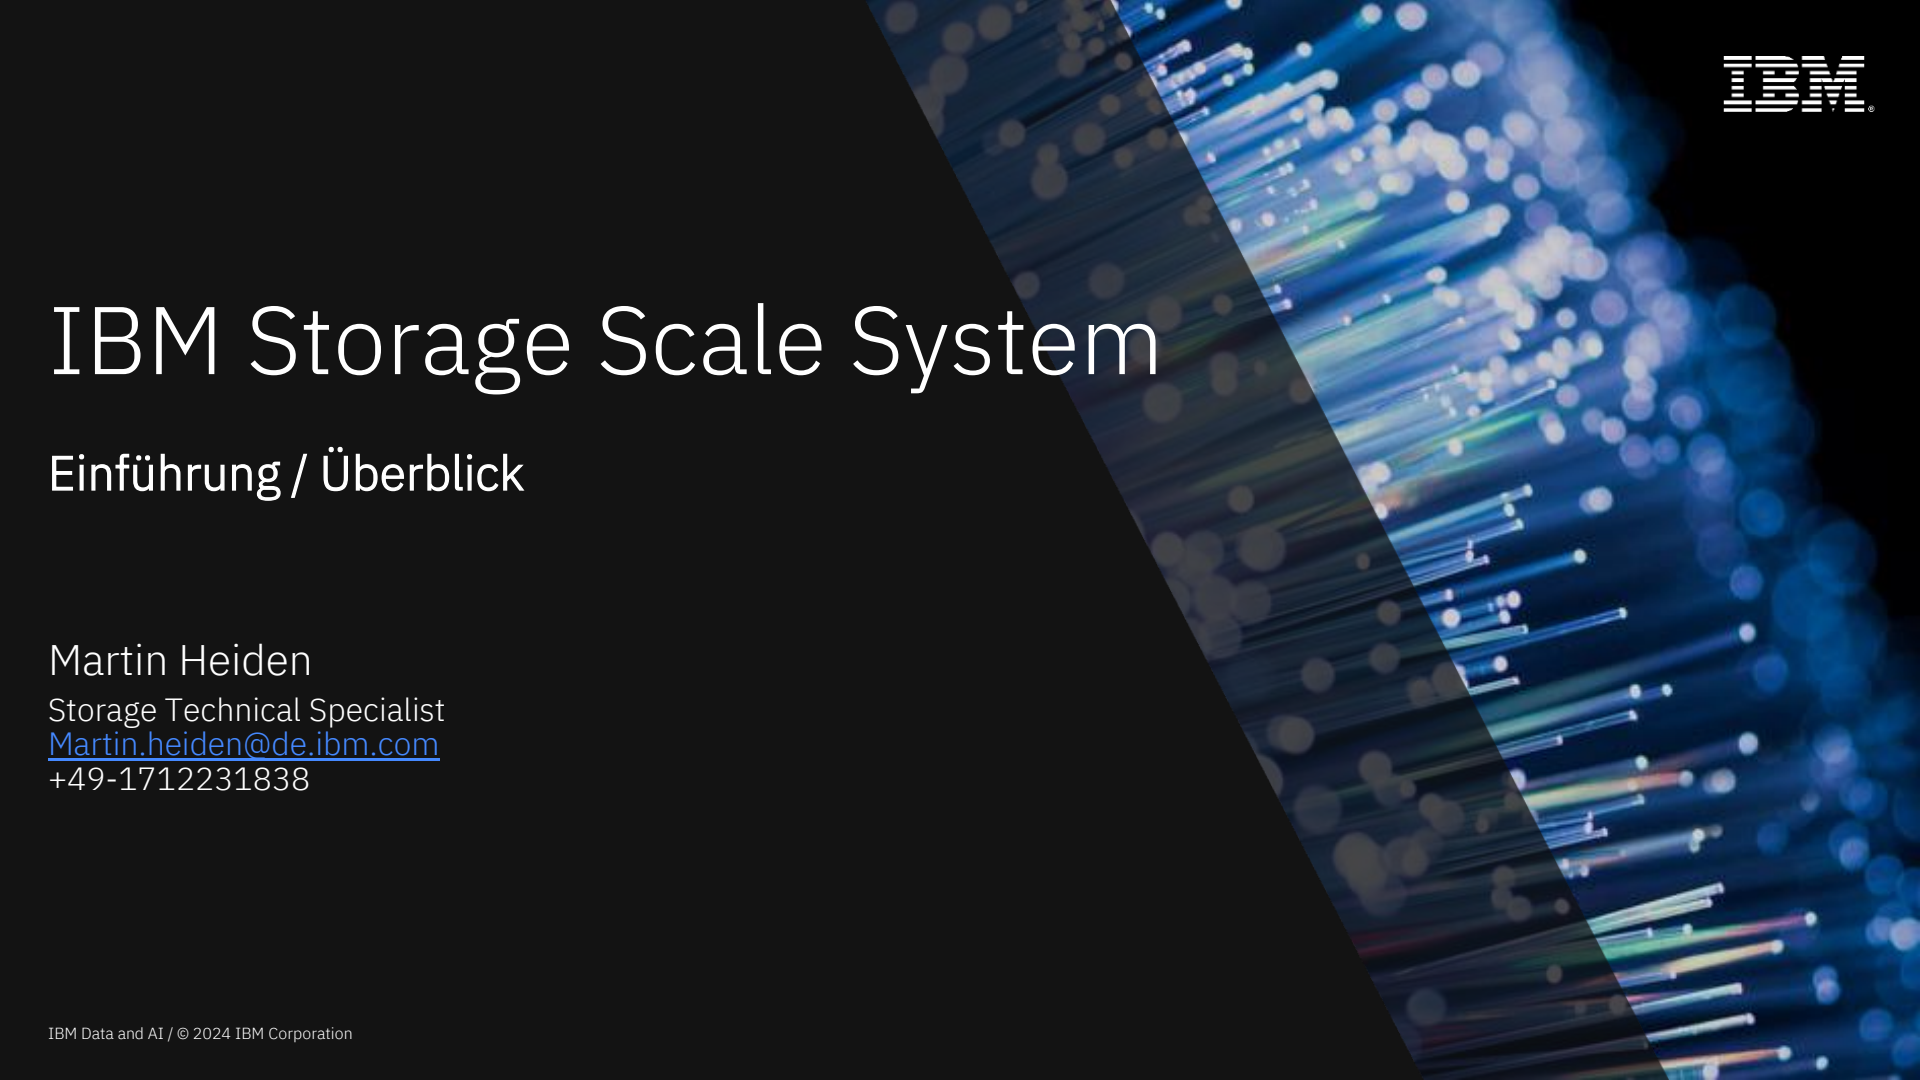

IBM Storage Scale System

Einführung / Überblick

Martin Heiden
Storage Technical Specialist
Martin.heiden@de.ibm.com
+49-1712231838

Disclaimer

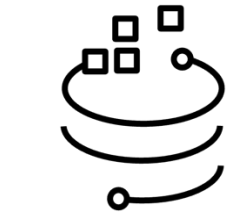

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

Benefits of integrated solutions

Building and testing is time consuming and complex

Sample tasks:

Sourcing
Source the storage, servers, networking, power, space on the data center floor

Installation
Operating system, patches, firmware, hardening, LDAP
Installing Spectrum Scale software

Test, Optimize, Deploy
Performance tuning

Maintenance
Patching of individual components (OS, server firmware, switch firmware, software update)

Storage Scale System is integrated and ready to deploy

- Storage Scale is integrated, tested, and factory preloaded, leveraging the latest software release
- Scale Systems are building blocks ready to deploy

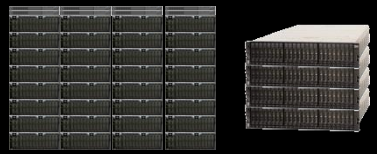

IBM Storage Scale System solution packaging

Integrated solution

Storage Scale is integrated, tested, and factory preloaded

Leverage the latest IBM Storage Scale releases

Data Management Edition and Data Access Edition

Optimal storage capacity & economy

Storage Scale System has various models providing NVMe, NL-SAS or both

Choose from various sizes of NVMe or HDD

Rack-mountable solution

Non-disruptive upgrades

Capacity upgrades can be performed without application disruption

Software automatically rebalances data across all drives

High performance connectivity

Supports EDR/HDR/NDR InfiniBand or 100 GbE / 200 GbE Ethernet high-performance networking

Storage Scale System is validated with NVIDIA POD architectures for AI modeling.

IBM differentiator: Storage Scale RAID

Erasure code software provides enterprise storage performance using standard, inexpensive disk drives

Faster rebuild times
- More disks are involved during rebuild
- Approximately 3.5x faster vs RAID-5

Minimal impact of rebuild on system performance
- Rebuilds are done by many disks
- Rebuilds can be deferred with enough protection

Better fault tolerance
- End to end checksum reduces or eliminates file system checks
- Stable, consistent high performance

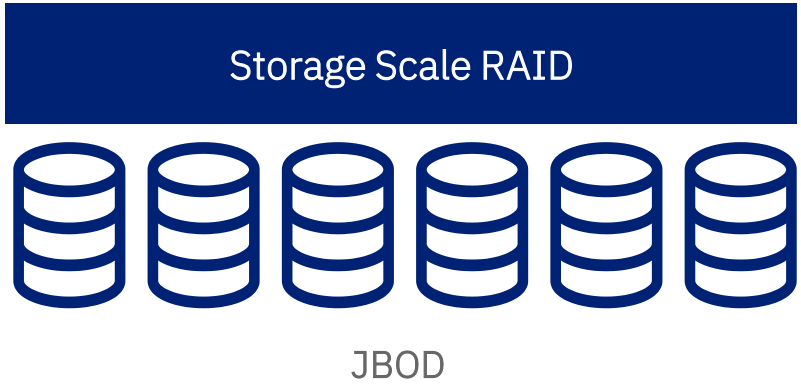

IBM Storage Scale Raid advantages

Faster and smarter rebuild operations compared to RAID arrays

- Uses many drives in parallel, distributes work across many nodes
- Normal rebuilds have minimal impact on system performance
- Critical rebuilds complete in minutes
- Rebuilds can be deferred with sufficient protection

Improved storage efficiency compared to replication

- 8+2P and 8+3P offer 25% - 38% overhead vs 100% - 200% for replication
- 4+2P and 4+3P also supported
- Spare capacity is also distributed across all drives and nodes

Higher performance than traditional erasure code implementations

- Patented strategies optimize IO data paths, read and multi-layer write caching
- Suitable for analytics, AI, and demanding read/write workloads, not just read-heavy workloads or cool data

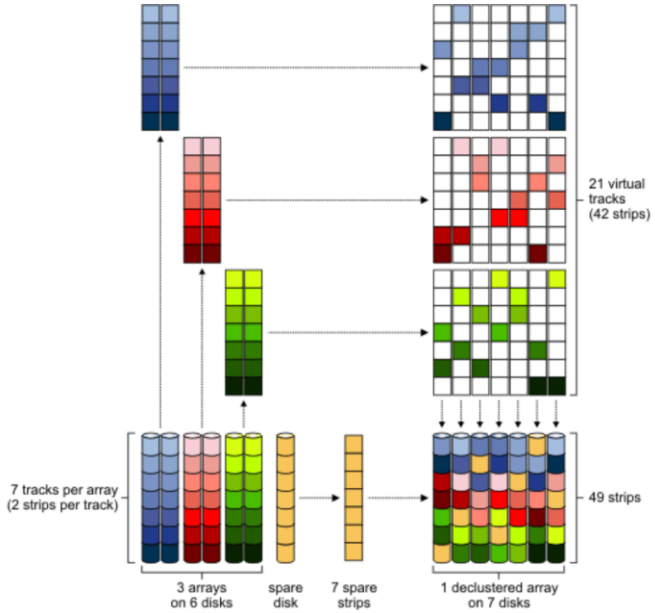

IBM Storage Scale System Evolution

2014 2017 2019 2020 2021 2022 2023 2024

Performance Use Case

Hybrid Use Case

Capacity Use Case

ESS Gen 1
GS/GL Models
(HMC required)

ESS Gen 2
GLxS, GSxS
Models

ESS 3000
(containerized deployment)

ESS 5000

ESS 3200

FCM 2.1

Hybrid Flash and Disk

ESS 3500

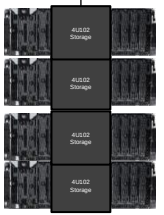

(up to 8 JBODs)
[Daisy-chain 5-8]

6000 – PCIe Gen5/24G SAS

(up to 9 JBODs)

FCM 4.0
(with compression)

2U x86 utility node with support for EMS, GUI, Callhome, Protocol AFM, GKLM, Protect Plus and Discover

P8/9 Management Server

Optional Protocol Node(s)

Scale System models are built for speed and capacity

Storage Scale

Speed

Capacity

Hybrid

IBM Storage Scale System 6000

4U48 Enclosure
24 or 48 NVMe / FCM drives

Performance
Machine type 5149-F48

Up to 310+GB/s
IOR 100% read, InfiniBand

IBM Storage Scale System 3500

2U24 Enclosure
12 or 24 NVMe drives

Performance
Machine type 5141-FN2

Up to 125 GB/s
IOR 100% read, InfiniBand

IBM Storage Scale System 3500 Hybrid (NVMe + HDD)

Up to 8x 4U102 JBOD

3500 Hx
Machine type 5141-FN2

Up to 91 GB/s – NVMe
IOR 100% read InfiniBand
Up to 48 GB/s – HDD
IOR 100% read InfiniBand

IBM Storage Scale System 6000 Hybrid (NVMe + HDD)

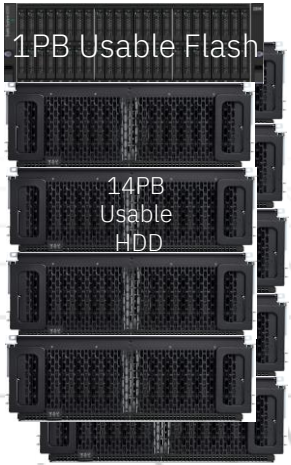

Up to 9x 4U91 JBOD

6000 Hx
Machine type 5149-F48

Up to 280 GB/s – NVMe
IOR 100% read InfiniBand
Up to 100 GB/s – HDD
IOR 100% read InfiniBand

IBM Storage Scale System 3500 Capacity (HDD-Only)

Up to 8x 4U102 JBOD

3500 Cx
Machine types
5141-FN2

Up to 48 GB/s
IOR 100% read C4+ model, InfiniBand

IBM Storage Scale System 6000 Capacity (HDD-Only)

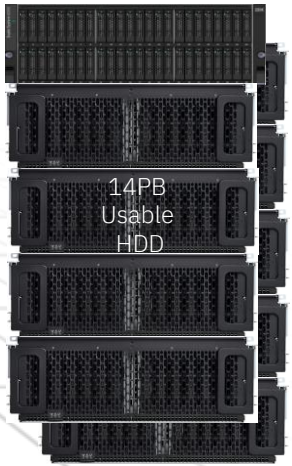

Up to 9x 4U91 JBOD

6000 Cx
Machine types
5141-FN2

Up to 90 GB/s
IOR 100% read C9 model, InfiniBand

Current IBM Scale System Models at a glance

Storage Scale

Models	6000 NVMe / FCM	3500 3500Hx NVMe	6000Hx 6000Cx	3500Hx 3500Cx first external enclosure	3500Hx 3500Cx additional external enclosures
	Half or full 2U48 24 or 48 drives	Half or full 2U24 12 or 24 drives	Full only @4U91 Storage enclosure 91 drives	Half or full @4U102 Storage enclosure 52 or 102 drives	Full only 2 to 8 @ 4U102 Storage enclosure 102 drives
Drive sizes	NVMe 3.84 TB 7.68 TB 15.36 TB 30.74 TB FCM 19.2 TB 38.4TB	NVMe 3.84 TB 7.68 TB 15.36 TB 30.74 TB	HDD 12 TB - SED 16 TB - SED 20 TB – SED 22TB - SED	HDD 10 TB - SED 14 TB 18 TB 20 TB – SED 22TB - SED	

*all Flash media is SED capable

6000Cx and 3500Cx has 4 NVMe

IBM Storage Scale System 3500

Dual redundant side-by-side control canisters in 2U
- Dual path to all drives
- Reed-Solomon 8+2P or 8+3P encoding
- Redundant Power Supplies, Fans, Network
- All FRUs are Hot Swap with Status LEDs

12 or 24 High performance PCI Gen4 NVMe drives
- 3.84 TB, 7.68 TB,15.36 TB or 30 TB drives
- 46 TB to 730 TB raw capacity

Four x16 PCIe Gen4 PCI adapter slots per canister
- 2x IB NDR / RoCE per canister
- Up to 8x 4U x102 SAS HDD enclosures

Support for Self Encrypting Drives

Scale from 1 to 1000s of nodes

up to 126 GB/s per BB

up to 300M IOPS per rack

48 TB to 738 TB raw Flash capacity

510 TB to 16.3 PB raw HDD capacity

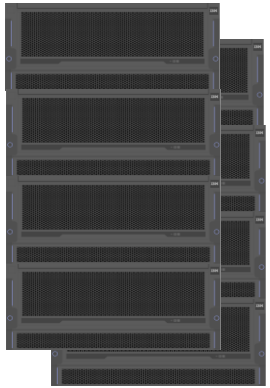

What the Scale System 6000 looks like under the covers

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 9454 48C 3.8Ghz

Memory per Canister

- 24 x 32GB (768GB) – default base
- 24 x 64GB (1536GB) – optional

Storage

- 48 U.2 G4 NVMe (24 and 48 drives)
- NVMe: 3.84TB, 7.6TB, 15TB, 30TB
- FCM: 19.2TB, 38.4TB

Networking

- NVIDIA CX7 supported cards:
- 400Gb single port (IB only) x16 Gen5
- 200Gb VPI dual port (IB/ETH) x16 Gen5

IBM Storage Scale System 6000

Capacity and Hybrid Support

- 310+ GB/s Read performance in a 4u48 form factor
- 1.5PB of High performance NVMe / 4U (15 PB High perf NVMe in a rack)
- Up to 5.5PB FCM Flash Capacity (3:1 Compression)
- VM Integrated Protocol support (SMB/NFS) on controller nodes
- Sixteen PCIe Gen5 adapter slots (SAS HBA and CX-7 VPI)
- Four Boot drives
- Two integrated 1GbE network ports (separation of MGMT and BMC)
- Four 10GbE network ports

- Introduction of 4U91 24Gb SAS enclosure (up to 9 JBODs per rack)
  - Increased SAS performance relative to 12Gb SAS enclosures
  - HDD only - 12 TB, 16 TB, 20TB SED, 22TB SED drives
  - 18PB of SAS HDD Capacity per Rack

ISS 6000 Performance: 24 or 48 NVMe (3.84 TB, 7.68 TB, 15.36 TB, 30.75TB or 19/38TB FCM)

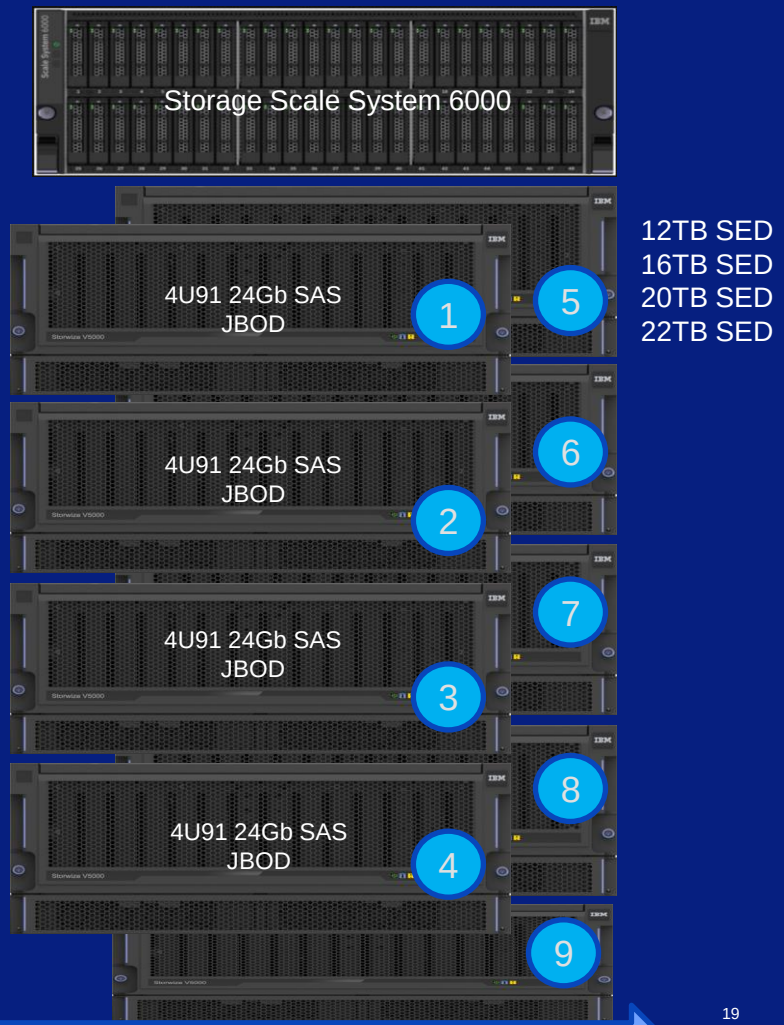

Performance → **Capacity**

Storage Scale System 6000 canister Ports and LEDs

Dual Boot Drives (mirrored)

2 x 10GbE Mgmt Network

BMC (1GbE)/1GbE host

Dual CXL(not used)

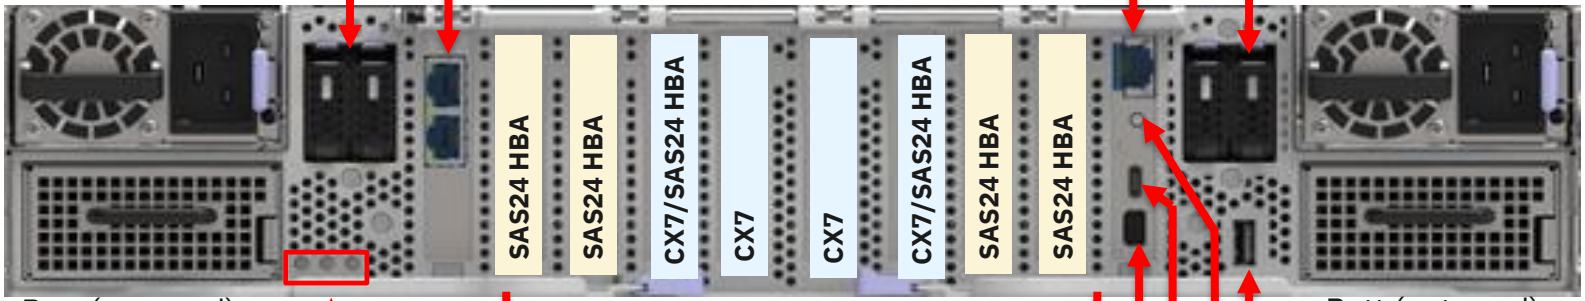

Batt (not used)

Canister LEDs

Batt (not used)

PCI Adapter Placement
(Showing hybrid/capacity Configuration)

USB

Push Button

Serial

Display

Storage Scale System 6000 slot usage

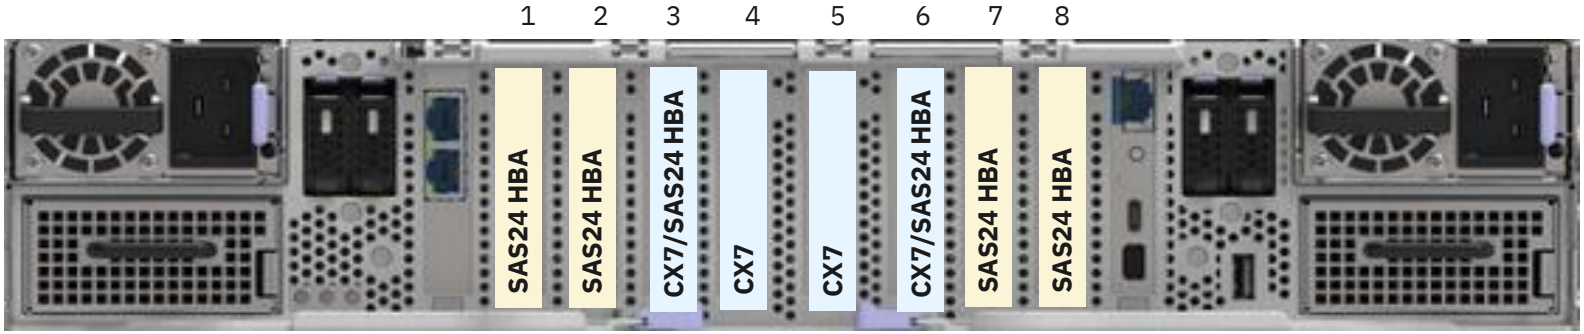

PCIe Slot	Usage	PCIe Slot	Usage
1	SAS4 G4 (Hybrid/Capacity) *)	5	**Nvidia CX7 Network x16 G5** (Performance)
2	SAS4 G4 (Hybrid/Capacity)	6	***Nvidia CX7 Network x16 G5*** *(Performance)* SAS4 G4 (Hybrid/Capacity)
3	**Nvidia CX7 Network x16 G5** (Performance) *SAS4 G4 (Hybrid/Capacity)*	7	SAS4 G4 (Hybrid/Capacity)
4	**Nvidia CX7 Network x16 G5** (Performance)	8	SAS4 G4 (Hybrid/Capacity)

*) Statement of Direction: IBM intends to deliver SAS G4 and HDD Enclosures support in 2024

IBM FlashCore™ Module 4

Capacity and Performance

2.5” dual ported U.2 NVMe Gen 4 PCIe
Industry leading density at 38.4 TB per drive
Inline hardware FIPS 140-3 encryption
Inline hardware 3:1 compression = 116 TB!

Internally tiered storage
-> MRAM -> SLC -> 3D QLC

Industry leading QLC endurance
15K Program/Erase cycles
Compared to 1500 for enterprise QLC

IBM Unique QLC management (100+ patents)
read calibration, heat binning, health binning, error correcting codes, optimized voltage

2U X86 Utility Node

All-purpose, powerful and fully integrated utility node, supporting multiple use cases and compatible with existing building blocks

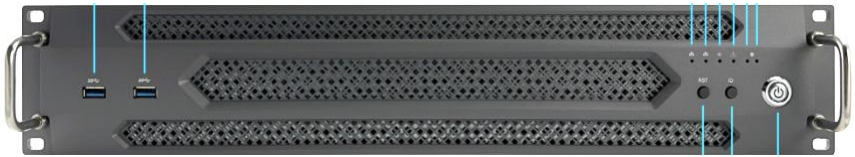

Replaces existing power-based EMS and Protocol node and adds support for additional storage use cases

System Config

- Processor: AMD EPYC (single/dual docket)
- Memory: 128GB – 512GB
- 2x internal boot drives
- High-Speed Network: 1-4 CX-6 adapters
- 1Gb/10Gb network

Versatility and Flexibility

- Support for EMS, GUI and Callhome
- Support for Protocol node functions
- Support for AFM gateway
- Support for GKLM (orderable via AAS)
- Support for IBM Storage Protect and Discover

IBM Storage Scale System 6000 and 3500 highlights

Fast time to value

Preconfigured, fully tuned

Easy to install or update by sysadmin or developer

Perfect for growing GPU workloads

Linear performance scaling

Operational efficiency

Containerized install and update software for a fast and easy out-of-the-box experience

6000 High performance, high density: 310 GB/sec and up to 845 TB usable NMVe capacity per 4U system

3500 High performance, high density: 126 GB/sec and up to 435 TB usable NMVe capacity per 2U system

High performance, high capacity: up to 55 GB/sec and 11 PB HDD usable capacity per 5000 (SC9 model)

Reliability

IBM Storage Scale erasure coding

Fast, non-disruptive data rebuild

Automated monitoring of key hardware components

Deployment flexibility

High-performance tier for any Storage Scale / Scale System cluster

Start as small as 46 TB, scale out to exabyte capacity

Loosely-coupled edge data management as component of global unified data solution

Supported with AI solutions based on Intel, IBM Power, Z, and LinuxONE

Jülich Supercomputing Center - Exascale HPC with IBM Storage Scale System

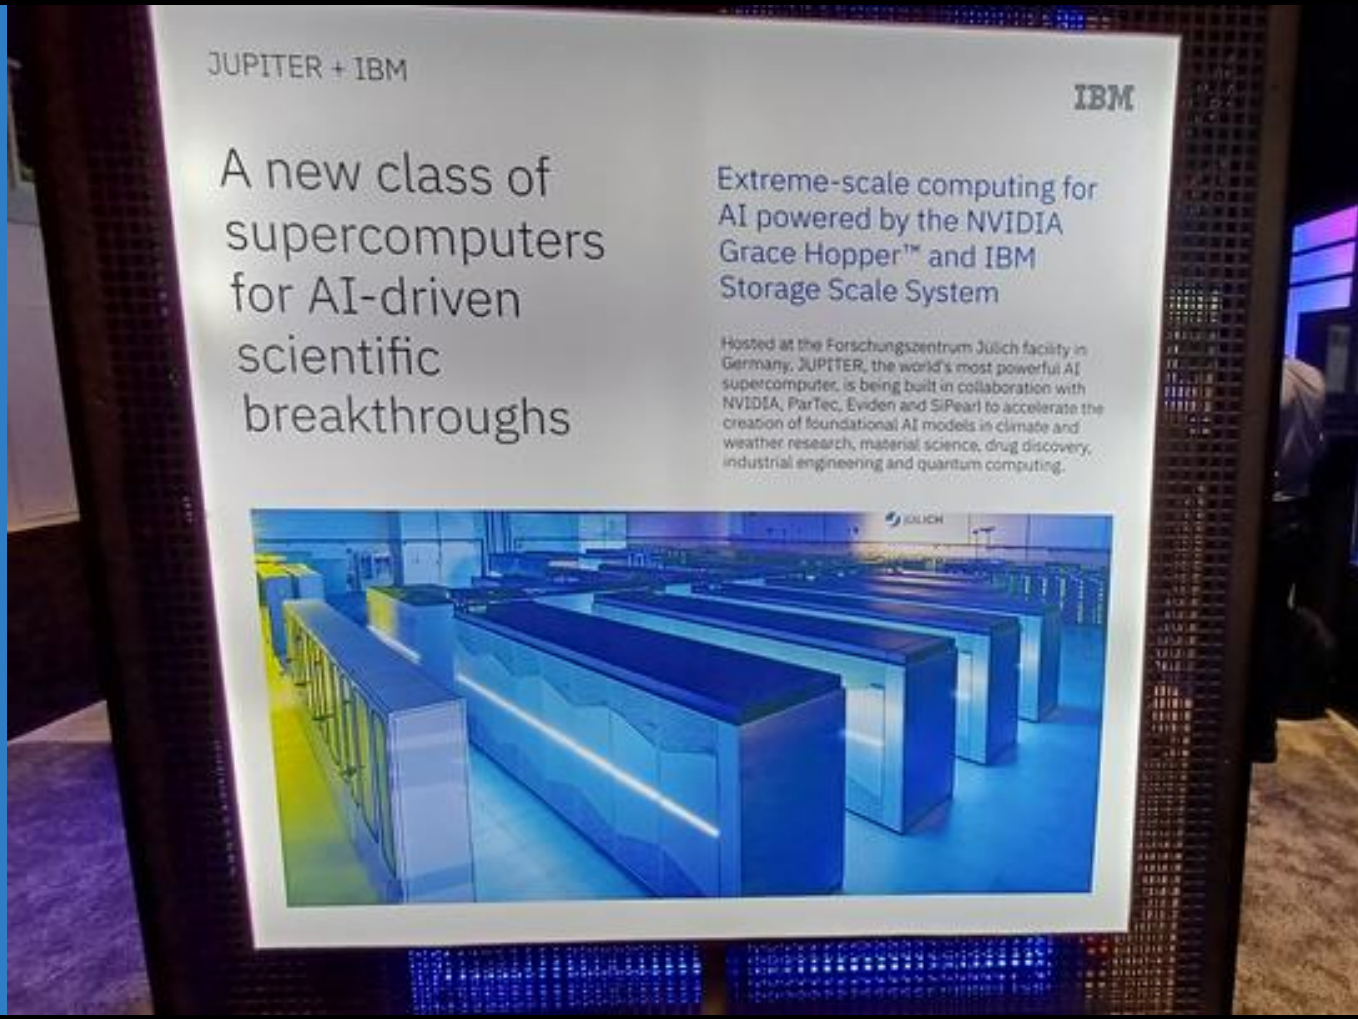

IBM Storage Scale System 6000

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 330GB/s throughput; up to 13M IOPS using NVMeoF, up to 3.4PBe (effective) in a 4U rack

The rapid rise of artificial intelligence (AI), high-performance computing (HPC), analytics, and hybrid cloud is transforming the business world. Unlike traditional applications that rely on structured databases, these modern workloads and data lakes process vast amounts of unstructured data, including documents, audio, images, and videos.

To stay competitive, organizations must rethink their storage strategies and leverage AI to unlock the value of their data—wherever it resides. IT leaders face key challenges, including:
- Accessing and analyzing data and workloads scattered across the globe.
- Managing the growing AI infrastructure and ensuring scalability for evolving workloads.
- The increasing time needed by AI training and inferencing workloads.

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.

"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

Storage Scale achieves this with active file management (AFM), which abstracts other storage systems, including customers' existing on-prem or public cloud storage, while providing a transparent caching layer that improves performance by unifying data into a single global namespace, accessible simultaneously via multiple protocols.

Built on a massively parallel file system, Storage Scale supports deployment across x86, IBM Power, IBM Z, ARM-based POSIX clients, virtual machines, and Kubernetes. It delivers high-performance access to unstructured data across any protocol, location, or format—helping customers accelerate AI pipelines and unlock the value of their data estate.

Unlocking AI Potential with Content-Aware Storage

Very little enterprise data has been indexed for generative AI applications, which prevents AI assistants from providing accurate, up-to-date answers. The content-aware storage capabilities in Storage Scale address this challenge by extracting the semantic meaning hidden inside unstructured data so that AI assistants can automatically generate smarter answers. Storage Scale enriches data using embedded compute and data pipelines that minimize data movement and latency to help reduce costs and improve performance.

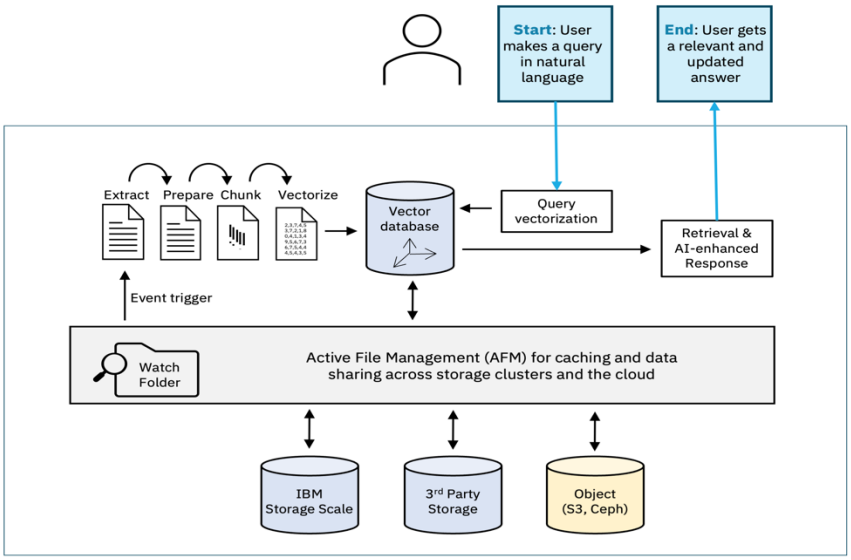

Figure 1. Storage Scale automates data extraction, vectorization, and storage updates, enabling seamless retrieval via a vector database. When users submit natural language queries, AI enhances search results for optimized responses.

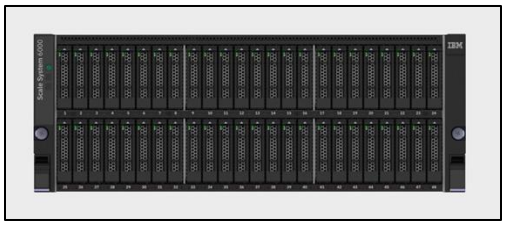

Figure 2 – The IBM Storage Scale System 6000 can deliver up to 330GB/s throughput, up to 13M IOPS, and up to 3.4PBe (effective capacity) in a 4U rack.

IBM Storage Scale System 6000

IBM Storage Scale System 6000 is a high-performance hardware platform designed to build a global data platform around file and object data. Combining IBM Storage Scale software with NVMe flash technology, it delivers extreme throughput and scalability for AI, data analytics, and storage-intensive workloads. Additionally, it supports up to nine SAS hard disk drive expansion enclosures, enabling flexible capacity growth.

330 GB/s

High-performance throughput
for demanding workloads

3.4 PBe

Massive effective capacity
in a compact 4U rack space

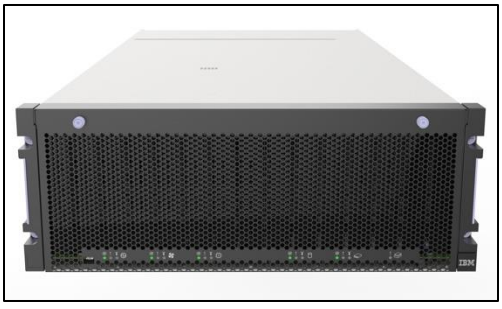

Figure 2 – The IBM Storage Scale System Expansion Enclosure enables organizations to cost-effectively deploy workloads operating on massive data sets.

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 330 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.

Storage Scale System 6000 is designed for massive scalability, supporting thousands of nodes and yottabytes of capacity. It leverages IBM Storage Scale RAID erasure coding for enhanced data efficiency, hardware failure mitigation, and intelligent monitoring with dynamic data tuning. Installations and updates are delivered by means of containerized software that speeds and simplifies the maintenance process.

For data resilience, Safeguarded Copy enables cyber-resilient, point-in-time snapshots, creating isolated, immutable copies on a scheduled basis. This ensures rapid and secure data recovery in the event of cyberattacks.

Accelerating AI Workloads
As AI and analytics datasets grow, slow data loading can bottleneck performance, limiting the full potential of high-speed GPUs. To eliminate IO constraints and optimize processing, organizations need storage built for AI-scale workloads.

IBM Storage Scale System 6000 enhances globally distributed data processing with NVIDIA GPUDirect Storage, enabling a direct data path between GPU memory and NVMe storage, including NVMe over Fabric (NVMe-oF). By bypassing the host CPU and DRAM, this architecture reduces latency, maximizes GPU utilization, and accelerates AI-driven insights.

IBM FlashCore Modules
Storage Scale System 6000 now supports 19.2TB and 38.4TB IBM FlashCore Modules (FCMs) as an alternative to standard NVMe flash drives, enhancing data density and resilience.

FCMs are IBM’s patented high-performance NVMe flash drives, designed for low latency, reliability, and efficiency. With a PCIe Gen4 U.2 interface and high-speed NAND memory, they deliver high throughput and IOPS with consistent and predictable latency.

IBM-developed custom FPGAs eliminate computational overhead, enabling compression and encryption without performance impact. Advanced flash management further extends endurance without compromising latency.

IBM Storage Scale Expansion Enclosure
Designed for enterprise-scale data growth, the IBM Storage Scale Expansion Enclosure seamlessly extends the capacity of the Storage Scale System 6000. Each unit supports up to 91 self-encrypting SAS HDDs with 20TB or 22TB of storage, ensuring robust security and performance.

With support for up to nine enclosures per system, it scales to 18PB per rack over 24Gb SAS, allowing organizations to efficiently manage multi-petabyte workloads with speed and reliability.

Specifications – IBM Storage Scale System 6000

System Features	• 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 5.4PBe[1]
Software	• IBM Storage Scale for Storage Scale System software • Red Hat® Enterprise Linux® (RHEL)
Software Features	• 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 muti-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	• Dual socket AMD EPYC™ Genoa • Sequential read performance up to 330GB/s and 13M IOPs[2]

Networking / Adapters	• 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:
	  ○ 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	• 24 or 48 U.2 NVMe SSDs (30.72TB) • 24 or 48 U.2 NVMe FCM (19.2TB or 38.4TB)
Environmental – Drives	• Nominal power: 20-25 watts per SSD • Nominal power: 30 watts per FCM
Environmental – Controller	• 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	• 4RU; H:7” (175 mm), W: 19” rack (483 mm), D: 33.5” (850 mm) without bezel
Size - Enclosure	• 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	330GB/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.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
March 2025