

NVIDIA DGX SuperPOD with IBM Storage Scale and IBM Storage Scale System 6000

NVIDIA DGX H100, H200, and B200 Storage Reference Architecture

Abstract

The [NVIDIA DGX SuperPOD™](#) with [NVIDIA DGX™ H200](#), or [NVIDIA DGX B200](#) systems is an artificial intelligence (AI) supercomputing infrastructure, which provides the computational power necessary to train today's state-of-the-art deep learning (DL) models and to fuel future innovation. The DGX SuperPOD delivers groundbreaking performance, deploys as a fully integrated system, and is designed to solve the world's most challenging computational problems.

This DGX SuperPOD storage reference architecture (RA) is the result of collaboration between DL scientists, application performance engineers, and system architects to build a system capable of supporting the widest range of DL workloads. The performance delivered by the DGX SuperPOD with DGX systems enables the rapid training of DL models at great scale. The integrated approach of provisioning, management, compute, networking, and fast storage enables a diverse, multi-tenant system that can span data analytics, model development, and AI inference.

In this paper, we will describe how to use the [IBM Storage Scale System 6000](#) to supporting DL workloads when connected to the DGX SuperPOD. IBM Storage Scale System 6000 is a storage appliance offering low latency NVMe physical storage, advanced erasure coding, and supports either InfiniBand or Ethernet networking. Multiple IBM Storage Scale System 6000's can be aggregated to create a high-performance cluster filesystem, or connected to multiple clusters for geographic and cross platform data sharing in a single global data platform. The Scale System 6000 is a 4U building block that makes it easy to deploy, manage, and grow fast storage for AI with NVIDIA DGX systems.

The DGX SuperPOD is a turnkey solution validated at scale with scale-out Scale System 6000s. Joint testing and integration ensure the NVIDIA DGX SuperPOD is a rapidly deployed and a robust solution for scalable AI development. NVIDIA and IBM jointly test, plan, and install the system, with the storage backed by IBM global deployment and support services.

As configured, tested, and deployed in the NVIDIA SuperPOD, the IBM Storage Scale System 6000 can be used for all DL workloads including:

- Training models efficiently with directly from IBM Storage Scale.
- Automatically leverage local resources as cache to minimizing rereading data across the network.
- Workspace for long-term storage (LTS) of datasets.
- A centralized repository for the acquisition, manipulation and sharing of results using standard protocols like NFS, SMB, and S3.

Table of Contents

Storage Overview	4
Storage Caching Hierarchy	4
Storage Performance Requirements	5
IBM Storage Scale System 6000 Overview.....	6
Planning	7
Hardware Overview	7
Multi-Protocol Access	8
Data Tiering and Caching.....	9
Integrated Lifecycle Management (ILM)	9
Active File Management (AFM).....	9
Scale System Management Server	10
Networking	10
Storage Fabric	11
IPMI/BMC Network	11
Management Network.....	11
Other Networking Considerations	12
Power	12
Sizing.....	12
Summary	13

Storage Overview

Training performance can be limited by the rate at which data can be read and reread from storage. The key to performance is the ability to read data multiple times, ideally from local storage. The closer the data is cached to the GPU, the faster it can be read. Storage must be designed considering the hierarchy of different storage technologies, either persistent or nonpersistent, to balance the needs of performance, capacity, and cost.

Storage Caching Hierarchy

The storage caching hierarchy of the DGX system is shown in Table 1. Depending on data size and performance needs, each tier of the hierarchy can be leveraged to maximize application performance.

Table 1. DGX system storage and caching hierarchy

Storage Hierarchy Level	Technology	Total Capacity ¹	Performance ¹
RAM	DDR4	2 TB per system ²	> 200 GB/s
Internal Storage	NVMe	30 TB per system ³	> 50 GB/s
<div>1. Total capacity and performance values are per system.</div> <div>2. Shared between the operating system, application, and other system processes</div> <div>3. PCIe NVMe SSD storage</div>			

Caching data in local RAM provides the best performance for reads. This caching is transparent after the data is read from the filesystem.

While local storage is fast, it is not practical to manage a dynamic environment with local disk alone in multi-node environment. Functionally, centralized storage can be as quick as local storage on many workloads.

Storage Performance Requirements

Performance requirements for high-speed storage greatly depend on the types of AI models and data formats to be used. The DGX SuperPOD has been designed as a capability-class system that can manage any workload both today and in the future. However, if systems are going to focus on a specific workload, such as natural language processing (NLP), it may be possible to better estimate performance needs of the storage system.

To enable customers to characterize their own performance requirements, some general guidance on common workloads and datasets is shown [Table 2](#).

Table 2. Characterizing different I/O workloads

Storage Performance Level Required	Example Workloads	Dataset Size
Good	NLP	Most all datasets fit in cache
Better	Image processing with compressed images, ImageNet/ResNet-50	Many to most datasets can fit within the local system's cache.
Best	Training with 1080p, 4K, or uncompressed images, offline inference, ETL	Datasets are too large to fit into cache, massive first epoch I/O requirements, workflows that only read the dataset once

Performance estimates for the storage system necessary to meet the guidelines in Table 2 are in:

- [Table 4](#) of the *NVIDIA DGX SuperPOD Reference Architecture—DGX H100 Systems*.
- [Table 6](#) of the *NVIDIA DGX SuperPOD Reference Architecture—DGX B200 Systems*.

Achieving these performance characteristics may require the use of optimized file formats such as TFRecord, RecordIO, or HDF5.

The IBM Storage Scale System 6000 can be used as shared storage to provide a single view of an organization's entire data set. It provides high peak system performance and high aggregate performance and meets the 'best' category requirements for both the H100 and B200 architectures.

IBM Storage Scale System 6000 Overview

The IBM Storage Scale System 6000 (Figure 1) combines the performance of NVMe storage technologies with the reliability and the rich features of IBM Storage Scale, along with several high-speed attachment options such as 400 Gb/s Ethernet and HDR InfiniBand (IB) — all in a powerful 4U storage system that scales out for performance and capacity.

Figure 1: IBM Storage Scale System 6000



IBM Storage Scale System on NVMe is designed to be the market leader in all-flash performance, and scalability with a bandwidth of 310 GB/s per NVMe all-flash appliance with low latency. Providing data-driven multicloud storage capacity, the NVMe all-flash appliance is deeply integrated with the software defined capabilities of IBM Storage Scale to seamlessly plug it into an analytics, scalable cluster, or AI workload.


Available with multiple drive options and advanced erasure coding, the Scale System 6000 provides options to optimize costs for different installation sizes. As with all IBM Storage Scale solutions, capacity and performance can be scaled. Combining Scale System 6000 systems provides nearly linear performance scalability. Scale System 6000 solutions may also be used as an all-flash NVMe performance tier combined with slower, more cost effective storage, including tape or object storage.

IBM Storage Scale is an industry leader in high-performance file systems. The underlying general parallel file system (GPFS) provides scalable throughput and low-latency data access, as well as superior metadata performance. Unlike other systems that can easily bottleneck, the distributed architecture of a parallel filesystem provides reliable performance for multi-user sequential and random read or write. This is particularly important in AI clusters where multiple compute nodes may need to read or write to the same file. IBM Storage Scale provides Container Native Access and Operators to support Kubernetes driven DevOps and Data Ops practices. In addition, IBM Storage Scale provides enterprise features such as call-home proactive support, encryption, and audit file logging that works with enterprise [security information and event management \(SEIM\)](#) platforms. IBM Storage Scale Systems integrate well with NVIDIA Base Command Manager to streamline administration and configuration of the solution.

Planning


Hardware Overview

The IBM Scale System 6000 is a fully redundant, centrally managed solution. The system is deployed in building blocks, with up to 48 U.2 form factor NVMe drives per building block.



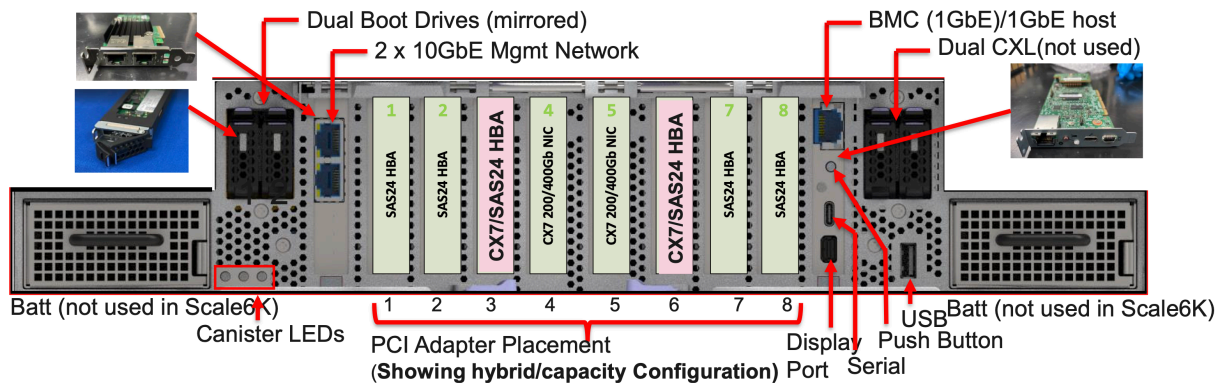
Front View

The rear of the system contains 4 power supplies and 2 redundant compute canisters



Rear View

The Scale System 6000 is a flexible solution that can be configured for either high capacity, using external SAS drive enclosures, or for performance. To meet the performance requirements for the NVIDIA SuperPOD, it is recommended to use the 'performance' configuration, populating 4 PCI gen5 slots per canister (slots 3-6 in the figure below) with NVIDIA ConnectX7 adapters. The Scale System 6000 should also be configured with NVME drives as outlined in the sizing section below.



The Scale System 6000 can be configured as a part of a solution to meet performance, capacity, and cost requirements for any need.

Multi-Protocol Access

An IBM Scale System 6000 utilizes a high-speed, proprietary protocol to provide access to data. This protocol provides high speed, consistent, redundant access to data. To access data, a client requires special software to provide access to the data. NVIDIA DGX systems will have the Storage Scale client installed for high speed access.

For user home directories and some cluster specific data, the NVIDIA DGX SuperPOD architecture recommends [NFS for data access on the management servers](#). To meet this need, a minimum of 2 storage scale protocol nodes are required. Each protocol node provides access to the file system via NFS, and optionally supports SMB, HDFS, and Object storage. Two nodes provide redundancy of access, and up to 32 protocol nodes can be added to the system. Additional protocol nodes can provide higher speeds, or can support clients external to the SuperPOD for easier data access and movement.

Each protocol node can run either on a Storage Scale Utility node or on standard x86 or IBM Power system running a standard OS such as RHEL or Ubuntu. See the Storage Scale FAQ for the latest OS's and releases supported.

Each protocol node is connected to the Storage Fabric with 2 200 Gb NDR adapters. These adapters can be connected via a splitter cable to the 400 Gb Storage Fabric for greater port density and to reduce the number of Infiniband ports required for the solution. In addition, protocol nodes will connect to the

Management Network for NFS access and can optionally connect to additional networks to support clients outside of the DGX SuperPOD.

Data Tiering and Caching

IBM Storage Scale offers both the ability to tier data within the file system, or to cache data from external systems.

The Integrated Lifecycle Management (ILM) functionality of Storage Scale moves data seamlessly between various storage mediums such as NVME, hard drives, and tape drives. By placing data on the appropriate storage type, Storage Scale allows for high-speed access to data while offering cost-effective capacity expansion.

The Active File Management (AFM) function caches storage from external sources such as Object, NFS, or other Storage Scale file systems. By caching storage on local storage, users are given high-speed local access to data even if the source copy resides on external storage.

Integrated Lifecycle Management (ILM)

The IBM Storage Scale ILM functionality combines multiple storage tiers, or pools, such as NVME, disk, or tape, into a single namespace. Data can be moved between the storage tiers seamlessly to an end user at any time. In addition, a robust policy syntax allows for automatic movement of data in certain conditions – for example once the hard disk pool reaches a certain capacity, automatically migrate the least recently used data to tape.

IBM Storage Scale 6000 systems can be configured with optional SAS adapters and external spinning disk enclosures in order to extend capacity. Storage Archive Enterprise Edition (<https://www.ibm.com/docs/en/storage-archive-ee/1.3.4?topic=overview-introduction-storage-archive-enterprise-edition>) can be used to connect to external tape enclosures using additional nodes.

Active File Management (AFM)

IBM Storage Scale AFM seamlessly caches data from external data stores. These data stores can be remote Storage Scale, NFS, or Object stores.

To provide connectivity, a Storage Scale cluster requires a minimum of 1 AFM gateway node, however at least 2 are needed for redundancy. The hardware for the gateway nodes can be a Storage Scale Utility node, or standard x86/Power hardware. Please see the IBM Storage Scale FAQ for the latest guidelines.

AFM gateway nodes require 2 connections to the SuperPOD data network, with NVIDIA ConnectX-7 400 Gb IB connectivity being recommended. In addition, each gateway node requires at least one connection to a high speed network that can connect to the storage to cache. The external network may vary depending on the speed and latency to the external storage source.

When sizing the number of gateway nodes and network connectivity, 3 factors should be considered:

1. The size of the active data set being cached
2. The number of files/objects being cached
3. The change rate of the data that is cached

Depending on these factors, additional gateway nodes may be needed in to meet the given workload. The IBM Storage Scale FAQ and the planning section of the knowledge center (<https://www.ibm.com/docs/en/storage-scale/5.2.1?topic=planning-afm>) provide additional guidance on gateway node configuration.

Scale System Management Server

The IBM Scale System 6000 requires a management server (EMS) to provision and manage the storage system. Each EMS can manage multiple Scale System 6000 systems in a single cluster. Typically, the management server is deployed on a dedicated 2U Scale System utility node, with 512 GB of RAM and 1 ConnectX-7 dual-port VPI adapter. The management server requires one connection from the ConnectX-7 to the SuperPOD storage fabric.



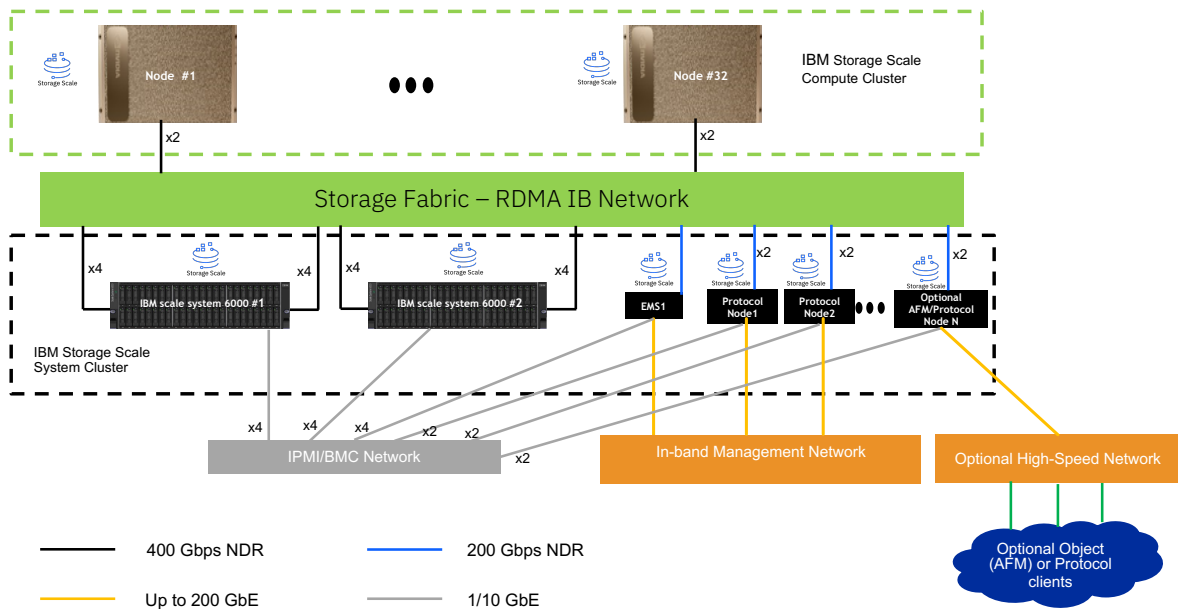
Front View of the Scale System Utility Node

In addition, a dedicated management switch from the management server is required, to provision and manage the system. This switch may be customer-owned; however an IBM switch is recommended as it is pre-configured with the appropriate configuration to manage the Scale System 6000.

For all configurations in the sizing section, a single management server and switch meeting these requirements will be necessary and is included in the rack units described.

Networking

The IBM Scale System 6000 uses multiple networks described below.



Sample architecture with 2 Scale System 6000 and 1 SU

Storage Fabric

The Storage Fabric connects the Scale System 6000 systems to the DGX Superpod. In an NVIDIA SuperPOD solution, it is recommended to use 4 400 Gb Infiniband adapters per canister, for a total of 8 adapters per Scale System 6000. These adapters should directly be connected to MQM9700 switches. The storage network can be scaled as required to meet any requirements for DGX compute systems, as well as storage capacity and performance. The standard SuperPOD spine-and-leaf configuration is recommended, as described in the NVIDIA SuperPOD [reference architecture](#).

In addition to the Infiniband connectivity for the Scale System 6000, the required EMS and a minimum of 2 protocol nodes require connectivity to the Storage Fabric. The EMS requires a single 200 Gbps NDR connection, while the protocol nodes should have 2 200 Gb connections each.

IPMI/BMC Network

The Scale System 6000 requires a dedicated management network to manage the systems, perform software upgrades, monitor, and configure the systems. While any network can be used for this task, it is recommended to use an IBM provided switch, which will be pre-configured and managed by the system.

Management Network

To manage the system, one connection from the management node to the management network is


required. This connection provides management access and can optionally be configured for call-home.

Other Networking Considerations

Optional features described earlier in this document such as multi-protocol access or AFM may require additional networking capabilities to connect to external storage systems, object storage, or clients. The quantity and speed of these connections will vary depending on existing infrastructure, the size of the solution and other factors.

Power

Each Scale System 6000 enclosure contains 4 redundant power supplies. The power supplies are designed to provide redundant power to each canister in the system.



In a 48 NVME drive configuration, the following power measurements represent the maximum draw for the system

Product	kVA	Amps	Inlet	Watts	Input Power
Scale System 6000 48 NVME drive	4.50	22.5	C20 (x4)	4800*	200 V to 240 V single phase 50 Hz or 60 Hz 13 A (x4)

* This value represents the absolute maximum power draw. Actual usage is affected by system load, ambient temperature, and several other factors

Refer to the Scale System 6000 Hardware Planning Guide (https://www.ibm.com/docs/en/5149-F48/pdf/ess6000_hwg.pdf) for additional information regarding system power, cooling, and other considerations.

Sizing

The IBM Storage Scale System 6000 has a flexible architecture and can scale to meet nearly any performance or capacity requirement. Each 4U Scale System 6000 enclosure can support up to 48 NVME devices and offers up to 310 GB/sec seq. read speeds. Both capacity and performance can be scaled

nearly linearly by adding more Scale System 6000 enclosures.

In addition to NVME drives, each Scale System 6000 system can support up to 9 SAS-attached enclosures to expand capacity with low-cost storage. The Scale System 6000's Integrated Lifecycle Management (ILM) functionality can seamlessly move data between the high performing NVME devices and slower drives to scale capacity in a cost-effective way. This reference only contains NVME solutions, however mixed NVME and SAS solutions can be deployed to meet additional capacity requirements.

The following table contains guidance from NVIDIA regarding the storage performance required for many AI workloads.

Performance Characteristic	Good (GBps)	Better (GBps)	Best (GBps)
Single-node read	4	8	40
Single-node write	2	4	20
Single SU aggregate system read	15	40	125
Single SU aggregate system write	7	20	62
4 SU aggregate system read	60	160	500
4 SU aggregate system write	30	80	250

(Source: <https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-b200/latest/storage-architecture.html#>)

The following table shows how the Scale System 6000 can be scaled in order to meet the 'best' requirement depending on the number of scalable units deployed. The rack units and data network (400 Gb Infiniband) connections required include the necessary management node, 2 protocol nodes, and internal switch for the configuration. These systems do not account for optional data movement via AFM, or additional protocol nodes.

SU	Scale System 6000	Read GB/s	Write GB/s	Rack Units	Storage Fabric Ports	Usable Capacity ¹		
						7.68 TB	15.36 TB	30.74 TB
1	1 ²	310	155	11	12 ²	0.25 PB	0.5 PB	1 PB
2	1	310	155	11	12	0.25 PB	0.5 PB	1 PB
4	2	620	310	15	20	0.5 PB	1 PB	2 PB

1. Usable capacity is an approximation assuming 8+2P and standard spare space. Actual capacity may vary
2. For the single SU configuration, it is possible to only ½ of the NVME and ConnectX-7 adapters (with reduced performance and capacity) and still meet the best category. Fully populating the system will allow for more flexible expandability to meet future needs. The guide assumes a fully populated system.

Summary

NVIDIA evaluations show that the IBM Storage Scale System 6000 with Storage Scale file system meets the 'best' DGX SuperPOD performance and functionality requirements. It is a great choice to pair with a

DGX SuperPOD to meet current and future storage needs.

The Scale System 6000 can tier data to hard disk, tape, and object storage to deliver a cost effective solution. The robust integrated lifecycle management (ILM) engine automatically moves data to the appropriate storage type to deliver high performance while moving unused data to a more cost-effective form of storage. In addition, global file sharing using the active file management (AFM) technologies allows for an organization to seamlessly share data across the world.

As storage requirements grow, IBM Scale System 6000 building blocks can be added to seamlessly scale capacity, performance, and capability. The combination of NVME hardware and IBM Storage Scale parallel file system architecture provides excellent random read performance, often just as fast as local storage for sequential read patterns. Testing has validated that each IBM Scale System 6000 can deliver the highest levels of per node performance and meet all our application performance requirements. The IBM Storage Scale parallel file system provides a platform that is fully supported with the NVIDIA DGX SuperPOD and has been deployed in production.