

IBM Storage Scale System 3500

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



Highlights

Integrated hardware and software for building a global data platform

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

Up to 126GB/s throughput;
up to 1.2M IOPS;
up to 500TB usable in a standard 2U rack space

NVIDIA GPUdirect support to accelerate AI training and inferencing performance

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 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 3500** is a hardware implementation of Storage Scale software designed for AI, HPC, analytics, and hybrid cloud workloads;
- **IBM Storage Scale System 6000** builds on the Storage Scale System 3500 with increased capacity and performance capabilities designed to address the requirements of the most demanding workloads.

IBM Storage Scale

The unstructured and semi-structured data from AI workloads, analytics, data lakes, and other data-intensive apps must be stored in a distributed file and object storage system 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 3500 can deliver up to 126GB/s throughput, up to 1.2M IOPS, and up to 500TB usable in a standard 2U rack space.

“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

IBM Storage Scale

The unstructured and semi-structured data from AI workloads, analytics, data lakes, and other data-intensive apps must be stored in a distributed file and object storage system 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.

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 3500

IBM Storage Scale System 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.

IBM Storage Scale System 3500 uses a simple building-block approach with performance that scales linearly, i.e.: a cluster of 10 Storage Scale System 3500 systems is capable of more than a terabyte per second throughput. It supports up to eight SAS hard disk drive expansion enclosures.

126

Up to 126 gigabytes per second throughput

500

Up to 500 terabytes capacity in a standard 2U rack space

Key features of the Storage Scale System 3500 include:

- A single 2U node with active-active controllers and redundant hardware to maximize uptime;
- Up to 126 gigabytes per second (GB/s) throughput with low latency;
- Up to 1.2 millions IOPS;
- Up to 500TB usable in a standard 2U rack space
- Supports up to 48 3.84TB, 7.68TB, 15.36TB, or 30TB 2.5" NVMe flash drives

To optimize performance, Storage Scale System 3500 is engineered to scale to thousands of nodes and yottabytes of capacity. Storage Scale System 3500 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 3500 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 3500 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 Storage Scale HDD Disk Enclosure

The IBM Storage Scale HDD Disk Enclosure is an enterprise-class, fully redundant storage expansion enclosure optimized for the Storage Scale System 3500. Each enclosure can contain up to 102 self-encrypting SAS hard disk drives (HDDs).

Up to eight disk enclosures can be attached to the Storage Scale System 3500, delivering up to 17.95PB of HDD storage capacity per rack using 12Gb SAS. The ability to add multiple expansion enclosures enables the Storage Scale System 3500 to address a wide new range of workloads that operate on multi-petabyte data sets.

Storage Scale System 6000

Designed for workloads with the most demanding performance and capacity requirements, IBM Storage Scale System 6000 has many of the same capabilities as the Storage Scale System 3500 but with enhanced storage and compute hardware.

For additional information, download the Storage Scale System 6000 [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 8 direct-attached JBODs | Up to 9 direct-attached JBODs |
| Data transfer | 12Gb SAS | 24Gb SAS |

Specifications – IBM Storage Scale System 3500

| | |
|-----------------------------------|--|
| System Features | <ul style="list-style-type: none"> ▪ Active-active dual 1-socket storage controllers ▪ 1024 GB memory ▪ De-Clustered RAID supporting erasure coding schemas: 3-way replication, 4-way replication ▪ Capacity – Up to 500TB usable |
| 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"> ▪ AMD 7642 48 core processor ▪ Sequential read performance up to 126GB/s¹ |
| Networking / Adapters | <ul style="list-style-type: none"> ▪ HDR 200Gb/s InfiniBand, up to 8 ports ▪ 100Gb/s Ethernet, up to 8 ports ▪ Two integrated networking ports (separation of MGMT and BMC) ▪ Four x16 PCIe Gen4 adapter slots |
| Drive Support | <ul style="list-style-type: none"> ▪ 12 or 24 NVMe SSDs (3.84TB, 7.68TB, 15.36TB, or 30.72TB) ▪ Up to 816 HDD drives (10TB, 14TB, 18TB, 20TB, or 22TB) |
| 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: 1,350 W (empty); 2000 W (max) ▪ Power supplies: 2 hot swappable, redundant |
| Environmental – Enclosure | <ul style="list-style-type: none"> ▪ Dual 1600W, 80+ Platinum ▪ 200-240V AC input, auto ranging, 50-60Hz ▪ Operating temperature: 5 to 40°C, Non-op temperature: -40 to 60°C ▪ Humidity: 5 to 85% relative humidity |
| Size – Controller | <ul style="list-style-type: none"> ▪ 2RU; H:3.5" (88 mm), W: 19" rack (483 mm), D: 33.5" (850 mm) without bezel |
| Size - Enclosure | <ul style="list-style-type: none"> ▪ 4RU Height: 175mm (6.89") ▪ Width: 447mm (17.61"); Depth: 1047mm (41.25") ▪ Depth in rack: Max of 1197mm (47.13") w/ dual CMA – includes 24 SAS cables ▪ 4U102 w/o drives: 31.8kg (70lbs) ▪ 4U102 w/ 102 HDDs: 118.8kg (262lbs) |

¹ 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.

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.

