


IBM Unleashes the Potential of Data and AI with its Next-Generation IBM Storage Scale System 6000

New cloud-scale, high performance storage appliance for unstructured data to power the most demanding AI and capacity-intensive workload



ARMONK, N.Y., Oct. 31, 2023 /PRNewswire/ -- Today, IBM (NYSE:IBM) introduced the new IBM Storage Scale System 6000, a cloud-scale global data platform designed to meet today's data intensive and AI workload demands, and the latest offering in the [IBM Storage for Data and AI portfolio](#).

For the seventh consecutive year and counting, IBM is a [2022 Gartner Magic Quadrant for Distributed File Systems and Object Storage Leader](#), recognized for its vision and execution.¹ The new IBM Storage Scale System 6000 seeks to build on IBM's leadership position with an enhanced high performance parallel file system designed for data intensive use-cases. It provides up to 7M IOPs and up to 256GB/s throughput for read only workloads per system in a 4U (four rack units) footprint².

To leverage the economic value of both foundation and traditional AI models, businesses must focus on the *data* – their current capacity and growth forecasts, where the data resides, how it's secured and accessed, and how to optimize future data storage investments.

"The potential of today's new era of AI can only be fully realized, in my opinion, if organizations have a strategy to unify data from multiple sources in near real-time without creating numerous copies of data and going through constant iterations of data ingest," said Denis Kennelly, general manager, IBM Storage. "IBM Storage Scale System 6000 gives clients the ability to do just that – brings together data from core, edge, and cloud into a single platform with optimized performance for GPU workloads."

The IBM Storage Scale System 6000 is optimized for storing semi-structured and unstructured data including video, imagery, text, instrumentation data, etc., that is generated daily and accelerates an organization's digital footprint across hybrid environments. With the IBM Storage Scale System clients can:

Expect greater data efficiencies and economies of scale with the addition of IBM FlashCore Modules (FCM), to be incorporated in 1H 2024:

- New maximum capacity NVMe FCM will provide capacity efficiency with 70% lower cost and 53% less energy per TB vs. IBM's previous maximum capacity flash drives for IBM Storage Scale System.³ This can help clients realize the full performance of NVMe with the cost advantages of Quad-level Cell (QLC).⁴
- Powerful inline hardware-accelerated data compression and encryption to help keep client data secured even in multi-user, multi-tenant environments.
- Storage Scale System 6000 with FCM will support 2.5x the amount of data in the same floor space than the previous generation system.⁵

Accelerate the adoption and operationalization of AI workloads with IBM watsonx:

- Engineered with a new NVMeoF turbo tier, new parallel multi-tenant data isolation and IBM patented computational storage drives, this is designed to provide more performance security and efficiency for AI workloads.
- Storage Scale software, the global data platform for unstructured data that powers the Scale System 6000, connects data with an open ecosystem of multi-vendor storage options including AWS, Azure, IBM Cloud and other public clouds, in addition to IBM Storage Tape.

Gain faster access to data with over 2.5x the GB/s throughput and 2x IOPs performance of market leading competitors⁶:

- High-processing throughput and access speed with multiple concurrent AI and data-intensive workloads that can be run to meet a range of use cases.

The Benefits of IBM Storage Scale in Applied AI at the University of Queensland

The University of Queensland (UQ) is a world-class research institution that has significantly accelerated a wide range of workloads providing faster access to data and improved efficiency and capabilities using the IBM Storage Scale global data platform and IBM Storage Scale System. [Examples](#) of some of the research where IBM Storage is used include applied AI for the characterization of neurodegenerative diseases and in the search for more effective and flexible vaccine technologies.

"With our current Storage Scale Systems 3500, we are helping decrease time to discovery and increase research productivity for a growing variety of scientific disciplines. For AI research involving medical image analysis, we have decreased latency of access by as much as 60% compared to our previous storage infrastructure. For genomics and complex fluid dynamics workloads, we have increased throughput by as much as 70%," said Jake Carroll, Chief Technology Officer, Research Computing Centre, The University of Queensland, Australia. "We get all the benefits of a high-speed parallel file system inside our supercomputing resources with the data management transparency and global data access that the IBM Storage Scale software provides."

Carroll added, "IBM's Storage Scale System 6000 should be a gamechanger for us. With the specs that I've seen, by doubling the performance and increasing the efficiency, we would be able to ask our scientific research questions with higher throughput, but with a lower TCO and lower power consumption per IOP, in the process."

Accelerating AI with the IBM Storage Scale System and NVIDIA Technology

The Storage Scale System 6000 has the ability to create an information supply chain from an NVIDIA AI solution to other AI

workloads independent of where they are located. IBM's new NVMeoF turbo tier has been engineered for small files like those collected from remote devices or to provide access to smaller transactions like data lake or lakehouse analytics so they can be integrated into an NVIDIA solution.

The Storage Scale System 6000 supports NVIDIA [Magnum IOTM GPUDirect® Storage](#) (GDS) with a direct path between GPU memory and storage. It has also been designed to increase performance with data movement IO when GDS is enabled. Utilizing [NVIDIA ConnectX-7™ NICs](#), the Scale System 6000 supports up to 16 ports of 100Gb RDMA over Converged Ethernet (RoCE), 200Gb/s and/or 400Gb/s InfiniBand, or a combination of both to increase performance between nodes or directly to NVIDIA GPUs. To learn more about how IBM Storage can accelerate access and provide a global data platform with NVIDIA AI solutions, visit <https://www.ibm.com/storage/nvidia>.

For more information on IBM Storage Scale System, please visit our website [here](#).

Clients can select the level of desired support at the time of sale through IBM Expert Care with a choice between Basic and Advanced tiers. Read [here](#) for more.

Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

About IBM

IBM is a leading provider of global hybrid cloud and AI, and consulting expertise. We help clients in more than 175 countries capitalize on insights from their data, streamline business processes, reduce costs, and gain the competitive edge in their industries. More than 4,000 government and corporate entities in critical infrastructure areas such as financial services, telecommunications and healthcare rely on IBM's hybrid cloud platform and Red Hat OpenShift to affect their digital transformations quickly, efficiently, and securely. IBM's breakthrough innovations in AI, quantum computing, industry-specific cloud solutions and consulting deliver open and flexible options to our clients. All of this is backed by IBM's long-standing commitment to trust, transparency, responsibility, inclusivity, and service.

Visit www.ibm.com for more information.

¹ **Gartner disclaimers:** Gartner, Magic Quadrant for Distributed File Systems and Object Storage, 19 October 2022, Julia Palmer, et. Al. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

GARTNER is a registered trademark and service mark, and MAGIC QUADRANT is a registered trademark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and are used herein with permission. All rights reserved.

² The performance considerations were obtained by running sample programs in a controlled environment with standard hardware/software configurations and testing procedures. Performance for IOPs is obtained using non protected tier for capability supported via special request. 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.

³ Capacity efficiency improvements based on comparison of new 38TB Flash Core Modular (FCM) drives with up to 2:1 inline

compression vs. previous 30TB max capacity drives. Cost per TB based on standard list price of drives (list pricing available upon request) and energy usage based on specification of drives (<https://www.ibm.com/downloads/cas/JBVQYVXB>).

⁴ Adshead, A. (2023, March 14). QLC edges into mainstream, but what are its advantages? *ComputerWeekly.com*.
<https://www.computerweekly.com/feature/QLC-edges-into-mainstream-but-what-are-its-advantages>

⁵ The 4 rack unit Storage Scale System 6000 capacity is based on a configuration using the 38TB FCM drives with up to 2:1 inline compression supported in 1H 2024 with 900 TB / rack unit of floor space vs previous generation 2 rack unit Scale System 3500 using 30TB Flash drives with 360 TB / rack unit of floor space.

⁶ The performance considerations were obtained by running sample programs in a controlled environment with standard hardware/software configurations and testing procedures. Performance for IOPs on IBM Storage Scale System 6000 is obtained using non protected tier for capability available via special request. Competitor comparisons are extrapolated from publicly available information comparing market leading vendors/competitors defined by Gartner magic quadrant (<https://community.ibm.com/community/user/storage/blogs/david-wohlford1/2022/10/21/ibm-is-a-leader-in-2022-gartner-magic-quadrant>) 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.

Media Contact

Alexandra Demetriades

Alexandra.Demetriades@ibm.com

SOURCE IBM

Additional assets available online:



<https://stage.mediaroom.com/ibmnewsroom/2023-10-31-IBM-Unleashes-the-Potential-of-Data-and-AI-with-its-Next-Generation-IBM-Storage-Scale-System-6000>