

IBM Storage Scale System

Highlights

- A single storage platform for multiple applications and high performance workloads
- Create a cyber secure solution and restore data with a single command using data security services
- Data caching services provide consistency and global data access to existing storage resources and increases data collaboration
- Automate data lifecycle and application performance tuning with data management services
- Create a more sustainable infrastructure with a single node and a system that can turn off and on data when not used and eliminate duplicate data
- Combine with other IBM and non-IBM storage including the public cloud
- Access the same data from a file and an object application simultaneously with data access services

Easiest way to build a global data platform



All flash configuration with turbo tier option



Flash + capacity configuration 1-8 storage enclosures

- Break performance barriers for AI
- Designed capacity from 46TB to 633YB
- Reduce inefficiencies and connect data silos
- Access data globally from edge to core to cloud
- Secure data assets with multiple security services
- Up to 16M IOPs and 91GB/s¹ performance per node

IBM® Storage Scale System is designed to be the simplest way for users to deploy IBM Storage Scale and the global data platform. IBM Storage Scale System comes installed with IBM Storage Scale and is managed as a single storage system. Installations and updates are delivered by means of containerized software that speeds and simplifies the maintenance process. The system is designed to easily extend an existing IBM Storage Scale cluster or be set up quickly as a new standalone environment. Leveraged by 1000s of customers from banks to government entities to some of the largest organizations in the world. Our scale-out node architecture is also a great fit for smaller organizations with configurations starting at 48TB in a 2u single node configuration. Each system is engineered with robust enterprise data services that enable organizations to access data with patented parallel performance and optimize data placement and access on low cost media such as cloud or tape.

¹ Performance numbers based on max engineering numbers obtained from internal testing.

The IBM Storage Scale System is designed to be easy to order, easy to install, easy to upgrade, easy to scale, and most importantly easy to use with a lower cost customer experience.

Application Workloads for the Storage Scale System



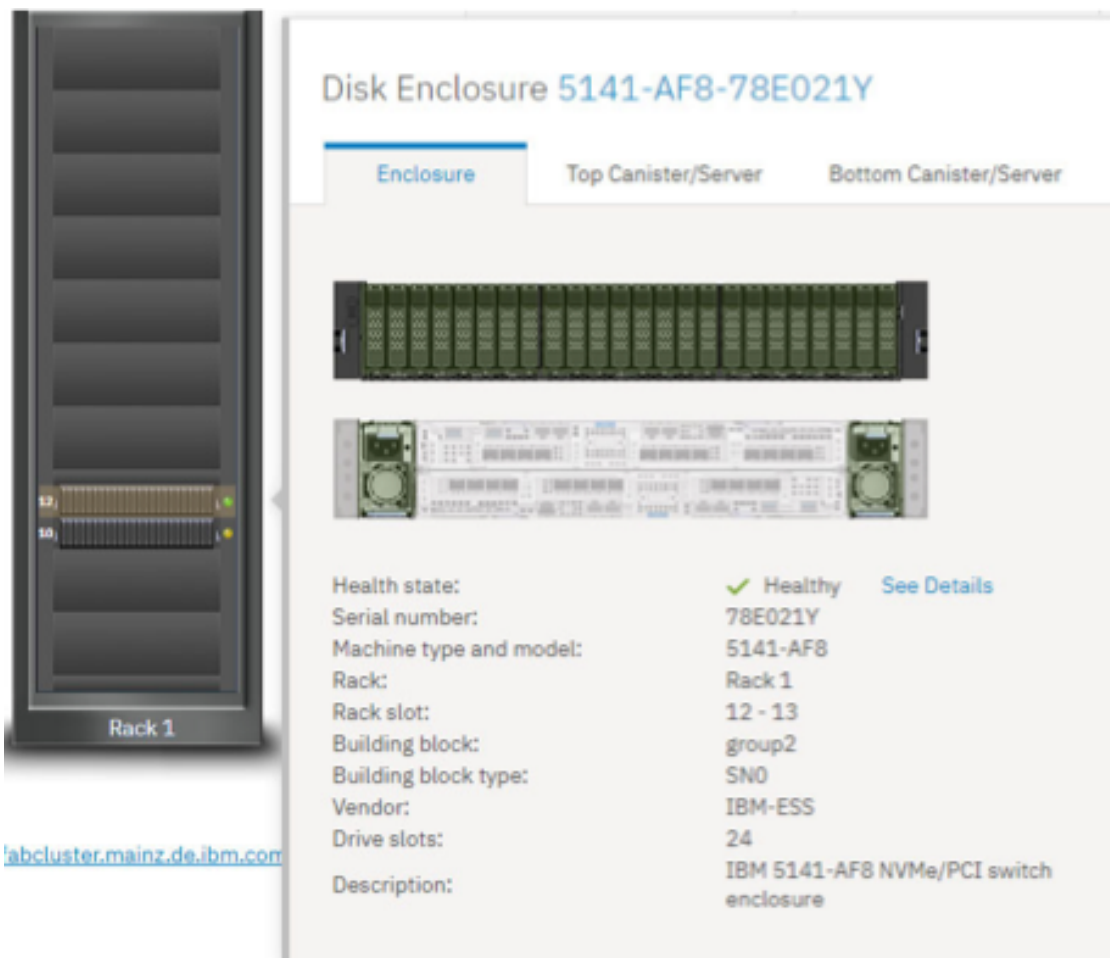
IBM Storage Scale System is not only easy but is powerful and provides an extreme high-performance tier of file and object storage including a turbo tier with up to 91GB/s of throughput and up to 16M IOPs, for a broad variety of applications. The Storage Scale System is designed to keep GPUs active to solve AI problems faster and running at peak performance. The system is also engineered to scale to 1000s of nodes and Yattabytes of capacity. Like all previous generations, IBM Storage Scale System runs the proven IBM Storage Scale RAID erasure coding, which provides superior consistent high-performance, mitigation of storage hardware failures, and intelligent monitoring / management / dynamic tuning of IBM Storage Scale System and IBM Storage Scale data.

IBM Storage Scale RAID manages the physical NVMe media. IBM Storage Scale RAID is a major differentiator in the Storage Scale System versus all of its competitors. Many competitors still need to rely on conventional disk storage technologies in order to provide storage hardware resilience and performance. As a result, competitive devices may have much lower usable capacity and access to data can be slow especially data that is not local to the storage node. IBM innovation optimizes the overall environment bringing much improved experience for the users and the applications.

IBM Storage Scale System is based on proven IBM Storage 2U24 hardware that can be expanded with additional storage enclosures upto 15.4PB of total capacity in a single node and 633YB in a single cluster. You can start with 48TB using half-populated flash nodes or create a fully-populated NVMe flash solution with 24, 2.5” drives in capacities of 3.84TB, 7.68TB, 15.36TB or 30TB. Using the largest capacity 30TB NVMe drives, up to 720TB total flash capacity, in a 2U form factor, along with associated low weight and low power consumption. Adding storage enclosures is easy as up to 8 enclosures (each 4u with 102 drives) can accommodate up to 816 drives of 10TB, 14TB or 18TB or 14.6PB of total raw HDD capacity.

Contact your IBM Sales representative or value added business partner for details about pricing and your specific configuration.

The simplicity of the Scale System



The image shows a server rack labeled 'Rack 1' with a callout window displaying the configuration and health status of a 'Disk Enclosure 5141-AF8-78E021Y'. The callout window has three tabs: 'Enclosure', 'Top Canister/Server', and 'Bottom Canister/Server'. The 'Enclosure' tab is selected, showing a top-down view of the enclosure and a front view of the enclosure. Below the views is a table of metadata.

Property	Value
Health state:	✓ Healthy See Details
Serial number:	78E021Y
Machine type and model:	5141-AF8
Rack:	Rack 1
Rack slot:	12 - 13
Building block:	group2
Building block type:	SNO
Vendor:	IBM-ESS
Drive slots:	24
Description:	IBM 5141-AF8 NVMe/PCI switch enclosure

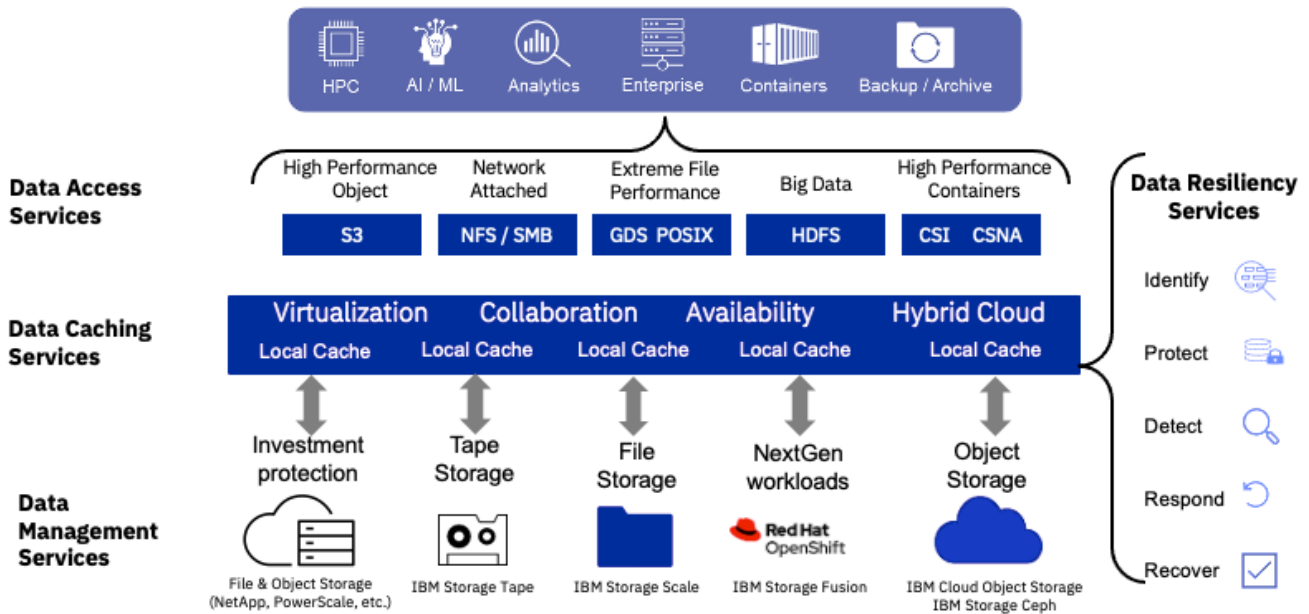
To address the challenges of managing today's data, IBM Storage Scale System delivers a new generation of software-defined storage. It builds on years of experience and couples proven IBM Storage Scale software with lightning-fast NVM storage technology to offer industry-leading file management capabilities and low latency data access. These build on and extend a track record of meeting the needs of the smartest, most demanding organizations. The Scale System is up to 300% faster than the previous generation of the Scale System and 26% lower power consumption making a great price performance leader.

Whatever the current scope of your data management needs, you can start as small or as big as you want and easily scale out by adding more units when needed. IBM Storage Scale System offers:

- Up to 300% better performance vs previous models and combines flash and capacity data with up to 15PB capacity in only 18u¹
- Investment protection by expanding an existing or build a new Global Data Platform and use current storage even if from another vendor
- Green sustainable data with 100% fewer nodes and 26% less power and enabling 3x more performance in a more consolidated node than previous generations
- Always-on upgrades and expansion with enhanced non disruptive upgrades for scale-up and scale-out

IBM Storage Scale System helps simplify ITops with an intuitive user interface that provides a picture and graphic view of problems should they arise. Operations can protect data directly to cloud or to a remote IBM Storage Scale System and then even archive to tape.

The IBM Global Data Platform



Our storage solutions are designed to optimize, secure and unlock customers data. We accomplish these goals with two software products and 4 differentiated data services. Our data services include:

Data Access Services with Multi-protocol Performance that connects directly to your applications.

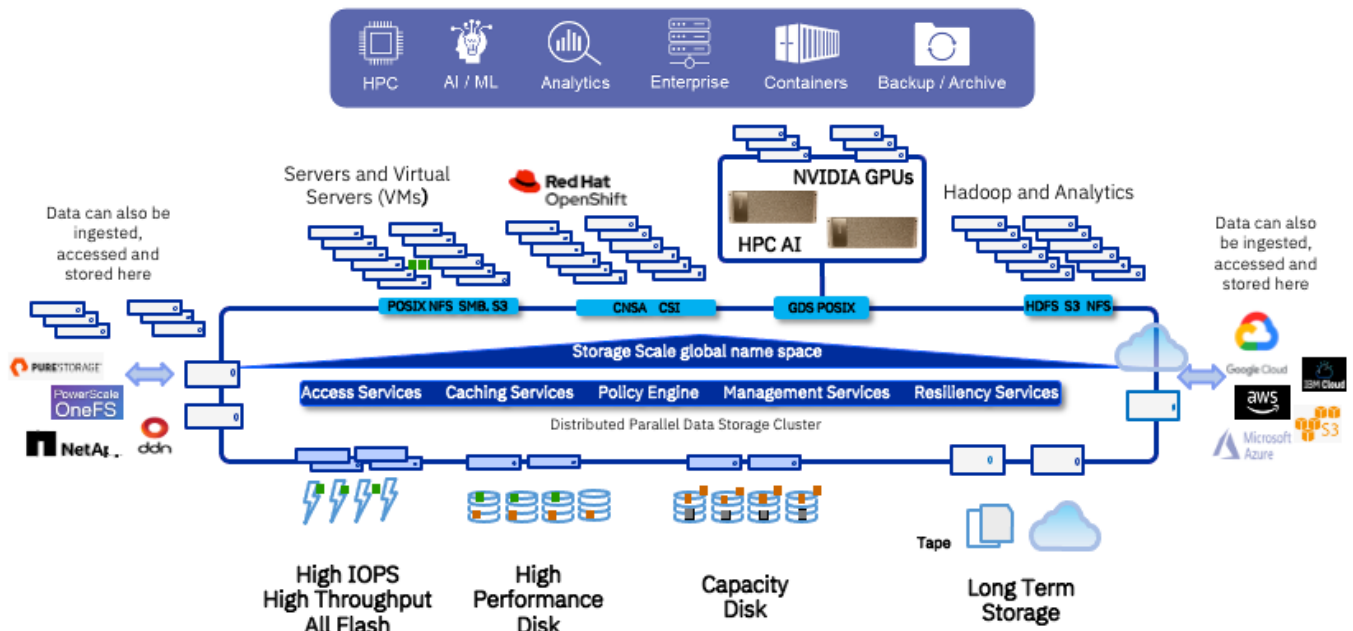
Data Caching or Core Services that provide Global Connectivity from multiple data sources and multiple locations to bring together data from IBM and non-IBM storage environments.

Data Management Services with policy automation that transparently helps you manage the flow of data and take much of the complexity out of data to day data management.

Data Resiliency Services that provide cyber secure automation to ensure your data is protected and safe and quickly recoverable when needed.

The Global Data Platform is composed of two software products, IBM Storage Scale and IBM Storage Ceph. IBM has engineered a solution that takes multiple data sources that can be used for multiple workloads from AI/ML and Analytics/HPC to any file and object or cloud native application and brings maximum value to the organization.

How a Global Data Platform Works




Specifications

Fully compatible with previous generations of IBM Storage Scale System (IBM ESS 3000, 3200, 5000 GLxS, GLxC, GSxS and GHxy) the system is capable of leveraging the same Scale System Management Servers, Protocol Nodes and Storage Scale Cluster/Name Space.

Embedded Red Hat Enterprise Linux 8.x operating system simplifies RHEL install, management and upgrade by removing the need to register with Red Hat to download errata and patches.

Standard three-year warranty, IBM on-site 9x5, next business day support. Optional upgrade for additional on-site.



2u Controller + NVMe Disks

4u HDD Disk Enclosure

Hot-swappable Adapters, power supply, fans and drives

System Features	<ul style="list-style-type: none"> • Dual 1-socket Storage Controllers, Active/Active • 1024 GB memory • De-Clustered RAID supporting erasure coding schemas: 3-way replication, 4-way replication, 4+2P, 4+3P, 8+2P, 8+3P
Software	<ul style="list-style-type: none"> • Storage Scale System software 6.1.4.0 • IBM Storage Scale for Storage Scale System 5.1.4 • Red Hat® Enterprise Linux® (RHEL) 8.4
Software Features	<ul style="list-style-type: none"> • Data Access Services with POSIX, GPU Direct, HDFS, NFS v4, SMB, HTTP, S3 • Data Management Services with integrated lifecycle management to optimize data from Memory, NVMe flash, HDD, public cloud, external storage and tape • Data Security Services with FIPS 140-3, 256 bit encryption and IBM Safeguarded Copy • Data Caching Services that include cloud and non IBM storage and multi-site asynchronous and synchronous replication options to 2 or 3 sites with concurrent parallel access

IBM Storage for File and Object

Data Sheet

Performance	<ul style="list-style-type: none"> • AMD 7642 48 core processor • Sequential read performance up to 91GB/s
Networking	<ul style="list-style-type: none"> • HDS 200Gb/s Infiniband®, up to 8 ports • 100Gb/s Ethernet, up to 8 ports • Two integrated networking ports (separation of MGMT and BMC)
Adapters	<ul style="list-style-type: none"> • Four x16 PCIe Gen4 adapter slots
Drive Support	<ul style="list-style-type: none"> • 12 or 24 NVMe SSDs (3.84TB, 7.68 TB, 15.36 TB or 30TB) • Up to 816 HDD drives (10TB, 14TB or 18TB)
Environmental Controller	<ul style="list-style-type: none"> • Input Voltage: 200-240V 50/60 Hz • Nominal Power: 1,350 W (empty); 2000 W (max) • Nominal Heat: 4,606 BTU/hr (empty); 6,825 BTU/hr (max) • Power Supplies: 2 hot swappable, redundant • Acoustical: 8.1 bels (idling or operating) • Environment Operating temperature (optical networking: 5°C to 32°C) ((copper networking: 5°C to 35°C) from 0 to 3048 m (0 to 10,000 ft); Above 900 m, de-rate maximum air temperature 1 degree per 300 m; 8%-80% humidity range
Environmental Enclosure	<ul style="list-style-type: none"> • Dual 1600W, 80+ Platinum • 200-240V AC input, auto ranging, 50-60Hz • 4 main enclosure fans, front-to-rear system cooling with zero-loss backflow prevention • 2 IO module fans • Dual PSUs with built-in fans • Fan speed tuning via system management • Operating Temperature: 5 to 40°C • Non-op Temperature: -40 to 60°C • Humidity: 5 to 85% relative humidity • Operating Altitude: 40°C @ 3,000ft • Sound Power: < 7.2Bels @ 23±2°C
Size / Weight 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 / Weight 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)

Why IBM?

Data matters. When planning high performance infrastructure for new or existing applications it's easy to focus on compute resources and applications without proper planning for the data that will drive the results for the applications. As customers strive to unlock IT, de risk IT and enable green IT, IBM storage is all about solving those difficult challenges. Our products are all about solving hard problems faster with data so that customers can react faster to competition and to market dynamics by accelerating insights and providing seamless access to data that is not limited by borders or data silos. Our story is not just about another storage product but is about innovation and a storage portfolio that is powered by our global data platform.

For Further Information

For further information on IBM Storage file and object products please contact your IBM Business Partner:

Logicalis

917669069 | marketing-es@es.logicalis.com

<https://www.es.logicalis.com/>

Download a trial of our software:

[Download Free Software](#)



Copyright IBM Corporation 2023 IBM Systems
Route 100
Somers, New York 10589 U.S.A.
Produced in the United States of America,
02/2023

IBM, the IBM logo, IBM Cloud, IBM Cloud Paks, IBM Elastic Storage System, Power, and Storage Scale are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom. InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®,

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.

This information is provided "as is" without warranty of any kind, express or implied, and is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this document. Nothing contained in this document is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.