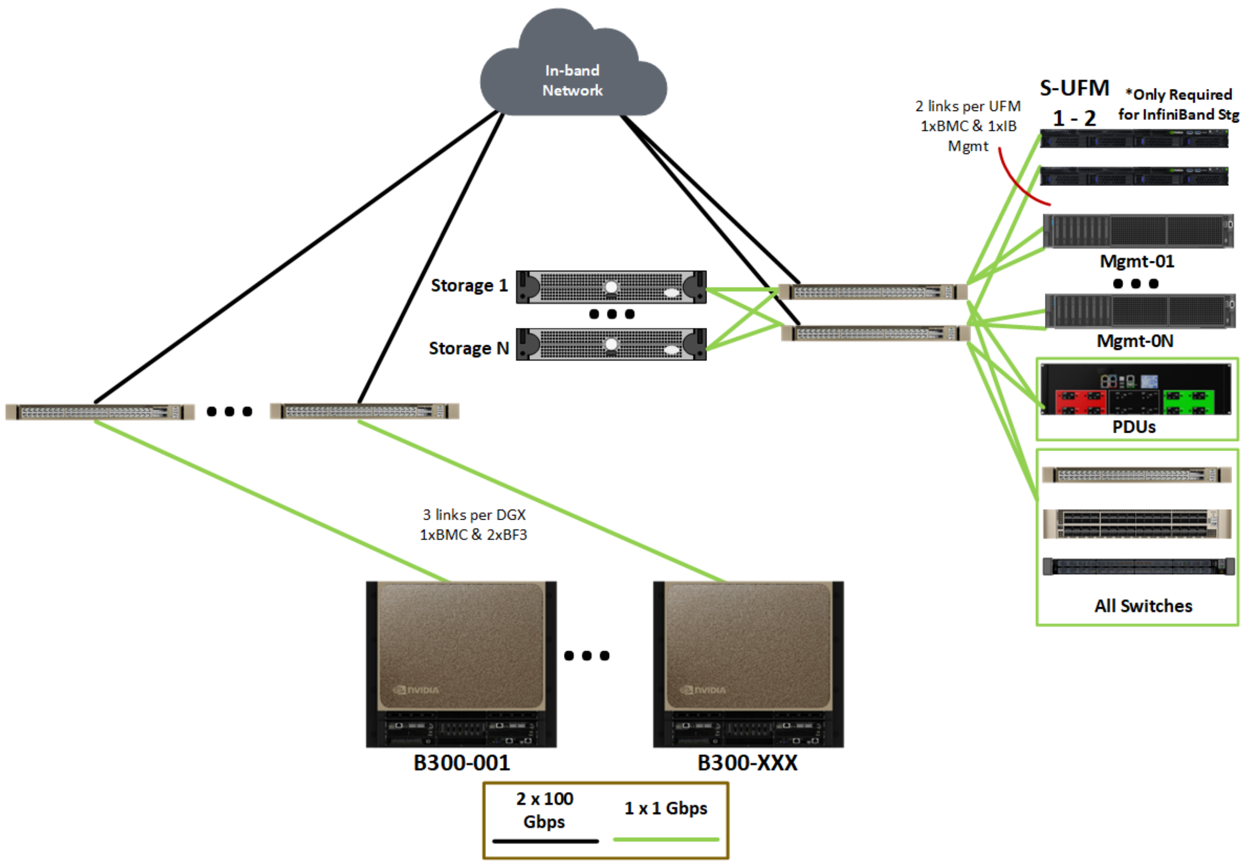

Figure 4.9: Logical OOB management network layout

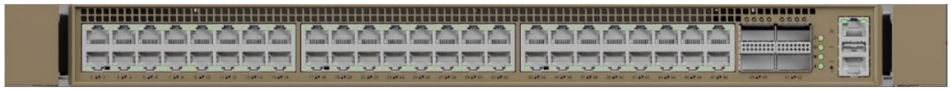

Figure 4.10: SN2201 switch

Customers, who cannot provide DR1-based connectivity, are encouraged to use a pair of dedicated border leaf switches to enhance the connectivity (not as part of the SuperPOD).

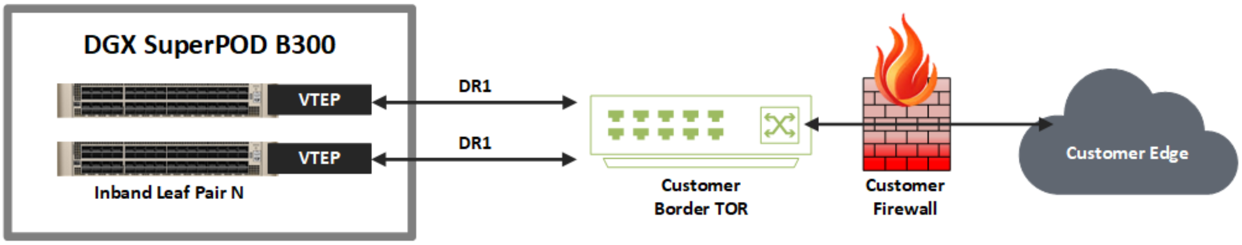

Figure 4.11: Customer Edge Example

4.4.2.2 User Storage Connectivity

For the operation with DGX SuperPOD, customers are required to provide an NFS-based home storage / configuration storage, to be integrated with NVIDIA Mission Control.

SuperPOD with DGX B300 supports storage integration with single-mode DR1 (or faster) connectivity for user storage. User storage is connected to the leaf-layer SN5600Ds.

If customer provided User Storage does not support DR1 connectivity, we recommend they implement border TORs that would be able to connect to the SN5600Ds.

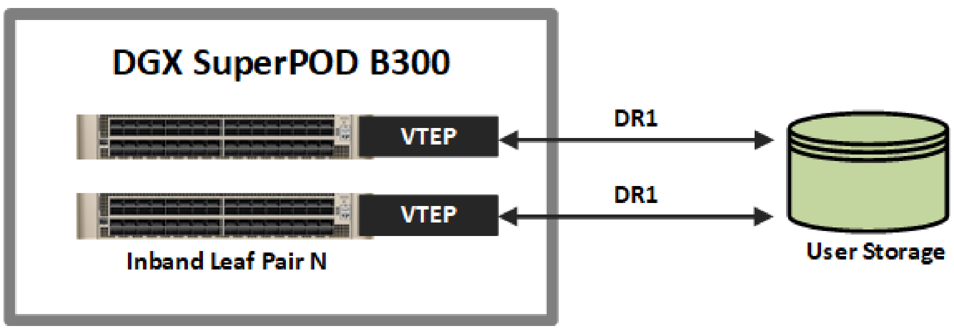

Figure 4.12: User Storage Example

Chapter 5. High Performance Storage Architecture

Data, lots of data, is the key to the development of accurate deep learning (DL) models. Data volume continues to grow exponentially, and data used to train individual models continues to grow as well. Data format, not just volume can play a key factor in the rate at which data is accessed so storage system performance must scale commensurately.

The key I/O operation in DL training is re-read. It is not just that data is read, but it must be reused again and again due to the iterative nature of DL training. Pure read performance still is important as some model types can train in a fraction of an epoch (ex: some recommender models) and inference of existing can be highly I/O intensive, much more so than training. Write performance can also be important. As DL models grow and time-to-train, writing checkpoints is necessary for fault tolerance. The size of checkpoint files can be terabytes in size and while not written frequently are typically written synchronously than blocks and block forward progress of DL models.

Ideally, data is cached during the first read of the dataset, so data does not have to be retrieved across the network. Shared filesystems typically use RAM as the first layer of cache. Reading files from cache can be an order of magnitude faster than from remote storage. In addition, the DGX B300 system provides local NVMe storage that can also be used for caching or staging data.

DGX SuperPOD is designed to support all workloads, but the storage performance required to maximize training performance can vary depending on the type of model and dataset. The guidelines in Table 5.1 and Table 5.2 are provided to help determine the I/O levels required for different types of models.

Table 5.1: Storage performance requirements

Level	Work Description	Dataset Size
Standard	Multiple concurrent LLM or fine-tuning training jobs and periodic checkpoints, where the compute requirements dominate the data I/O requirements significantly.	Most datasets can fit within the local compute systems' memory cache during training. The datasets are single modality, and models have millions of parameters.
Enhanced	Multiple concurrent multimodal training jobs and periodic checkpoints, where the data I/O performance is an important factor for end-to-end training time.	Datasets are too large to fit into local compute systems' memory cache requiring more I/O during training, not enough to obviate the need for frequent I/O. The datasets have multiple modalities and models have billions (or higher) of parameters.

Table 5.2: Guidelines for storage performance

Performance Characteristic	Good (GBps)	Better (GBps)
Single SU aggregate system read	80	250
Single SU aggregate system write	40	124
8 SU aggregate system read	640	2000
8 SU aggregate system write	320	992

High-speed storage provides a shared view of an organization's data to all nodes. It must be optimized for small, random I/O patterns, and provide high peak node performance and high aggregate filesystem performance to meet the variety of workloads an organization may encounter. High-speed storage should support both efficient multi-threaded reads and writes from a single system, but most DL workloads will be read-dominant.

Use cases in automotive and other computer vision-related tasks, where high-resolution images are used for training (and in some cases are uncompressed) involve datasets that easily exceed 30 TB in size. In these cases, 4 GBps per GPU of read performance is required.

While NLP and LLM cases often do not require as much read performance for training, peak performance for reads and writes are needed for creating and reading checkpoint files. This is a synchronous operation and training stops during this phase. If you are looking for best end-to-end training performance, do not ignore I/O operations for checkpoints. Consider at least ½ of the read performance as recommended write performance for LLM and large model use cases.

The preceding metrics assume a variety of workloads, datasets, and need for training locally and directly from the high-speed storage system. It is best to characterize workloads and organizational needs before finalizing performance and capacity requirements.

Chapter 6. Management Servers

For DGX SuperPOD with B300, we provide 2 means of Cluster Access for End Users (AI Practitioners):

- With SLURM – best fit for pretraining-type of workflow, with bare-metal or containerized access to the compute infrastructure.
- With Kubernetes and NVIDIA Run:AI.

To support the operations, monitoring, and installation of the DGX SuperPOD, a set of management servers is required.

6.1. Management Server Quantities and Connectivity

These nodes are used for the following purposes:

- Base Command Manager in High Availability (HA): 2 Nodes, connects to Inband and OOB
- K8s Management Server: 3 Nodes, connects to Inband and storage
- SLURM Login Nodes, 2 Nodes, connects to Inband and storage
- All devices connect to OOB with 1GbE for IPMI/Redfish as well.

Table 6.1: Control Plane Node Form Factor Requirement

System	Power Supply, Form Factor
B300 PDU	AC, EIA
B300 Busbar	DC, MGX

Chapter 7. DGX SuperPOD Software

DGX SuperPOD is an integrated hardware and software solution. The included software (Figure 7.1) is optimized for AI from top to bottom. From the accelerated frameworks and workflow management to system management and low-level operating system (OS) optimizations, every part of the stack is designed to maximize the performance and value of DGX SuperPOD.

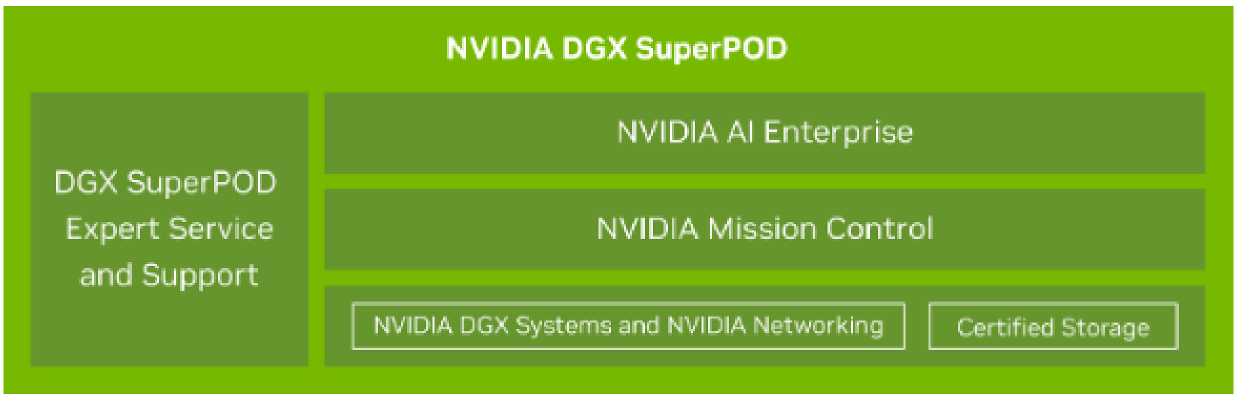

Figure 7.1: DGX SuperPOD high-level software architecture

7.1. NVIDIA Mission Control

NVIDIA Mission Control is the standard for every DGX SuperPOD with DGX B300. It streamlines AI operations, from workloads to infrastructure, with world-class expertise delivered as software powering AI data centers, bringing instant agility for inference and training while providing full-stack intelligence for infrastructure resilience. Every enterprise can run AI with hyperscale efficiency, simplifying and accelerating AI experimentation.

NVIDIA Mission Control includes NVIDIA Base Command Manager and NVIDIA Run:ai functionality as part of integrated software delivery across configuration, validation, and operations.

7.2. NVIDIA Base Command Manager

NVIDIA Base Command Manager offers fast deployment and basic end-to-end management for heterogeneous AI and HPC clusters. It automates provisioning and administration of DGX SuperPOD from hundreds to thousands of nodes.

7.3. NVIDIA Run:ai

NVIDIA Run:ai is cloud native AI workload and GPU orchestration platform that simplifies and accelerates AI and machine learning with DGX SuperPOD through dynamic resource allocation, comprehensive AI lifecycle support, strategic resource management and advanced scheduling. Run:ai maximizes GPU efficiency and workload capacity. Its policy engine, open architecture, and visibility into AI workloads foster strategic alignment with business objectives.

7.4. NVIDIA NGC

NVIDIA NGC provides software to meet the needs of data scientists, developers, and researchers with various levels of AI expertise.

Software hosted on NGC undergoes scans against an aggregated set of common vulnerabilities and exposures (CVEs), crypto, and private keys.

Software from the NGC catalog is tested and ensured to scale to multiple GPUs and in some cases, to scale to multi-node, ensuring users maximize the use of their DGX SuperPOD.

7.5. NVIDIA AI Enterprise

NVIDIA AI Enterprise is the end-to-end software platform that brings generative AI into reach for every enterprise, providing the fastest and most efficient runtime for generative AI foundation models developed with the NVIDIA DGX platform. With production-grade security, stability, and manageability, it streamlines the development of generative AI solutions. NVIDIA AI Enterprise is included with DGX SuperPOD for enterprise developers to access pretrained models, optimized frameworks, microservices, accelerated libraries, and enterprise support.

Chapter 8. Summary

DGX SuperPOD with NVIDIA DGX B300 systems is the next generation of data center scale architecture to meet the demanding and growing needs of AI training. This RA document for DGX SuperPOD represents the architecture used by NVIDIA for our own AI model and HPC research and development. DGX SuperPOD continues to build upon its high-performance roots to enable training of the largest NLP models, support the expansive needs of training models for automotive applications, and scaling-up recommender models for greater accuracy and faster turn-around-time.

DGX SuperPOD represents a complete system of not just hardware but all the necessary software to accelerate time-to-deployment, streamline system management, proactively identify system issues. The combination of all these components keeps systems running reliably, with maximum performance, and enables users to push the bounds of state-of-the-art. The platform is designed to both support the workloads of today and grow to support tomorrow's applications.

Chapter 9. Major Components

Major components for the DGX SuperPOD configuration are listed in Table 9.1. These are representative of the configuration and must be finalized based on actual design.

Table 9.1: Major components of the 4 SU, 256-node DGX SuperPOD

Count	Component	Recommended Model
DGX Compute and Racks		
256	DGX B300 DC Busbar Appliances	DGXB300-DC
72	19" MGX rack with bubar	754-24972-0046-001
8	19" MGX rack without busbar	754-24972-0047-001
160	rPDU for MGX rack	PX4-57A7I2U-C8E7V2
288	33KW power shelf for 19" MGX rack with busbar	755-0129-000
1696	Blank panel	7B32C0800-600-G
7	DC Busbar X86 Control Nodes for Mission Control	SYS-221GE-FNB-A-NC24B
Storage		
1	Home storage	AFF A90A
32	DDN AI400 X3 Chassis	
32	100GbE Ports	
128	400GbE	
130	1GbE OOB Ports	
2	1GbE Ports in In-band	
Compute Fabric		
64	SP4 (SN5600) - Compute Leafs - DC	920-9N42F-00RI-KC0
32	SP4 (SN5600) - Compute Spines - DC	920-9N42F-00RI-KC0
2048	Transceivers on DGX	980-9I51A-00NS00
2048	Twin port OSFP Transceivers on Leaf <-> DGX	980-9I510-F4NS00

continues on next page

Table 9.1 – continued from previous page

Count	Component	Recommended Model
4096	Twin port OSFP Transceivers on Leaf <-> Spine	980-9I510-F4NS00
4096	DGX-Leaf Cable	980-9I570-00N030
4096	Leaf-Spine Cable	980-9I570-00N030
Ethernet Storage Fabric		
8	SP4 (SN5600) - Storage Leafs DGX Side - DC	920-9N42F-00RI-KC0
2	SP4 (SN5600) - Storage Leafs Storage Side - DC	920-9N42F-00RI-KC0
4	SP4 (SN5600) - Storage Spines - DC	920-9N42F-00RI-KC0
512	400G Transceiver (for BF3240)	980-9I693-F4NS00
128	400G Transceiver (for DDN)	980-9I693-F4NS00
4	400G Transceiver (for Control Plane Nodes)	980-9I693-F4NS00
289	Twin port OSFP Transceivers on Leaf <-> DGX/DDN/Mgmt	980-9I510-F4NS00
576	Twin port OSFP Transceivers on Leaf <-> Spine	980-9I510-F4NS00
512	DGX-Leaf Cable	980-9I570-00N030
128	DDN, Mgmt - Leaf Cable	980-9I557-00N030
576	Leaf-Spine Cable	980-9I570-00N030
InfiniBand Storage Fabric		
16	NDR Switches, Stroage Leaf DGX Side	920-9B210-00RN-0M6
4	NDR Switches, Stroage Leaf Storage Side	920-9B210-00RN-0M6
8	NDR Switches, storage spines	920-9B210-00RN-0M6
512	400G Transceiver (for BF3240)	980-9I693-00NS00
128	400G Transceiver (for DDN)	980-9I693-00NS00
8	400G Transceiver (for Control Plane Nodes and UFM)	980-9I51S-00NS00
324	Twin port OSFP Transceivers on Leaf <-> DGX/DDN/Mgmt	980-9I510-00NS00
512	Twin port OSFP Transceivers on Leaf <-> Spine	980-9I510-00NS00
2	UFM Appliance	920-9B020-00RA-0D0
2	UFM <-> Inband Transceiver	980-9I693-F4NS00
648	Leaf Cable	980-9I570-00N030
512	Leaf-Spine Cable	980-9I570-00N030
Inband Fabric		
4	SP4 (SN5600) - Inband Leafs	920-9N42F-00RI-KC0
2	SP4 (SN5600) - Inband Spines	920-9N42F-00RI-KC0

continues on next page

Table 9.1 – continued from previous page

Count	Component	Recommended Model
512	400G Transceiver (for BF3240)	980-9I693-F4NS00
24	400G Transceiver (for Control Plane Nodes)	980-9I51S-F4NS00
134	Twin port OSFP Transceivers on Leaf <-> DGX, Management	980-9I510-F4NS00
32	Twin port OSFP Transceivers on Leaf <-> Spine	980-9I510-F4NS00
269	DGX, Mgmt - Leaf Cable	980-9I557-00N030
32	Leaf-Spine Cable	980-9I570-00N030
10	800G Transceivers (between leaf -> NFS & Uplink)	980-9I30H-F4NM00
20	DR4 -> 4x DR1 Fabric Splitter Leaf->NFS, SN2201, Uplink	NV08M31P2X1BHM020
4	100G DR1 (Uplink, NFS -> SP4)	
32	DR1 Storage	980-9I042-00C000
2	Inband <-> DDN Inband, using 25GbE port on SN5600, QSA 1/10G	980-9I71G-00J000
2	Inband <-> DDN Inband, using 25GbE port on SN5600, BaseT RJ45	980-9I251-00IS00
Out of Band Management		
33	SN2201	920-9N110-00R1-0C0
33	SN2201	920-9N110-00R1-NC0
15 38	Total RJ45 Count	Cat6
66	100G DR1 Transceiver (SN2201 @OOB)	980-9I042-00C000
18	800G optics (SN5600 Inband Spines -> SN2201)	980-9I30H-F4NM00

Chapter 10. Notices

10.1. Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

10.2. Trademarks

NVIDIA, the NVIDIA logo, NVIDIA DGX, and NVIDIA DGX SuperPOD are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

10.3. Copyright

Copyright