

# NVIDIA DGX BasePOD with IBM Storage Scale and IBM Storage Scale System 6000

NVIDIA DGX H100, H200, and B200 Storage Reference Architecture

# Abstract

[NVIDIA DGX BasePOD™](#) 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 fuel future innovation.

In this paper, we will describe how to use the [IBM Storage Scale System 6000](#) to support DL workloads when connected to NVIDIA DGX BasePOD. IBM Storage Scale System 6000 is a storage appliance offering low latency NVMe physical storage, advanced erasure coding, and supports either [NVIDIA Quantum-2™ InfiniBand](#) or [NVIDIA Spectrum™ 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 storage system that makes it easy to deploy, manage, and grow fast storage for AI with NVIDIA DGX systems.

NVIDIA DGX BasePOD with scale-out Scale System 6000s has been rigorously tested and validated with NVIDIA to ensure a robust solution for scalable AI development.

As configured, tested, and deployed in NVIDIA DGX BasePOD, the IBM Storage Scale System 6000 can be used for all ML workloads and includes:

- Efficient training and checkpointing of AI models with data directly accessed from IBM Storage Scale.
- Automatic caching of local resources to minimize rereading of 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

<b><i>IBM Storage Scale System 6000 Overview</i></b>	<b>3</b>
<b>Multi-Protocol Access</b>	<b>4</b>
<b>Data Tiering and Caching</b>	<b>4</b>
Integrated Lifecycle Management (ILM)	4
Active File Management (AFM)	5
<b>Scale System Management Server</b>	<b>7</b>
<b><i>Design</i></b>	<b>8</b>
<b>Compute</b>	<b>8</b>
<b>Storage</b>	<b>8</b>
<b>Networking</b>	<b>9</b>
Compute Fabric	9
Storage Fabric	9
In-Band Management Fabric	9
Out-of-Band Management Fabric	9
<b>Cabling</b>	<b>10</b>
<b>Power</b>	<b>11</b>
<b><i>Performance</i></b>	<b>12</b>
<b><i>Summary</i></b>	<b>13</b>

# 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 NVIDIA Spectrum 400 Gb/s Ethernet and NVIDIA Quantum-2 400 Gb/s InfiniBand — 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 connect to compute clusters supporting analytics or AI workloads.

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 excellent performance scalability. Scale System 6000 solutions may also be used as an all-flash NVMe performance tier combined with higher latency, 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 IBM Storage Scale's 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

(SIEM) platforms. IBM Storage Scale Systems integrate with [NVIDIA software](#) such as [NVIDIA Base Command™ Manager](#) to streamline administration and configuration of the entire solution.

## Multi-Protocol Access

An IBM Scale System 6000 utilizes a high-speed, proprietary protocol to provide access to data. This protocol provides parallel, consistent, and redundant access to data concurrently from multiple systems. To access data using this protocol, clients require special software to be installed to provide access to the data. In the reference architecture, NVIDIA DGX systems require the Storage Scale client to be installed for high-speed access.

For external access to data stored on the Scale System 6000 by other users, the Storage Scale client can be used, or the solution can be configured with optional protocol nodes to support NFS, SMB, HDFS, and low-latency S3 object access to data. This allows external users to access the data using standard protocols, and to read, write, or view data directly.

To take advantage of multi-protocol access, a minimum of 2, and up to 32 protocol nodes can be installed depending on the number of users, speed of access required, and protocols used. These nodes can either be Storage Scale Utility nodes, or any standard x86 or IBM Power system running RHEL or Ubuntu Linux. See the [IBM Storage Scale FAQ](#) for the latest OS's and releases supported.

## 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, IBM 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, the least recently used data is automatically migrated to lower cost, higher latency storage.

To extend capacity with lower-cost storage, IBM Storage Scale 6000 systems can be configured with optional SAS adapters and spinning disks. In addition, products such as [IBM Storage Archive Enterprise Edition](#) can be used to connect to external tape enclosures using additional nodes.


ILM can be used to extend the capacity in a single namespace. However, any data that is actively being used by NVIDIA DGX systems or GPUs should reside on NVMe storage for optimal performance.

## 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. This capability allows for the file system to be extended beyond the Scale System 6000 to public or private clouds using the S3 protocol, to additional IBM Storage Scale clusters, or to other NFS-compliant storage systems.

The ability to tier data to S3-compliant clouds allows applications to create and modify data directly. That data is then cached to the Scale System 6000, which provides high-speed access to NVIDIA DGX BasePOD for high-speed training or other jobs. Deep-learning or AI algorithms then have the benefit of training with the most recent and up-to-date data automatically.

To provide connectivity, a Storage Scale cluster requires a minimum of one AFM gateway node, however at least two 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. The following diagram contains an example of how the system is connected.



Each AFM gateway node requires 2 connections to the storage fabric, with NVIDIA ConnectX®-7 400 Gb/s InfiniBand connectivity recommended. In addition, each gateway node requires at least one connection to a high-speed external 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:

- The size of the active data set being cached
- The number of files/objects being cached
- The change rate of the data that is cached

Depending on these factors, additional gateway nodes may be needed to meet the given workload. The [IBM Storage Scale FAQ](#) and the [planning section](#) of the knowledge center 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 1 NVIDIA ConnectX®-7 dual-port adapter, providing 400 Gb/s InfiniBand or 400 Gb/s Ethernet connectivity. The management server requires one connection from the ConnectX-7 to the storage fabric.



Front View of the Scale System Utility Node

In addition, a dedicated management switch from the management server is recommended, 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.



## Design

### Compute

IBM and NVIDIA have worked to ensure that the IBM Storage Scale System 6000 meets the performance requirements for NVIDIA DGX systems. The IBM Storage Scale System 6000 has been qualified for the following DGX models.

Supported NVIDIA DGX Models			
DGX A100,	DGX H100,	DGX H200,	DGX B200

### Storage

The IBM Storage Scale system offers flexible deployment, supporting both NVIDIA Quantum InfiniBand and Spectrum Ethernet networks. In addition to gaining capacity, performance scales when adding additional Storage Scale 6000 building blocks to the system. In addition to the IBM Storage Scale 6000 building blocks, 1 Scale System utility node is required to manage the solution. A single utility node can manage multiple Storage Scale 6000 systems.

A dedicated management switch for the Storage Scale 6000 and management node can be ordered from IBM and is preconfigured for ease of use and deployment. However, the system can use existing 1 Gb switches in an out-of-band management network if they are configured appropriately. This configuration may require configuring VLANs on the switches to separate traffic as needed.

When configuring the storage, it is recommended to have at least 1 GB/sec read capacity per GPU installed in the DGX BasePOD. Some workloads may benefit from additional throughput or additional capacity. The following table shows the required building blocks for various DGX BasePOD sizes.

DGX Systems	1	4	8
GPUs	8	32	64
IBM Storage Scale 6000	1*	1*	1

\* A ½ populated system Scale System 6000 with 24 NVMe drives, can still meet the requirements of many workloads

## Networking

The IBM Storage Scale 6000 reference architecture recommends splitting the management and storage fabric into 2 separate networks. Therefore, 4 networks are recommended:

- Compute Fabric
- Storage Fabric
- In-band management Fabric
- Out-of-band management fabric

### Compute Fabric

**Purpose:** Provides high-speed, low latency communications between GPUs.

**Recommendation:** NVIDIA Quantum-2 400 Gb/s InfiniBand

**Hardware:** NVIDIA Quantum-2 QM9700 InfiniBand switches

**Connections:** 4 400 Gb/s OSFP ports from each DGX system

### Storage Fabric

**Purpose:** Provides high-speed communication between the IBM Storage Scale 6000, DGX compute nodes, and management node infrastructure.

**Recommendation:** NVIDIA Quantum-2 400 Gb/s InfiniBand. The IBM Storage Scale also supports NVIDIA Spectrum Ethernet, with and without RDMA over Converged Ethernet (RoCE). InfiniBand or RoCE are required for [NVIDIA GPUDirect](#) communication

**Recommended Hardware:** NVIDIA Quantum-2 QM9700 InfiniBand switches

**Connections:** 8 ports from each IBM Storage Scale 6000, 2 ports from each DGX, 2 ports from management node, 1 port from Storage Scale management node

### In-Band Management Fabric

**Purpose:** Cluster management, scheduling, and provisioning of the NVIDIA DGX BasePOD.

**Recommendation:** 100 Gb/s Ethernet

**Recommended Hardware:** NVIDIA Spectrum SN4600C Ethernet switch

**Connections:** 2 ports from each DGX, 2 port for each NVIDIA DGX BasePOD management node

### Out-of-Band Management Fabric

**Purpose:** Management and monitoring of the NVIDIA DGX BasePOD. Optionally can be used to manage and provision the IBM Storage Scale 6000 if the IBM-provided switch is not used.


**Recommendation:** 1 Gb Ethernet

**Recommended Hardware:** NVIDIA Spectrum SN2201 switch

**Connections:** 1 port from each DGX, 1 port for each NVIDIA DGX BasePOD management node, optionally 4 ports from each IBM Storage Scale 6000, 3 ports from the IBM Storage Scale management node


## Cabling

The following diagram shows a sample validated network configuration for an IBM Storage Scale 6000 and 4 NVIDIA DGX B200 systems. The compute and storage fabrics are shown.



## Power

Each IBM 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](#) for additional information regarding system power, cooling, and other considerations.

## Performance

The IBM Storage Scale System 6000 provides excellent performance to meet AI training needs. As models increase in complexity, the Scale System 6000's superior write bandwidth allows for efficient checkpointing, allowing GPUs to spend more time training and less time waiting for data.

For a standard NVIDIA DGX BasePOD, a single Scale System 6000 can meet most performance needs as shown below:

<b>DGX Systems</b>	1	4	8
<b>GPUs</b>	8	32	64
<b>Read Performance (GB/sec)</b>	32	122	219
<b>Read performance per GPU (GB/sec)</b>	4	3.8	3.4
<b>Write Performance (GB/sec)</b>	23	92	144
<b>Write Performance per GPU (GB/sec)</b>	2.9	2.9	2.3

Due to the Scale System 6000's scalability, adding additional units can easily increase performance to meet nearly any requirement.

## Summary

NVIDIA and IBM have collaborated to ensure that the IBM Storage Scale System 6000 meets the requirements for NVIDIA DGX BasePOD deployments. The IBM Storage Scale System 6000 has been rigorously tested and validated by NVIDIA and IBM to ensure a seamless experience when paired with NVIDIA DGX systems.

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 nearly any application performance requirement. The IBM Storage Scale parallel file system provides a platform that is fully supported with the NVIDIA DGX BasePOD and has been deployed in production.