

IBM Storage Scale System 6000

Accelerated infrastructure for artificial intelligence, high-performance computing, analytics, and hybrid cloud

Highlights

Integrated hardware and software for building a global data platform

NVIDIA GPUDirect support to accelerate AI training and inferencing performance

Provides distributed file and object storage for the most demanding AI, HPC, analytics, and hybrid cloud workloads

Up to 310GB/s throughput; up to 13M IOPS using NVMeoF, up to 3.4PBe (effective) in a 4U rack

The business world is rapidly being transformed by artificial intelligence (AI), high-performance computing (HPC), analytics, and hybrid cloud technologies. But unlike traditional applications that work with *structured* data stored in databases, these new workloads and data lakes operate on a vast ocean of *unstructured* data, such as documents, audio, images, and videos.

As a result, organizations are having to re-think their data storage strategies and adapt to new ways of working, including how to leverage AI to unlock the value of their data – wherever it might reside.

IT leaders face many new challenges including:

- Accessing and analyzing data and workloads scattered across the globe;
- The increasing time needed by AI training and inferencing workloads;
- The cost and scarcity of resources, especially NVIDIA graphic processing units (GPUs).

Addressing these challenges requires specialized software and hardware:

- **IBM Storage Scale** is software-defined file and object storage for both structured and unstructured data;
- **IBM Storage Scale System 6000** is a hardware implementation of Storage Scale software and is optimized for the most demanding AI, HPC, analytics, and hybrid cloud workloads.
- **IBM Storage Scale System 3500** is for customers requiring an enterprise-ready entry-level or mid-level system.

IBM Storage Scale

The unstructured and semi-structured data from AI workloads, advanced analytics, data lakes, and other data-intensive apps must be stored in distributed file and object systems to make it accessible to geographically dispersed applications, services, and devices.

IBM Storage Scale software is designed to address these requirements with global data abstraction services that provide connectivity from multiple data sources and multiple locations to bring together data wherever it lives, including non-IBM storage environments.



Figure 1 – The IBM Storage Scale System 6000 can deliver up to 310GB/s throughput, up to 13M IOPS, and up to 3.4PBe (effective capacity) in a 4U rack.

Storage Scale achieves this with active file management (AFM), which abstracts other storage systems, including customers’ existing on-prem or public cloud storage, and provides a transparent caching layer to improve storage performance. The data from the different storage systems is presented in a single global namespace and can be simultaneously accessed via multiple protocols.

Storage Scale is based on a massively parallel file system and can be deployed on multiple hardware platforms including x86, IBM Power, IBM zSystem mainframes, ARM-based POSIX client, virtual machines, and Kubernetes.

By delivering high-performance access to customers’ entire unstructured data estate regardless of protocol, location, or format, including those on their existing storage, Storage Scale enable customers to accelerate their AI pipeline and unlock the value of their structured and unstructured data.

IBM Storage Scale System 6000

IBM Storage Scale System 6000 is a hardware platform that’s designed to be the fastest and most flexible way for organizations to build a global data platform around their file and object data. It leverages the power of IBM Storage Scale software combined with NVMe flash technology to deliver high-performance storage for AI, data analytics and high-performance file and object use cases.

Storage Scale System 6000 uses a simple building-block approach with performance that scales linearly, i.e.: a cluster of 10 Storage Scale System 6000 systems is capable of more than 3 terabytes per second throughput. It supports up to nine SAS hard disk drive expansion enclosures.

Key features of the Storage Scale System 6000 include:

- A single 4U node with active-active controllers and redundant hardware to maximize uptime;
- Up to 310 gigabytes per second (GB/S) throughput with low latency;
- Up to 13 millions IOPS using NVMeoF;
- Up to 3.4PBe (effective capacity) in a standard 4U rack space;
- Supports up to 48 3.84TB, 7.68TB, 15.36TB, or 30TB 2.5” NVMe flash drives;
- Supports 19.2TB and 38.4TB FlashCore Module 4 NVMe drives.

“We fundamentally believe that core to the competitiveness of every company going forward will be their ability to use AI to unlock real-time value from their data wherever the data resides.”

Arvind Krishna – CEO, IBM

310

Up to 310 gigabytes per second throughput

3.4PBe

Up to 3.4PBe (effective capacity) in a standard 4U rack space

To optimize performance, Storage Scale System 6000 is engineered to scale to thousands of nodes and yottabytes of capacity. Storage Scale System 6000 runs IBM Storage Scale RAID erasure coding, which provides data efficiency, mitigation of storage hardware failures, and intelligent monitoring and dynamic data tuning. Installations and updates are delivered by means of containerized software that speeds and simplifies the maintenance process.

For data resilience, Storage Scale System 6000 supports Safeguarded Copy, a protection mechanism that supports the ability to create cyber-resilient, point-in-time copies of volumes of data. Safeguarded Copy creates isolated immutable snapshots of data on a regular schedule so that operational data can be rapidly recovered in the event of a cyber-attack or other potential data loss event.

Accelerating AI Workloads

Distributed file and object storage systems are well suited to AI and analytics workloads. But as AI and analytics datasets continue to increase in size, the time spent loading data begins to impact application performance. To unlock the full potential of AI and ensure that fast GPUs aren't being starved by slow IO, organizations may need to improve their existing IT infrastructure.

To accelerate the processing of globally distributed data, Storage Scale System 6000 supports the NVIDIA GPUdirect Storage protocol, which enables a direct data path between GPU memory and local or remote storage, such as NVMe or NVMe over Fabric (NVMe-oF). This GPUdirect architecture removes the host server CPU and DRAM from the data path, so the IO path between storage and the GPU is shorter and faster.

IBM FlashCore Modules

As an alternative to the current industry-standard NVMe flash drives, Storage Scale System 6000 can now be configured with 19.2TB or 38.4TB IBM FlashCore Modules for increased data density and data resilience.

FlashCore Modules (FCMs) are IBM's patented high-performance NVMe flash drives, designed for low latency, density, and reliability. They feature a PCIe Gen4 U.2 interface and high-speed NAND memory to provide high throughput and IOPS with consistent and predictable latency.

IBM researchers have developed custom field programmable gate arrays (FPGAs) that free up computational overhead so that compression and encryption can be performed with zero effect on performance. FlashCore Modules also implement advanced flash management, which improves flash endurance over standard implementations without sacrificing latency.

IBM Storage System Scale Expansion Enclosure

The IBM Storage Scale Expansion Enclosure is an enterprise-class, fully redundant storage expansion enclosure optimized for the Storage Scale System 6000. Each enclosure can contain up to 91 20TB or 22TB self-encrypting SAS hard disk drives (HDDs).

Up to nine enclosures can be attached to the Storage Scale System 6000, delivering up to 18PB of HDD storage capacity per rack using 24Gb SAS. The ability to add multiple expansion enclosures enables the Storage Scale System 6000 to address a wide new range of workloads that operate on multi-petabyte data sets.




Figure 2 – The IBM Storage Scale System Expansion Enclosure enables organizations to cost-effectively deploy workloads operating on massive data sets.

Specifications – IBM Storage Scale System 6000

System Features	<ul style="list-style-type: none"> ▪ Active-active dual controller storage system: Dual socket AMD EPYC™ Genoa CPUs (48 cores) ▪ 768 GB to 1536 GB memory per canister (system total is either 1536 GB or 3072 GB) ▪ De-clustered RAID supporting erasure coding schemas: 3-way and 4-way replication, 4+2P, 4+3P, 8+2P, 8+3P ▪ Capacity 96TB to 5.4PBe¹
Software	<ul style="list-style-type: none"> ▪ IBM Storage Scale for Storage Scale System software ▪ Red Hat® Enterprise Linux® (RHEL)
Software Features	<ul style="list-style-type: none"> ▪ Data access services with POSIX, NVIDIA GPUDirect, Container Native Storage Access (CNSA), CSI, HDFS, NFS v4, SMB, HTTP, S3 ▪ Data abstraction services including cloud and non-IBM storage and multi-site asynchronous and synchronous replication ▪ Data management services with integrated lifecycle management to optimize data from memory, NVMe flash, HDD, public cloud, external storage and tape ▪ Data resilience services with FIPS 140-2, 256-bit encryption, cyber-protect and IBM Safeguarded Copy
Performance	<ul style="list-style-type: none"> ▪ Dual socket AMD EPYC™ Genoa ▪ Sequential read performance up to 310GB/s and 13M IOPs²
Networking / Adapters	<ul style="list-style-type: none"> ▪ NDR / HDR switch support ▪ 16 x PCIe Gen5 slots per system (8 slots per canister) ▪ Up to 12 24G SAS adapters (6 per canister) ▪ NVIDIA CX7 supported adapters: <ul style="list-style-type: none"> ▪ 400Gb single port (InfiniBand only) x16 Gen5 ▪ 200Gb VPI dual port (InfiniBand/Ethernet) x16 Gen5 ▪ 8 Max High Speed Network Adapters (CX7) ▪ Up to 16 network ports
Drive Support	<ul style="list-style-type: none"> ▪ 24 or 48 U.2 NVMe SSDs (30.72TB) ▪ 24 or 48 U.2 NVMe FCM (19.2TB or 38.4TB)
Environmental – Drives	<ul style="list-style-type: none"> ▪ Nominal power: 20-25 watts per SSD ▪ Nominal power: 30 watts per FCM
Environmental – Controller	<ul style="list-style-type: none"> ▪ Input voltage: 200-240V 50/60 Hz ▪ Nominal power: 1,500 W (empty); 4800 W (PSU max) ▪ Power supplies: 4 hot swappable, redundant
Size – Controller	<ul style="list-style-type: none"> ▪ 4RU; H:7" (175 mm), W: 19" rack (483 mm), D: 33.5" (850 mm) without bezel
Size - Enclosure	<ul style="list-style-type: none"> ▪ 4RU; H:7" (175 mm), W: 17.5" (445 mm), W: 19.0" rack (483 mm), D: 39" (1000 mm), L: 36" (926 mm) chassis length

¹ Effective capacity based on 3:1 in-line compression; compressibility is dependent on the characteristics of the data stored on each FCM drive.

² Disclaimer: Performance metrics were obtained by running sample programs in a controlled environment with standard hardware/software configurations and testing procedures. Since performance varies with configuration, program characteristics, and other installation and environment factors, results obtained in other operating environments may vary. IBM® does not represent, warrant, or guarantee that a user will achieve the same or similar results in the user's environment.

IBM Storage Scale System 3500

IBM Storage Scale System 3500 is for customers requiring an entry-level or mid-level system for less extreme product environments. It has many of the same capabilities as the Storage Scale System 6000 and is also capable of high capacity and performance.

To learn more, download the Storage Scale System 3500 [data sheet](#).

Parameter	Storage Scale System 3500	Storage Scale System 6000
Size	2 rack units (2U)	4 rack units (4U)
Maximum capacity	24 x 30.72TB NVMe	48 x 30.72TB NVMe 48 x 38TB FCM4 modules
Maximum throughput	126GB/s	310GB/s
Expansion	Up to 4 direct-attached JBODs	Up to 9 direct-attached JBODs
Data transfer	12Gb SAS	24Gb SAS

For more information

To learn more about IBM Storage Scale System, contact your IBM representative or IBM Business Partner, or visit ibm.com/products/storage-scale-system.

© Copyright IBM Corporation 2024
IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
May 2024

IBM, the IBM logo, and zSystems are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

