


IBM Storage for Data and AI

Accelerate AI innovation with faster access to more data



IBM Storage Scale 6000

Faster than any competitive solution to access data ... YB of data




Let's create ↻ smarter ways of putting your data to work

AI spending is expected to grow to **\$301 billion** by 2026

Data that provides the fuel for AI is expected to grow over **250%** in the next 5 years

Organizations Need to Unlock Value From Data...
Wherever it Resides



“ 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.**”

However, leaders are faced with unprecedented **data challenges**

Changing Technology



Traditional infrastructure limitations

Deploying infrastructure to support new generative AI and AI workflows is **top initiative** and traditional infrastructure is not designed for new requirements and maintaining current needs

Data Silos



Data is difficult to access across silos

82% of enterprises are inhibited by data silos and **87%** run multi-cloud environments, with **44%** running applications siloed on different clouds

Sustainability and Costs



AI requirements drive the need for more resources

80% of enterprises are working with or planning to adopt AI and **86%** have a sustainability strategy that strive to lower and not increase the number of new resources

Unknown Threats



As AI becomes mission critical, the data needs more protection

As of 2023, over **72 percent of businesses** worldwide were affected by ransomware attacks.

... Leading to more cost and complexities with with current infrastructure

We need to break free from the challenges that can restrict innovation

The goal is to accelerate AI and open more possibilities to achieve business results

Challenges



- Traditional infrastructure limitations
- Data is difficult to access across silos
- AI requirements drive needs for more resources
- AI data needs more protection



The Goals



- Accelerate new workloads to achieve faster results
- Single view of data to enhance workload productivity no matter where the location or what type of data
- Create a more sustainable infrastructure that uses less energy and does not overrun budgets
- Proactively protect data and quickly recovery when attacked with an always available data experience


IBM Storage for Data and AI

Unlock unstructured data's potential and propel forward business growth and AI innovation by accelerating discovery, eliminating data silos, support green and lower cost initiatives and protecting from the unknown



Storage Scale

The foundation for business growth and AI innovation creating a global data platform



Storage Ceph

Operationalize AI with a data lake designed to improve as it scales

NVIDIA DGX™ ... watsonx™ ... any AI workloads

High Performance Computing (HPC)

AI Training and AI Inference

Generative AI

Data Lakes and Analytics

Backup / Archive and Collaborate

IBM Storage Scale and Storage Scale System

The foundation for AI innovation and accelerated discovery

Expedite AI-discovery and workflows

256GB/s and 5M IOPS that can scale to 1000s of storage systems and YB of capacity¹



Collaboration with investment protection

Access data from anywhere that includes tape, cloud and your current storage resources²



Support green and lower cost initiatives

Lower resource requirements with configurations as low as 1.1 W/TB energy and 3.6PBe/4u density and up to 70% less cost per TB³



Safeguard data from the unknown

Support cyber secure initiatives with Safeguarded copy and new Cyber-protect with up to 99.9999 availability⁴



The global data platform with a storage database and data catalog

Consolidate workloads

Easy to start

Easy to grow

Multi-cloud support

Mix and match

Now the ultimate data acceleration storage system for Hybrid AI



Ultimate Performance

up to **5M IOPs** and
256 GB/s per node¹



Scale System 6000

Ultimate Efficiency

70% lower cost / TB

53% lower Watts / TB²

- Scales from 1 to 1000+ nodes and 633YB of capacity
- Connects to non-IBM and public cloud storage
- Integrates to current IBM Storage Scale platform
- Safeguard data with Safeguarded Copy, CyberVault and CyberProtect
- Up to 99.9999% availability⁴

Ultimate Density

96 TB to 2.6 PBe

of effective flash capacity per 4 rack units (4u)³

¹ Performance results are read access and based on internal testing using standard configurations

² Based on energy specifications and standard IBM list pricing (available upon request) of new max capacity 38TB Flash Core Module (FCM) drives available 1H 2024 vs specifications and list pricing of previous max capacity 30TB flash drives

³ 4 Max effective capacity based on new 38TB Flash Core Module (FCM) drives available with up to inline 2:1 compression 1H 2024

⁴ Disclaimer: 99.9999% availability requires a replicated and multi-site configuration installed by customer. Individual results may vary depending on specific client environment. Since availability 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 as the experimental results reported in this information.

What the Scale System 6000 looks like under the covers

A single 4u node with active active controllers and redundant hardware to maximize always on data



Engineered for performance and efficiency

Processor per canister

Dual AMD EPYC Genoa 48C

Memory per Canister

24 x 32GB (768GB) – default base

24 x 64GB (1536GB) - option

Storage

48 U.2 G4 NVMe (24 and 48 drives)

NVMe: 3.84TB, 7.6TB, 15TB, 30TB

FCM: 19TB and 38TB¹

Networking

NVIDIA CX7 supported cards:

400Gb single port (IB only) x16 Gen5

200Gb VPI dual port (IB/ETH) x16 Gen5

Consolidate workloads on a single global data platform

Easy to start



Scale System 3500



Scale System 6000

Easy to grow



Scale up



Scale out

9

Multi-cloud support



Mix and match



Faster time to insights with high-performance data storage for AI research

Customer Challenges

To speed up research collaboration across multiple sites, including for complex, AI-driven projects, UQ needed a storage solution that supports hundreds of terabytes of data generated daily with high-speed access to both active and archived data.

Organizational Outcomes


- 2 hours to achieve ROI that saves researchers hundreds of hours of processing time per week
- ~74% faster runtimes for medical imaging analysis to help speed time for critical research
- Exponentially increasing data volumes with online classification and cost-effective data platform

Technology Choices

- IBM Storage Scale
- IBM Storage Scale System
- IBM Storage Fusion HCI System
- IBM Storage Fusion data catalog

Scale System 6000 – “a gamechanger”

Seamless access to data at University of Queensland




“ IBM’s Storage Scale System 6000 should be a gamechanger for us. By doubling the performance and increasing the efficiency of the new Scale System 6000, we would be able to ask our scientific research questions with higher throughput - but do that with a lower TCO and lower power consumption per IOP, in the process.

Jake Carroll, Chief Technology Officer Research Computing Centre
The University of Queensland

Why IBM Storage and **NVIDIA** are better together

Accelerate AI workloads with more performance, access, efficiency and protection



Where does IBM Storage for Data and AI make a difference

AI Training/ Inference
/ HPC

Extreme performance critical

AI with NVIDIA
LLM creation
Digital twin
Foundational model creation
Genomics research
Computational chemistry
Weather prediction
Seismic imaging
Financial fraud/claims processing
Financial risk modelling
Autonomous driving

Generative AI
Data Lake /Analytics

Performance priority

watsonx.data
Customer analysis/support
Healthcare treatment/diagnosis
Customer care analysis
Engineering design
Research/business analysis
Large pattern analysis
High performance database
Computer aided engineering(CAE)
Public safety/smart cities

Backup, Archive and
Data Collaboration

Access and costs priority

DR backup data
Archival of less used data
Video/Voice/Image applications
Remote data collaboration
Design collaboration
Video surveillance/production

Scale System 6000 – “destined to be the best *accelerator for AI*”



“Sycomp has many customers who have successfully created a seamless high performance hybrid cloud solution for AI and analytics workloads using IBM Storage Scale and the global data platform. The IBM Storage Scale System 6000 is destined to be the best accelerator for AI acceleration and enable the perfect partnership to seamlessly connect to Sycomp Storage fueled by IBM Storage Scale in the cloud.”

John Zawistowski, Global System Solutions Executive
Sycomp

Scale System 6000 – “a real breakthrough for hybrid cloud HPC solutions”




“As IBM Cloud enables customers to tackle large-scale, compute-intensive challenges and speed time to insight with hybrid cloud HPC solutions. The IBM Storage Scale software is an important part of the data platform strategy of IBM Cloud. The new IBM Storage Scale 6000 provides a breakthrough for hybrid cloud HPC solutions as it provides the highest on-prem performance for customers that seamlessly connects to the high-performance global data platform in the IBM Cloud.”

Terry Fisher

Distinguished Engineer - IBM Cloud High Performance Computing
IBM

How a Global Data Platform Works

Its all about the data and creating an information supply chain for AI



IBM