IBM Storage Scale System Overview

Max Huber

Storage Technical Specialist - Software Defined Storage

max.huber@ibm.com

+49-172-2594998

Disclaimer

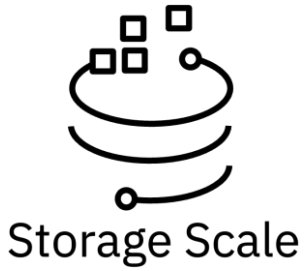

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. 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.

IBM Storage Scale deployment methods

Software

Scale System
integrated solution

Cloud services

IBM Storage Scale System for AI and Hybrid Cloud Workloads

AI: NVIDIA GPU

AI: Analytics

AI: IBM Data Fabric

Hybrid Cloud: Backup / Archive

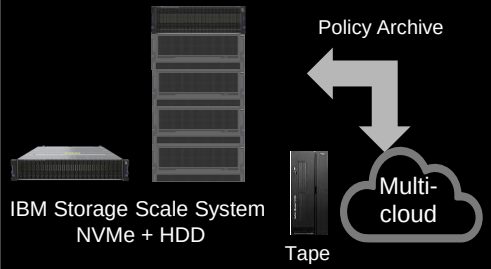

Hybrid Cloud: Data Lakes

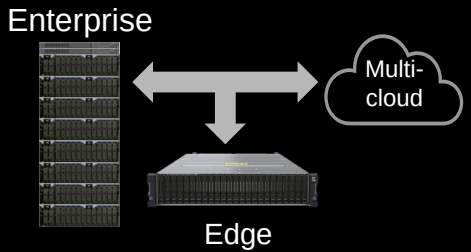

Hybrid Cloud: HPC

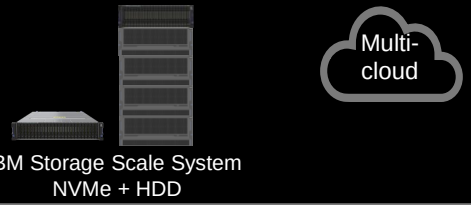

IBM Storage Scale System: Accelerating Enterprise Apps

140% improved AI training time using IBM Storage Scale (IBM Storage Scale System) and NVIDIA DGX system

Accelerate AI-enabled transformation up to **100%** while using **2x less** compute nodes for big data analytics

Achieve near native **126GB/s** IBM Storage Scale System 3500 performance

NVIDIA

hadoop Apache Spark

CLOUDERA

Red Hat OpenShift Container Storage

Results based on internal testing of 3500

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
- Easy to integrate with Public Cloud storage

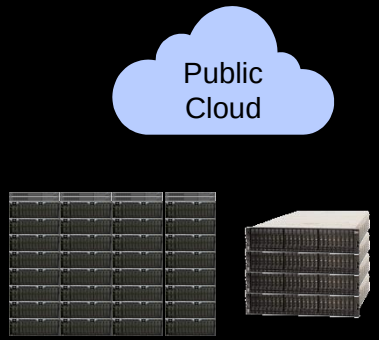

IBM Storage Scale System solution packaging

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

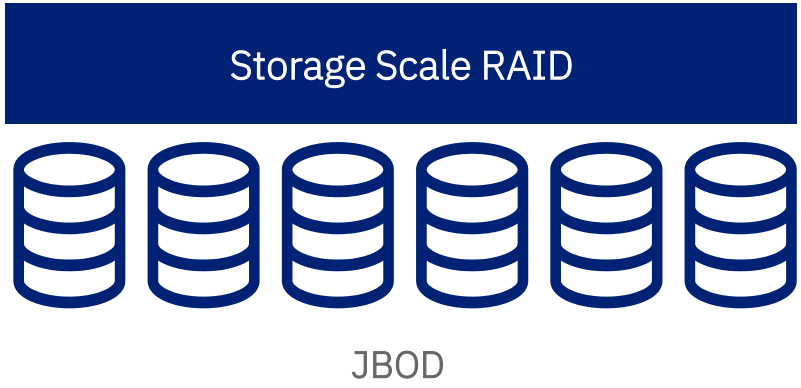

Storage Scale erasure code advantages

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

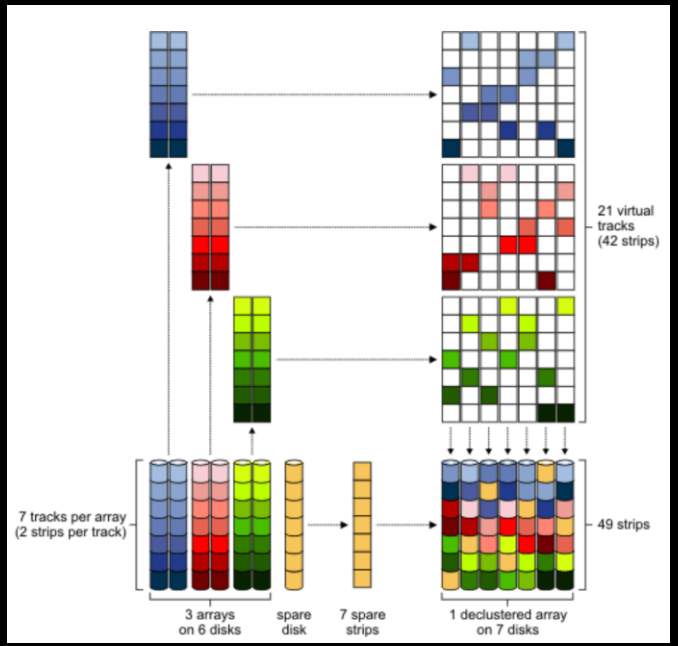

History of IBM Storage Scale (System)

Incorporating Decades of Storage Innovation

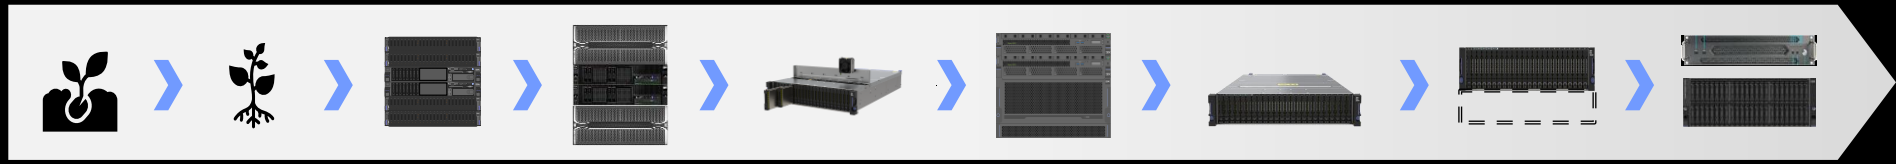

1998

GPFS
(original name of Storage Scale)

The seeds of the Global Data Platform are born

2012

AFM

Data Caching Services

2014

GSx & GLx

1st Gen with Security Services

2017

GSxS, GLxS, GLxC, & GHxy

2nd Gen with denser enclosures

2019

3000

3rd Gen with NVMe all flash cloud scale storage

2020

5000

3rd Gen with faster processor and greater capacity (1st 10PB+ node)

2021

3200

high performance all flash node

2022

3500

New Family of cloud scale innovation

2023

EMS Utility node

6000

4th Gen with NVMe

The easiest way to deploy a Global Data Platform

IBM Storage Scale System

All NVMe Flash 6000

- 48 TB to 1+ PB flash
- Up to 15+ PB capacity per 3500 Capacity
- Scales from 1 to 10000s+ of clients

3500 Hybrid NVMe Flash + HDD

Performance Optimized

- 310GB/s per 6000
- 13M+ IOPS per 6000
- Parallel access to data
- Locally cached global data with dynamic memory pools

Cost Savings

- Integrate existing non-IBM storage and cloud
- Turn off/turn on unused storage w/ tape
- Mix and match old/new

Ready for business-critical data

- End to end encryption with customer keys
- Disaster recovery with synchronous or asynchronous replication
- Ensure reliability and fast rebuild times using Storage Scale RAID's dispersed data and erasure code

Global Connectivity

- Break down silos by connecting remote systems, cloud data and non-IBM storage
- Connect any node to a global data platform online when needed
- Six 9s (99.9999%) of availability and online scalability and upgrades

Scale System building blocks and models

Each building block contains:

- A pair of Storage Scale I/O NSD data servers
  - 6000, a fully integrated 4U storage building block
  - 3500, a fully integrated 2U storage building block and up to 8 @ 4U external storage enclosures
- Uses a management server, one per Storage Scale cluster

Scale System has various models to suit a customer's needs

- Types of storage media: NVMe Flash, NL-SAS HDD
- Various sizes of NVMe (3.84, 7.68, 15.36, 30.74 TB) and HDD (10, 14, 18, 20 TB)
- Various sizes of IBM FlashCore Module (FCM) support expected in 1H2024

6000 Performance	3500 Performance	3500Hx Hybrid	3500Cx Capacity
Built for speed with NVMe Flash	Built for speed with NVMe Flash	NVMe Flash + capacity with HDD	High capacity with HDD

Scale System models are built for speed and capacity

Speed

Hybrid

Capacity

Scale System 6000

4U48 Enclosure
24 or 48 NVMe drives

Flash
Machine type 5149-F48

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

Scale System 3500

2U24 Enclosure
12 or 24 NVMe drives

Flash
Machine type 5141-FN2

Up to 126 GB/s
IOR 100% read, InfiniBand

Scale System Hybrid 3500H

3500 Hx
Machine type 5141-FN2

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

Scale System Capacity Hard Disk Drive (HDD)

3500 Cx
Machine types
5141-FN2
5147-106 or 5147-092
Storage
5105-22E POWER9 servers

Up to 48 GB/s
IOR 100% read C4+ model, InfiniBand
(pending testing C5+)

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 per node

510 TB to 16.3 PB raw HDD capacity per node

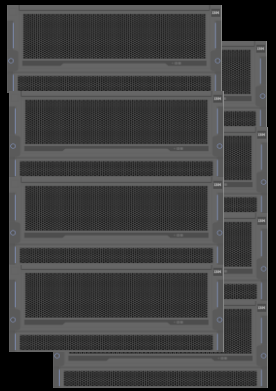

What the Scale System 6000 looks like under the covers

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

Networking

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

Comparision of IBM Storage Scale Systems

	Storage Scale System 3500	Storage Scale System 6000
Processor per canister	1 x AMD EPYC Rome 7642 CPU 48C, 3.2GHz	2 x AMD EPYC Genoa 9454 48C 3.8Ghz
Memory per canister	512GB (8 x 64GB DDR4-3200 RDIMM) For systems with more than 4 enclosures: 1,024GB (8 x 128GB DDR4-3200 RDIMM)	768GB (24 x 32GB DDR5-4800 RDIMM) 1,536GB (24 x 64GB DDR5-4800 RDIMM)
Number of SSD per system	12 or 24	24 or 48
Max. SAS Adapter per canister	2 x SAS3	4 x SAS4
Max. CX Adapter per canister	2 (hybrid), 4 (SSD only)	4 (hybrid), 4 (SSD only)

2U X86 Utility Node

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

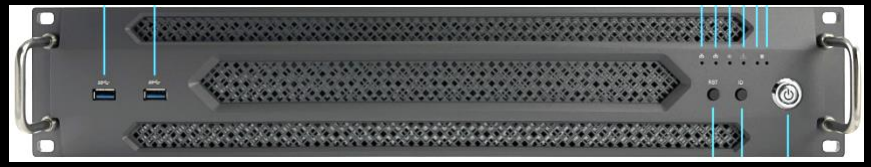

Replaces existing power-based EMS and Protocol node and adds support for additional storage use cases

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