


Introduction to IBM Storage Scale and Storage Scale System

Phillip Gerrard
Sagrika Arora
Nilton Carlos Dos-Santos
Jordi Caubet
Stieg Klein
Chris Maestas
Bis Purkayastha



Introduction

Purpose of this document

This IBM Redpaper publication provides an overview of the IBM Storage Scale System. These scalable, high-performance data and file management solutions are built on IBM Storage Scale technology. Providing reliability, performance, and scalability, IBM Storage Scale Systems can be implemented to support a wide range of diverse requirements.

The IBM Storage Scale System 6000 builds upon the innovations of the most recent Storage Scale System models. The Storage Scale System 6000 comprises two 2U canisters for an overall 4U footprint. Each of the Storage Scale System 6000 canisters provides eight Gen5 x16 Peripheral Component Interconnect Express (PCIe) slots.

This publication is designed to help you understand the solution and its architecture. It describes and provides an overview of the system components and tools that are used to order the best solution for your environment. This includes planning the installation and the integration of the solution into your environment and correctly maintaining your solution.

The solution is created from the following combination of physical and logical components:

- Hardware
- Operating system
- Storage
- Network
- Applications

Knowledge of the IBM Storage Scale System components is key when planning an environment.

Who should read this document

This paper is targeted toward technical professionals (consultants, technical support staff, IT Architects, and IT specialists) who are responsible for delivering cost-effective cloud services and providing AI and big data ready solutions.

You should have working knowledge of IBM Storage Scale and general familiarity with high-performance storage systems, networking, and Linux-based environments.

Section overview

Introducing the IBM Storage Scale System

Introduces the IBM Storage Scale System family of products, including their characteristics, value, and systems integration. Describes where solutions fit in today's business environments and provides an overview of IBM Storage Scale RAID software and the storage building block concepts of the solution.

IBM Storage Scale System architecture

Describes the architecture, hardware, and software of the IBM Storage Scale System, including the management server, utility nodes, I/O and data servers, and storage enclosures. Also covers the operating system, Ansible-based container software upgrades, IBM Storage Scale, and IBM Storage Scale RAID.

IBM Storage Scale System planning and integration

Provides guidelines and considerations for planning, installing, and configuring the IBM Storage Scale System. Covers roles and responsibilities, networking, installation, maintaining hardware and software, and upgrade procedures.

Appendix A: IBM Storage Scale System models

Describes the differences between the three generations of IBM Storage Scale Systems in use today, including detailed model specifications, IBM Storage Scale RAID erasure codes, and declustered RAID array technology.

Authors

This paper was produced by subject matter experts working with IBM Redbooks® (ITSO), Tucson Center.

Phillip Gerrard is a Project Leader for the International Technical Support Organization working out of Beaverton, Oregon. As part of IBM for over 20 years, he has authored and contributed to hundreds of technical documents published to IBM.com and worked directly with IBM's largest customers to resolve critical situations. As a team lead and Subject Matter Expert for the IBM Spectrum® Protect support team, he is experienced in leading and growing international teams of talented IBMers, developing and implementing team processes, creating and delivering education. Phillip holds a degree in computer science and business administration from Oregon State University.

Sagrika Arora is an IBM Storage Scale Technical Support Professional and has been focused providing her expertise to the IBM Storage Scale team for the last 3 years. She has worked extensively assisting clients across the globe with GPFS code, operational, performance, and networking issues they encounter. She has been responsible for guiding clients during critical and high priority situations to resolve complex issues with their storage and storage systems. Prior to IBM, Sagrika has worked as a system Admin in Wipro technologies for about 5 years where she has assisted customers in multiple regions to manage their servers.

Jordi Caubet is a Presales Storage Architect within IBM Technology, bringing more than two decades of expertise in high-performance computing (HPC) and advanced storage solutions. With a strong foundation in designing and delivering scalable, high-performance storage architectures, Jordi has helped organizations optimize data-intensive workloads across industries such as research, engineering, and enterprise IT. His background in HPC uniquely positions him to bridge the gap between compute demands and storage performance, ensuring that clients can accelerate innovation through resilient, efficient, and future-ready data platforms. At IBM, Jordi collaborates closely with customers to architect tailored storage strategies, leveraging deep technical knowledge and a passion for solving complex infrastructure challenges.

Nilton Carlos Dos-Santos is a Senior IT Specialist with more than 20 years of experience in several different roles. He joined IBM in 2008 and for the past 5 years he has been working in the IBM Storage space. He holds a bachelor's degree in computer science and a post-graduate degree in Computer Network Systems. Nilton has also co-authored other IBM Redbooks® publications, including IBM Software Defined Infrastructure for Big Data Analytics Workloads, IBM Platform Computing Solutions for High Performance and Technical Computing Workloads, IBM Redpaper publication IBM Platform Computing Cloud Services and Getting Started with Docker Enterprise Edition on IBM Z®.

Mike Kieran is a product marketer currently focusing on Storage for IBM watsonx.data. He has more than a decade of storage marketing experience at Pure Storage, NetApp, and Nimble Storage. Mike is the author of four books on digital color and the recipient of an Emmy for science and technology writing. He is an avid stargazer and astrophotographer.

Stieg Klein is a Senior Storage Technical Specialist in IBM Advanced Technology Group (ATG). He joined IBM in 2012 and is focused on IBM software-defined storage (SDS), including IBM Storage Scale System and IBM Storage Scale. He holds a Bachelor of Arts degree in Computer Science from the University of California, Berkeley.

Chris Maestas is the Chief Executive Architect for IBM File and Object Storage Solutions with over 25 years of experience deploying and designing IT systems for clients in various spaces. He has experience scaling performance and availability with various file systems technologies. He has developed benchmark frameworks to test out systems for reliability and to validate research performance data. He has also led global enablement sessions online and face to face, where he described how best to position mature technologies like IBM Storage Scale with emerging technologies in cloud, object, container, or AI spaces.

Bis Purkayastha is a Pre-sales Solution Architect with IBM's Technology Expert Labs, bringing over 20 years of specialized experience in enterprise file and HPC storage solutions. He has worked extensively across the storage industry, helping organizations design and implement scalable, high-performance storage architectures. Prior to IBM,

Bis served as a Storage Architect at NTT Data Services, where he led enterprise storage transformation initiatives for global clients. He also spent significant time at JP Morgan Chase as a NAS Storage Developer. Throughout his career, Bis has provided hands-on storage engineering and high-level architectural strategy, making him a trusted advisor to enterprise clients. Bis's expertise spans NAS, HPC storage, file systems, and enterprise storage design across financial services and IT consulting sectors.

Thanks to the following people for their contributions to this project

Larry Coyne - International Technical Support Organization

Rezaul Islam - IBM Systems

John Sing - IBM Global Sales

Thanks to the authors of prior editions

Dino Quintero, Puneet Chaudhary, Lee Nee Helgeson, Brian Herr, Steven Normann, Marcelo Ramos, Richard Rosenthal, Robert Simon, Marian Tomescu, Richard Wale

Feedback

This document reflects practical experience implementing these technologies in production environments. If you encounter issues not covered here, or if you have suggestions for improving the guidance, your feedback helps make this resource more valuable for the community.

[Contact IBM Redbooks](#)

Introducing the IBM Storage Scale System

This section introduces the IBM Storage Scale System family of products. This section includes their characteristics, value, and systems integration and describes where solutions fit in today's business environments.

The latest IBM Storage Scale System 6000 provides investment protection to expand or build a new Global Data Platform in the environment, by using either new or currently in use storage devices. The Scale Storage System 6000 extends and expands on the IBM Storage Scale System 3500 models features and capabilities. IBM Storage Scale System can scale up or out with different storage mediums in the environment and supports Nvidia's ConnectX-7 network interface card (NIC). With the appropriate transceivers and adapter cables the system is capable of 100 Gb, 200 Gb, or 400 Gb Ethernet and supports InfiniBand HDR, InfiniBand NDR connectivity, or a combination of both.

This section also describes some of the software and hardware characteristics of the IBM Storage Scale System, IBM Storage Scale Redundant Array of Independent Disks (RAID) software, and the storage building block concepts of the solution.

Note: In this IBM Redpaper publication, IBM Storage Scale System can refer to any of the following products:

- IBM Elastic Storage Server (the first and second generations of this solution)
- IBM Elastic Storage System which includes the following models:
 - IBM Elastic Storage System 5000 (IBM Storage Scale System 5000)
 - IBM Elastic Storage System 3500 (IBM Storage Scale System 3500)
 - IBM Elastic Storage System 3200 (IBM Storage Scale System 3200)
 - IBM Elastic Storage System 3000 (IBM Storage Scale System 3000)
 - IBM Storage Scale System 6000

For more information about all the current IBM Storage Scale System products, see IBM Storage Scale System documentation.

For more information, see the following IBM Redbooks publications:

- Implementation Guide for IBM Elastic Storage System 3000, SG24-8443
- Implementation Guide for IBM Elastic Storage System 5000, SG24-8498
- Implementation Guide for IBM Elastic Storage System 3200, SG24-8516
- Implementation Guide for IBM Elastic Storage System 3500, SG24-8538
- Implementation Guide for IBM Storage Scale System 6000, SG24-8566

The general architecture and capabilities of these solutions are similar. Where suitable, the differences between the models and generations are described.

IBM Storage Scale System

The IBM Storage Scale System is an integrated hardware and software storage solution for high-performance, reliable file and object storage. Storage Scale System 6000 is designed for massive scalability, supporting thousands of nodes and yottabytes of capacity. It uses 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 the use of containerized software that speeds and simplifies the maintenance process.




Figure: IBM Storage Scale end-to-end data strategy

IBM has engineered a global data platform to enable faster access to the right data in the right place at the right time for faster AI results




Figure: IBM Storage Scale with the Global Data Platform

Included in an IBM Storage Scale System is a pair of Storage Scale Data Servers, which are deployed as a fully integrated, tested storage building block solution (see the following figure). Instead of hardware-based disk RAID controllers, IBM Storage Scale System uses software-defined Scale Native RAID, providing de-clustered erasure code technology to deliver consistent high performance. This feature mitigates storage drive failures and RAID rebuild overhead and improves storage efficiency.




Figure: Third-generation IBM Storage Scale System integrated with IBM Storage Scale

IBM Storage Scale System with IBM Storage Scale RAID also provides superior data protection by reducing rebuild times to a fraction of the time required to recover using conventional hardware-based RAID controllers.

Data access to IBM Storage Scale System attached storage is available by using one of the following methods:

- IBM Storage Scale clients on user workstations, which provide a POSIX file system interface to data on IBM Storage Scale System.
- IBM Storage Scale clients on Nvidia GPU systems, which provide GPU Direct Storage through CUDA implementation, reducing latency by directly accessing storage from the GPU.
- IBM Storage Scale protocol nodes, which provide IBM Storage Scale System data access through Server Message Block (SMB), Network File System (NFS), or S3 object data access. In addition, Hadoop clusters can access IBM Storage Scale System data by using the IBM Storage Scale Hadoop Transparency Connector.
- IBM Storage Scale Container Native Storage Access (CNSA) or Container Storage Interface (CSI) drivers that run on Red Hat OpenShift or Kubernetes environments.

Multiple IBM Storage Scale System building blocks can also be integrated into an IBM Storage Scale cluster to suit the storage needs of the environment.

IBM Service supports IBM Storage Scale System as an integrated solution providing support for the full software and hardware stack.

IBM Storage Scale System and IBM Storage Scale

This section describes how the IBM Storage Scale System is part of an IBM Storage Scale cluster.

Within the IBM Storage Scale cluster, IBM Storage Scale System is an integrated IBM Storage Scale building block that reads and writes data for IBM Storage Scale users. IBM Storage Scale System combines storage hardware with IBM Storage Scale software to manage any data in storage.

The following sections include a brief overview of these topics:

- What is an IBM Storage Scale cluster
- What is an IBM Storage Scale Client and IBM Storage Scale Data Server
- How IBM Storage Scale System is an integrated IBM Storage Scale storage building block

This section includes a high-level IBM Storage Scale and IBM Storage Scale System solution overview. For more detailed information, see [IBM Storage Scale System Architecture](#).

IBM Storage Scale overview

IBM Storage Scale is IBM's strategic high-performance parallel file system, a shared storage platform for end-to-end collaborative common enterprise, data platform, big data analytics, and AI workflows .

How a Global Data Platform Works

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




Figure: IBM Storage Global Data Platform with IBM Storage Scale overview

IBM Storage Scale is designed to provide the following major value propositions:

- Simplified data management by supporting enterprise workflows on a single common enterprise data platform through various access methods to reduce data silos.
- A single global namespace that supports enterprise-level data over high-performance networks to accelerate your workloads to the maximum available bandwidth or POSIX transactions per second (File IOPS).
- Abstraction of intelligent automatic tiering of data between storage pools, externally to tape, to object-based resources, and to cloud resources. This delivers cost-effective storage economics by automatically managing and tiering data to different classes of storage including data targets on remote S3 or NFS storage that might already exist.

Although multiple types of IBM Storage Scale cluster configurations are available, the configuration into which IBM Storage Scale System is commonly deployed is the IBM Storage Scale Network Shared Disk (NSD) configuration, as shown in the following figure.

IBM Storage Scale System contains a pair of IBM Storage Scale NSD Data Servers, which are configured together as a tested, integrated, highly available (HA), and reliable IBM Storage Scale building block based solution.

As shown in the IBM Storage Scale cluster component overview, eight IBM Storage Scale nodes are workstations, servers, or users. Four nodes are IBM Storage Scale data server nodes. In this configuration, the user workstations are running the IBM Storage Scale client in this configuration. In the IBM Storage Scale clients and servers figure, the IBM Storage Scale client provides multi-threaded, highly parallel data access to the IBM Storage Scale Data Servers.

IBM Storage Scale System




Figure: IBM Storage Scale cluster component overview

IBM Storage Scale System




Figure: IBM Storage Scale clients and servers

The IBM Storage Scale client achieves high performance by performing simultaneous real-time parallel I/O to all IBM Storage Scale data servers, storage volumes, and NSDs simultaneously.

IBM Storage Scale clusters can grow by adding nodes, whether they are IBM Storage Scale clients or IBM Storage Scale data servers, as shown in the following figure.




Figure: IBM Storage Scale grows by adding scale-out nodes

An IBM Storage Scale cluster can provide 1–256 logical POSIX file systems to users and workstations. The IBM Storage Scale client provides the appearance of a mountable POSIX file system to the applications and users on the workstation where the IBM Storage Scale client is installed.

IBM Storage Scale users are unaware of the physical distribution of data in the IBM Storage Scale data server physical storage pools. The automatically balanced data distribution is seamlessly determined by the IBM Storage Scale policy engine at the time that the data is imported. The policy engine can also transparently move data from one storage pool to another storage pool while the data is accessed and active.

The IBM Storage Scale parallel file system provides an enterprise with the capability for data management over large amounts of data. It also performs constant auto-balance of workload and storage by equally distributing I/O and data within a storage pool or among different storage pools.

The preferred method of accessing IBM Storage Scale data is to install the IBM Storage Scale client on every workstation or server that accesses IBM Storage Scale data. The IBM Storage Scale client provides multiple threads and communication with data servers to provide high-performance parallel throughput. Concurrently, IBM Storage Scale manages full read/write data integrity between users who are working with the data in the file system.

Protocol nodes overview

As an optional part of the IBM Storage Scale solution, protocol nodes provide access to IBM Storage Scale data by using NFS, SMB, or object protocols without installing IBM Storage Scale software on a node.

The primary benefit of protocol nodes is that applications, workstations, and users that do not have the IBM Storage Scale client can still access IBM Storage Scale data through one of these configured protocols. IBM Storage Scale is often used as an enterprise data lake or central enterprise data repository.

Protocol nodes are IBM Storage Scale nodes that are designed, configured, and set up to provide the following benefits:

- NFS, SMB, or object protocols on the customer-facing side
- A full IBM Storage Scale client that provides parallel access to IBM Storage Scale data on the storage-facing side

The following figure shows an example of protocol nodes that were added to an IBM Storage Scale cluster.




Figure: IBM Storage Scale System protocol nodes

When IBM Storage Scale data is accessed through protocol nodes, the limiting factor on performance is the single-threaded NFS, SMB, or object protocols. These protocols cannot deliver the same highly parallelized performance of the IBM Storage Scale client. However, the value of accessing the IBM Storage Scale data through NFS, SMB, or object protocols is a flexible means for allowing enterprise-wide user access to data stored in an IBM Storage Scale System.

IBM Storage Scale System storage building block solution

IBM Storage Scale System with IBM Storage Scale is a pair of IBM Storage Scale Data servers that are cross-configured together into an integrated, highly available IBM Storage Scale building block. IBM Storage Scale System provides a fully tested, integrated, and supported IBM solution. You can use Storage Scale to deploy, manage, maintain, and expand the environment as needed, by using a Storage Scale building block approach.

Two IBM Storage Scale System models are deployed as a cluster to provide most of the IBM Storage Scale storage, as shown in the following figure.




Figure: IBM Storage Scale Systems as IBM Storage Scale cluster building blocks

IBM Storage Scale System can coexist with other IBM Storage Scale data servers in the same IBM Storage Scale cluster. IBM Storage Scale System provides a powerful, flexible, fully integrated, and supported option for deploying IBM Storage Scale storage in the most efficient manner possible.

An IBM Storage Scale System storage building block is designed as a unit of storage expansion for an IBM Storage Scale cluster. If more storage is needed, an IBM Storage Scale System I/O enclosure or additional IBM Storage Scale System building blocks can be dynamically added to expand an IBM Storage Scale cluster.

To provide high reliability and consistent high performance, IBM Storage Scale System runs IBM Storage Scale native RAID (erasure coding), which is designed to maintain high performance even while transparently recovering from storage media or storage data server failures. IBM Storage Scale native RAID also provides checksum and disk hospital functions to detect and correct silent data corruption issues and maintain high availability (HA) of exabyte scale file systems.

As part of an IBM Storage Scale System order, you can also order IBM provided servers to use as IBM Storage Scale protocol nodes. These protocol nodes are managed by the IBM Storage Scale System solution software stack and the IBM Storage Scale System GUI. When including IBM Storage Scale System and IBM Storage Scale protocol nodes this provides an integrated solution for an IBM Storage Scale cluster.

For more information about planning for the usage of protocol nodes, see [IBM Storage Scale System Utility Node and Protocol Node](#).

IBM Storage Scale and Storage Scale System solution value

The following sections describe why the IBM Storage Scale and IBM Storage Scale System bring advantages to the business.

IBM Storage Scale advantages

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.

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

IBM Storage Scale System advantages

Today's AI and big data applications require large, high-performance, manageable, and flexible storage to be successful. When more data is generated, datasets grow and new language models are created, it is also important to be able to allow that storage to grow with the needs of the enterprise.

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 use AI to unlock the value of their data wherever it resides. IT leaders face key challenges:

- Accessing and analyzing data and workloads that are scattered across the globe
- Managing the growing AI infrastructure and ensuring scalability for evolving workloads
- The increasing time that is 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.

The IBM Storage Scale System is designed to provide an integrated and tested IBM Storage Scale storage building block to meet these growing storage needs. Incorporating IBM Storage Scale System storage servers flexibly adds to the overall IBM Storage Scale capacity, bandwidth, and performance, all within a single global namespace.




Figure: IBM Storage Scale System integrated solution

As shown in the previous IBM Storage Scale System integrated solution diagram, IBM Storage Scale System is an integrated IBM Storage Scale solution that significantly accelerates the time-to-value of deploying IBM Storage Scale environments. IBM Storage Scale System provides the following benefits:

- An integrated and tested IBM Storage Scale storage building block
- Predictable, consistent high performance at any level of scale
- Tested and integrated with the functions of IBM Storage Scale
- Supported as an integrated solution worldwide by IBM Service and Support

For more information about the IBM Storage Scale System, see the following web pages:

- [IBM Storage Scale System product page](#)
- [IBM Storage Scale System documentation](#)

Integrated and reliable storage building block

Every part of the IBM Storage Scale System solution is engineered to provide high-performance enterprise-wide data availability and reliability, from an entry-level system up to massive scale multi-node configurations.

IBM engineering and testing teams work cohesively on designing, building, testing, and delivering an end-to-end IBM Storage Scale System solution that is integrated and tested for reliability. The complete IBM Storage Scale System solution is again verified during manufacturing.

In the deployment stage, IBM Storage Scale System-specific tools and scripts are used for deployment. IBM Systems Lab Services are available and recommended to provide installation of the IBM Storage Scale System solution and to assist you in integrating IBM Storage Scale System into your client environment.

IBM Storage Scale System provides optimized storage configurations that include the following benefits:

- Optimum adapters that are strategically placed on the servers
- HA and redundancy for disks, drawers, and adapters
- Optimal cabling performance
- Tested and integrated firmware and software versions
- Policy-managed optimized placement of data
- HA access to data
- Automated storage management

IBM Storage Scale System is capable of 100 GB, 200 GB, or 400 GB Ethernet or InfiniBand HDR 200, InfiniBand NDR 200/400 connectivity, or a combination of both.

Predictable and consistent high performance

IBM Storage Scale System uses IBM Storage Scale RAID, which is a de-clustered RAID erasure code technology that recovers from multiple disk failures in minutes, versus hours or days with older technology RAIDs. This reduced recovery time gives the IBM Storage Scale System solution predictable performance and data protection, which includes 8+2 and 8+3 RAID protection and platter-to-client data protection.

A high-performance parallel file system environment, such as IBM Storage Scale, can run only as fast as its slowest component. IBM Storage Scale System is designed to use IBM Storage Scale RAID to provide consistent high performance by masking and mitigating performance effects of storage hardware failures, even if multiple drive failures occur. In this way, IBM Storage Scale System helps ensure that the larger IBM Storage Scale parallel file system always runs at optimum efficiency and consistently delivers expected high performance, even if storage hardware or media failures occur.

This new version provides deeper flash capacity by increasing the number of QLC in the Scale System 6000. Starting with IBM Storage Scale 7.0.0.0, data acceleration tier (DAT) is supported on a fully populated IBM Storage Scale System 6000 with performance configuration (48 TLC NVMe drives). DAT for IBM Storage Scale System 6000 addresses the IOPS performance issues that constrain real-time AI inferencing.

As shown in the following Data Acceleration Tier (DAT) overview diagram, DAT uses the IBM Storage Scale System technology called asymmetric data replication, which provides a performance data replica for fast reads and a GNR data replica for data resiliency.




Figure: Data Acceleration Tier (DAT) overview

Combining QLC and TLC

AI initiatives, analytics platforms, content repositories, backup environments, and long-term retention requirements are driving storage growth that often outpaces infrastructure budgets, accelerating adoption of quad-level cell (QLC) flash.

While triple-level cell (TLC) flash remains the preferred choice for write-intensive applications, QLC is emerging as a compelling option for capacity-driven workloads because it delivers higher density and lower cost per gigabyte. Modern NVMe QLC technology has also closed much of the performance gap that once limited adoption. IBM Storage Scale System helps organizations realize the density and cost advantages of QLC through intelligent tiering and workload-aware data placement.

GPFS Native RAID (GNR) understands the characteristics of both TLC and QLC media and uses that awareness to make smarter data placement decisions. Rather than treating all flash storage identically, GNR can direct small random writes to TLC while reserving QLC for larger sequential operations better aligned with high-density flash. Metadata-intensive operations can also be placed on the most appropriate tier, helping maintain performance while reducing unnecessary wear on QLC devices. By reducing write amplification and aligning workloads with the right media type, IBM Storage Scale System helps organizations achieve the cost advantages of QLC without sacrificing reliability. Rather than forcing organizations to choose between performance and economics, IBM Storage Scale System is designed to optimize both within a single architecture.

GNR also incorporates workload-aware data placement strategies that distribute activity evenly across drives. By avoiding workload skew, the platform helps maintain consistent performance while promoting more uniform drive aging. Background validation and repair processes further improve reliability by identifying and correcting potential issues before they affect production workloads.

IBM Storage Scale software continuously monitors latency, throughput, endurance metrics, spare capacity, and overall SSD health. The platform also evaluates drive behavior to identify underperforming devices and emerging health concerns before they become operational problems. Combined with predictive monitoring, automated remediation, and failure-domain-aware data placement, these capabilities provide the resiliency required for enterprise-scale deployments.

The architecture also provides a flexible path for scaling capacity as requirements grow. Organizations can increase the proportion of QLC storage to maximize economics, maintain a larger TLC footprint for performance-sensitive

environments, or combine both approaches within a single deployment. This flexibility allows infrastructure teams to optimize for cost, performance, or a balance of both as business requirements evolve.

The result is a platform capable of delivering flash storage at a scale that would have been difficult to justify economically using TLC alone. In large-scale deployments, a single rack can provide more than 47 PB of raw flash capacity, enabling organizations to support rapidly growing AI and analytics environments without dramatically increasing storage costs. Configurations leveraging QLC NVMe flash enable up to 4.79 PB of total system capacity, more than three times the capacity of using only TLC flash.

Simplified IBM Storage Scale System installation

IBM Storage Scale System technology deployment methods use Red Hat Ansible driven container orchestration, which results in faster installation, upgrades, with less effort and skills that are required for deployment and maintenance, as shown in the following figure.




Figure: IBM Storage Scale GUI-based update container overview

As shown in the previous simplified installation diagram, IBM Storage Scale System provides a fully tested and integrated solution that is composed of Red Hat Enterprise Linux (RHEL), IBM Storage Scale, and IBM Storage Scale System integration code. Installation modules are done by using Ansible-driven container orchestration and Ansible playbook bundles (APBs). This creates a standard, fast, and manageable mechanism for faster automated deployment of IBM Storage Scale System.

For version 7.0, you can use the GUI to upgrade the IBM Storage Scale System, which is typically easier than using a series of commands. See [Setting up the IBM Storage Scale System API server container](#).




Figure: Current generation IBM Storage Scale System models

The use of the GUI to perform the upgrade includes the following steps:

- Initialize the deployment container
- Update the I/O nodes
- Update the management server virtual machine and bare metal
- Update the firmware

For more information, including prerequisites, see [Setting up the IBM Storage Scale System API server container](#).

Supported as a solution

The IBM Storage Scale System is supported as a solution. Clients who engage IBM Service and Support for assistance on IBM Storage Scale System receive support from a dedicated IBM team that is cross-trained on all IBM Storage Scale System (hardware and software) components for the solution.

This specialized IBM Storage Scale System Service and Support team can debug and advise on all aspects of IBM Storage Scale System.

If more detailed support is needed for an IBM Storage Scale System component, IBM Storage Scale System service and support records show and document the problem resolution progress for the overall IBM Service and Support team. The resulting fast time to resolution contributes to the added value of purchasing an integrated, tested IBM Storage Scale solution from IBM.

For more information about IBM service and support of the IBM Storage Scale System, see [IBM Storage Scale System Support Reference Guide](#) (IBMid log in required).

IBM Storage Scale System models

IBM Storage Scale System is an IBM Storage Scale storage building block that is available in various models. These models offer a flexible portfolio of IBM Storage Scale storage building block options with NVMe or hard disk drive (HDD) storage.

The general physical arrangement for an IBM Storage Scale System is an IBM rack that features the following components:

- A pair of IBM Storage Scale I/O Data Servers
- Storage enclosures
- Network switches

A timeline of some of the models of IBM Storage Scale System is shown in the following figure.




Figure: IBM Storage Scale System (SSS) current generation models timeline

IBM Storage Scale System architecture

This section describes the architecture, hardware, and software of the IBM Storage Scale System, including the following components:

- IBM Storage Scale Management Server (IBM EMS)
- IBM Storage Scale Utility Node (Services for Management, Protocol, Active file Management (AFM), Protect and in the future Data Cataloging)
- I/O and data servers
- Storage enclosures

This section also describes the following building block models, software components, and features:

- Operating system
- Ansible playbooks and container-based software upgrades
- IBM Storage Scale
- IBM Storage Scale RAID

Note: IBM Storage Scale System is an integrated packaged solution. Users should not install different kernel levels or drivers into the IBM Storage Scale System. The user should also not run any non-IBM

Storage Scale System (client) application or non IBM Storage Scale System workloads on the I/O nodes or IBM Storage Scale System Utility Node Servers.

Architecture overview

An IBM Storage Scale System is a fully integrated solution that combines hardware and software, designed and tested to operate as a single, unified system. The major components of an IBM Storage Scale System solution release include the following items:

- Server hardware
- Storage hardware
- IBM Storage Scale software including IBM Storage Scale RAID
- Embedded Red Hat Enterprise Linux (RHEL) operating system

The minor components, which are integration-tested with the major components, include (but not limited to) the following items:

- Server firmware
- Server host bus adapters (HBAs) and related firmware
- Nvidia network adapter, drivers, firmware, and network switches
- Storage enclosure and drive firmware
- IBM racks, power distribution units (PDUs), and cabling management

In addition, many of the hardware and software components of the IBM Storage Scale System solution communicate with each other over an IP network. IBM Storage Scale System requires all networking connections to be in place and provisioned to install the system. This IP network can use Nvidia network switches, which are tested with IBM Storage Scale System by IBM, or the network switches can be provided by the client. These IP network components consist of the following switches:

- Low-speed 1 Gb and 10 Gb Ethernet (GbE) network switches for management and service networks
- High-speed network switches for data read and write over a high-speed data network

IBM Storage Scale System hardware components

This section includes a description of the following hardware components:

- Solution models:
 - 6000
 - 3500
 - Server hardware

Current generation IBM Storage Scale System 6000 and 3500 use x86 based servers for storage.

- Storage enclosures
- Network interface cards (NICs) and network switches
- Rack

IBM Storage Scale Systems are a rack-mounted storage solution. The initial storage building block in an IBM Storage Scale System cluster includes the following components:

- One IBM Storage Scale Management Server: Required for every IBM Storage Scale cluster
- One or more IBM Storage Scale System, each of which includes redundant I/O data servers

Current generation IBM Storage Scale System models include various models that provide NVMe, serial-attached SCSI (SAS), or Serial Advanced Technology Attachment (SATA) hard disk drives (HDDs) in selected sizes.

The IBM Storage Scale System 6000 was announced and delivered in 4Q2023. The IBM Storage Scale System 6000 design builds on the IBM Storage Scale System 3500, with more powerful processors and additional PCIe Gen 5 x 16 lanes that are both faster and wider than those supported by the 3500. Additionally, several design updates improve serviceability and provide additional system configuration flexibility, such as an optional campus connection.

The IBM Storage Scale System 3500 was announced and delivered in 2Q2022. The IBM Storage Scale System 3500 design is based on prior IBM Storage Scale System models. The IBM Storage Scale System 3500C/H models support up to eight external storage enclosures.

Available IBM Storage Scale System models

At the time of writing, two IBM Storage Scale System models are available to order from IBM.

IBM Storage Scale System 6000

The IBM Storage Scale System 6000 machine type 5149-F48 for the 4U48 form factor includes the following features:

- Embedded NVMe TLC drive options with a choice of 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB TLC Gen5 drives.
- Embedded mixed NVMe TLC + QLC drive options with a choice of 30.72 TB, 61.44 TB, or 122.8 TB QLC drives.
- Embedded NVMe options IBM FlashCore® Modules with 38.4 TB drives. IBM FlashCore Modules have inline compression enabled.
- Optionally extended with HDD expansion enclosures (MTM 5149-091) supporting 91 SAS HDDs per enclosure with a choice of 12TB, 16 TB, 20 TB, 22 TB, or 24 TB Self-Encrypting Drives (SED) and up to 9 expansion enclosures per system.
- Optionally extended with all-flash expansion enclosures supporting 26 NVMe QLC drives per enclosure with a choice of 30.72 TB, 61.44 TB or 122.8 TB drives and up to 4 expansion enclosures per system.

IBM Storage Scale System 3500

The IBM Storage Scale System 3500 machine type 5141-FN2 for the 2U24 form factor includes the following features:

- Embedded NVMe TLC drive options with a choice of 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB TLC Gen5 drives.
- Optionally extended with up to 8 HDD expansion enclosures (MTM 5147-102). Each supports 102 SAS HDDs per enclosure with a choice of 10 TB, 14 TB, 18 TB, 20 TB, 22 TB, or 24 TB, with support for Self-Encrypting Drives (SED) capability.

For more information about previous IBM Storage Scale System models, see the IBM Redbooks publication

IBM Storage Scale System 6000

IBM Storage Scale System 6000 is an all-flash, integrated, end-to-end NVMe storage solution, providing industry-leading performance, scalability, and unrivaled data protection and availability. The 6000 provides twice the performance and storage capacity relative to the previous generation systems and is the most sustainable storage solution to date.

IBM Storage Scale System 6000 is a high performance, accelerated storage offering optimized for data intensive workloads for advanced AI, analytics, and high-performance computing. The 6000 delivers significant advancements to the data services provided by the global data platform.

The Storage Scale System 6000 uses the power of Storage Scale combined with NVMe and flash technology in an x86 PCIe Gen5 4U48 form factor to deliver high-performance storage for AI, data analytics, and any high-performance file and object use cases. With support for IBM Flash Core Module technology, the 6000 supports up to 1.8 PB of raw flash capacity and

5.5 PB of compressed flash storage in a single 4U building block, as shown in the following figure.



Figure: IBM Storage Scale System 6000 front view



Figure: IBM Storage Scale System 6000 rear view

The Storage Scale System 6000 can deliver more than 300 gigabytes per second (GBps) of performance with low latency to optimize all types of I/O patterns and data layouts. This performance delivers data to applications faster and helps ensure maximum GPU resource usage. Performance can vary because of disk technology, high-speed network adapters, workload, and configuration. To estimate performance for your environment, contact an IBM representative or IBM Business Partner.

The Storage Scale System 6000 supports up to 16 ports of 100 Gb RDMA over Converged Ethernet (RoCE), InfiniBand HDR 200, InfiniBand NDR 200/400, or a combination of these.

The two redundant high-performance controllers in a 2U platform provide high availability and reliability. An embedded baseboard management controller (BMC) with remote control capability allows administrators to manage the system remotely.

IBM Storage Scale System is an accelerated storage offering optimized for data intensive workloads for advanced AI, analytics, and high-performance computing.

IBM Storage Scale System 6000 can be an ideal storage platform for the following applications:

- Compute intensive AI model training and inferencing for GPU-accelerated workloads requiring low latency, and high OPS
- Accelerated AI, delivering performance, scalability, data virtualization, and data resiliency for IBM watsonx.data® and IBM watsonx.ai®
- Industry-leading¹ throughput for data intensive high-performance computing solutions

IBM Storage Scale System is designed so that you can break data access barriers and create new storage dynamics for Nvidia AI solutions. With IBM Storage for data and AI, you can create a global data platform with top-rated

performance to accelerate data access greater than 125 GB/s per node with parallel data access for GPU-accelerated computing with NVIDIA DGX systems.

IBM Storage Scale System 6000 Performance Model

IBM Storage Scale System 6000 Performance Model refers to the machine type 5149-F48 with no external enclosures. IBM Storage Scale System 6000 control enclosure provides 48x NVMe slots at the front redundantly accessible by the two I/O servers at the back of the enclosure.

A system can be configured with 24 drives in a half-populated configuration or 48 drives in a fully populated configuration. Two NVMe drive options can be configured:

- NVMe TLC drives with a choice of 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB TLC Gen5 drives.
- IBM FlashCore® Modules with a choice of 38.4 TB drives. IBM FlashCore Modules have inline compression enabled.




Figure: IBM Storage Scale System 6000 Performance Model options with TLC or FCM

IBM Storage Scale System 7.0.1 release introduces a 12-drives quarter-populated configuration providing a small starter configuration.




Figure: IBM Storage Scale System 6000 Performance model 12 drives option with

The following table provides an overview of the IBM Storage Scale System 6000 server canister specifications.

Number of DIMMs Per Server Canister	Total Memory Per Server Canister	Total Raw Capacity Per ISS6k unit	Server Canister Features
8, 16 or 24 (32 GiB DIMM)	Up to 768 GiB	Up to 1470 TB NVMe TLC / Up to 1800 TB IBM FlashCore / Up to 4978 TB mix NVMe TLC and QLC	Dual socket AMD EPYC 9454 SP5, 48 core processors.
8, 16 or 24 (64 GiB DIMM)	Up to 1536 GiB	Up to 1470 TB NVMe TLC / Up to 1800 TB IBM FlashCore / Up to 4978 TB mix NVMe TLC and QLC	Dual 960 GB NVMe boot drives.
8, 16 or 24 (128 GiB DIMM)	Up to 3072 GiB	Up to 1470 TB NVMe TLC / Up to 1800 TB IBM FlashCore / Up to 4978 TB mix NVMe TLC and QLC	Dual 960 GB NVMe boot drives.

- Although one third and two third memory configurations are possible, number of drives - type and size - and number of enclosures will force a minimum memory footprint

IBM Storage Scale System version 6.2.3 introduces the capability to mix TLC drives and QLC drives. A system can be configured with 24 TLC drives and 24 QLC drives with a tiered flash storage configuration in the same enclosure. IBM Storage Scale System 7.0.0 release includes support so that you can configure 12x TLC drives and 36x QLC drives. All drives on each tier must have the same capacity. QLC drive options include 30.72 TB, 61.44 TB, or 122.8 TB. Support for the 122.8 TB drive was introduced with version 7.0.0.¹




Figure: IBM Storage Scale System 6000 Performance model options with mixed TLC and QLC

By default, IBM Storage Scale System 6000 Performance Model can configure up to four adapter slots for use with high-speed network cards. The four adapter slots that are used are located in the middle of the device, as shown in the following figure. The other adapter slots are kept empty for future MES to connect external storage expansion enclosures.




Figure: IBM Storage Scale System 6000 Performance Model high-speed network adapter slots

IBM Storage Scale System 6000 with HDD storage expansion enclosures (MTM 5149-091)

IBM Storage Scale System 6000 (machine type 5149-F48) can include external HDD storage expansion enclosures (5149-091). A system that has internal NVMe drives plus storage expansion enclosures is a Hybrid Model. The number of internal NVMe drives can be 24 drives for half populated or 48 drives for fully populated. The system that has only storage expansion enclosures is a Capacity Model and has no internal drives for data storage. The Capacity Model includes internal data drives that are used exclusively for log tip and other internal IBM Storage Scale usage but are not used to store data.

At the time of writing, you can add storage expansion enclosures in only nonmixed NVMe configurations. For instance, you can add them in combination with NVMe TLC or IBM FlashCore modules but not in combination with mixed TLC+QLC configurations.

A system can include up to 9x external expansion storage enclosures. Each enclosure includes 91 SAS HDDs with a choice of 12 TB, 16 TB, 20 TB, 22 TB, or 24 TB. All enclosure drive sizes in all enclosures must be the same. Using the MES process, you can add enclosures that are fully populated with the same existing drive size.



Figure: IBM Storage Scale System 6000 external HDD storage expansion enclosure front view

External HDD storage enclosures are connected by using SAS adapters that are installed on the four available PCI slots that are not used for high-speed network adapters. You can use the four SAS adapters in each server to

connect up to 8 enclosures. To connect a ninth enclosure, six SAS adapters are required, which leaves room for only two high-speed network adapters.




Figure: IBM Storage Scale System 6000 with HDD storage expansion enclosures

IBM Storage Scale System 6000 with JBOF storage expansion enclosures (MTM 5149-F26)

IBM Storage Scale System version 7.0.0 introduces support for external Flash storage expansion enclosures (5149-F26). You can add these enclosures to only NVMe TLC-only configurations. Systems with IBM FlashCore modules or with mixed NVMe TLC and QLC drives do not support external Flash storage expansion enclosures.

Each 5149-F26 enclosure includes 26 NVMe QLC drives in one of three capacities: 30.72 TB, 61.44 TB, or 122.8 TB. Each enclosure ships fully populated with drives. You can attach up to four 5149-F26 enclosures per system. A system can use either HDD or JBOF external enclosures, but not both types on the same system.




Figure: IBM Storage Scale System 6000 JBOF enclosure front and rear view

External JBOF storage enclosures are connected by using NVMe over Fabrics (NVMe-oF) in direct-attached mode to the I/O canisters. In the four available PCI slots that are not used for high-speed network adapters, the system installs high-speed network adapter cards for such direct-attached connectivity instead of SAS adapters when you use HDD enclosures. These adapters, even if they are the same adapters used in high-speed adapter slots, are used only for backend connectivity, meaning direct attachment of I/O canisters to JBOF enclosures.




Figure: IBM Storage Scale System 6000 adapter slot allocation for JBOF enclosures

An existing Performance Model with TLC drives (24 or 48) can add 1, 2, 3, or 4 JBOF enclosures or it can be configured as an initial system. At the time of writing this Redbooks publication, the MES process to increase JBOF enclosures is not supported. However, always look to the latest documentation for further updates.




Figure: IBM Storage Scale System 3500 front view

IBM Storage Scale System 3500

IBM Storage Scale System 3500 is an all-flash, integrated, end-to-end NVMe storage solution that provides industry-leading performance, scalability, data protection, and availability.



Figure: IBM Storage Scale System 3500 rear view



Figure: IBM Storage Scale System 3500 models and expansion example

The Storage Scale System 3500 can provide over 125 gigabytes per second (GBps) of performance with low latency across all types of I/O patterns and data layouts. This delivers data to applications faster and helps ensure maximum GPU resource usage. Performance varies with disk technology and high-speed network adapters. For more information about estimating and sizing your performance, contact your IBM representative or IBM Business Partner representative.

The Storage Scale System 3500 supports up to 8 ports of 100 Gb RDMA over Converged Ethernet (RoCE), InfiniBand HDR 200, InfiniBand NDR 200, or a combination of both. Two redundant high-performance controllers in a 2U platform support high availability and reliability.

An embedded baseboard management controller (BMC) provides remote management capability.

IBM Storage Scale System is designed help you create new storage dynamics for Nvidia AI solutions. With IBM Storage for data and AI, you can create a global data platform with data access rates greater than 125 GB/s per

node. By using the Storage Scale System 3500, you can achieve these speeds by using parallel data access for GPU-accelerated computing with high-performance, local NVMe storage and SAS-connected capacity storage.

IBM Storage Scale System 3500 can include 12 or 24 NVMe drives of 3.84 TB, 7.68 TB, 15.36 TB, or 30.74 TB for user data. For more information, see [IBM Storage Scale System 3500](#).

IBM Storage Scale 3500 Capacity and Hybrid models can include 1–8 storage shelves (MTM 5147-102) of NL-SAS drives. When only one storage shelf is used, it can contain 52 or 102 drives. After expanding the first shelf to 102 drives, extra shelves that contain 102 drives can be added. Supported drive sizes are 10 TB, 14 TB, 18 TB, 20 TB, 22 TB, or 24 TB. For more information, see [IBM Storage Scale System 5147-102 Storage Enclosure](#).

IBM Storage Scale System 3500 supports nondisruptive capacity upgrades including the addition of up to a maximum of 8 storage shelves. For each shelf, all HDDs must be of equivalent size and performance.




Figure: IBM Storage Scale System 3500 Performance model options with TLC

IBM Storage Scale System 3500 Performance Model

IBM Storage Scale System 3500 Performance Model refers to the machine type 5141-FN2 with no external enclosures. The IBM Storage Scale System 3500 control enclosure provides 24 NVMe slots that are redundantly accessible by the two I/O servers at the back of the enclosure.

A system can be configured with 12 drives in a half-populated configuration or 24 drives in a fully populated configuration. The supported NVMe drives are 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB TLC Gen5 drives. On the Storage Scale System 3500, Gen5 drives operate as Gen4 drives because the internal PCI bus is Gen4.




Figure: IBM Storage Scale System 3500 external HDD storage expansion enclosure front view

IBM Storage Scale System 3500 with HDD storage expansion enclosures (MTM 5147-102)

IBM Storage Scale System 3500 (machine type 5141-FN2) can include external HDD storage expansion enclosures (5147-102). A system that has 12 or 24 NVMe drive and storage expansion enclosures is a Hybrid Model. The system that has only storage expansion enclosures is a Capacity Model. The Capacity Model includes four NVMe drives that are used exclusively for log tip and other internal IBM Storage Scale usage.

A system can include up to 8 external expansion storage enclosures, as shown in the preceding figures. Each enclosure includes 102 SAS HDDs with a choice of 10 TB, 14 TB, 18 TB, 20 TB, 22 TB and 24 TB. For each shelf or enclosure, all drives must have the same capacity and speed. By using the MES process, it is possible to add enclosures that are fully populated with the same drive size.



Figure: IBM Storage Scale System 3500 with HDD storage expansion enclosures

External HDD storage enclosures are connected by using SAS adapters that are installed in the two available PCI slots that are not used for high-speed network adapters. Two SAS adapters in each server provide the connection to up to 4 enclosures. You can connect a maximum of 8 enclosures that are connected by using daisy chain connections. After a 4th enclosure is added, any additional enclosures add only capacity, not additional performance.




Figure: General positioning of available IBM Storage Scale System models

Positioning various IBM Storage Scale System models

IBM Storage Scale System building blocks and models

Each building block contains:

- A pair of IBM Storage Scale System I/O NSD data servers
 - IBM Storage Scale System 6000, a fully integrated 4U storage building block containing 24 or 48 NVMe Flash drives providing performance up to 340GB/s and 13M IOPs.
 - IBM Storage Scale System 3500, a fully integrated 2U storage building block with the ability to add up to 8@4U storage enclosures.
- Uses a IBM Storage Scale System Utility Node server (one per cluster) which can be configured as a Management Node and/or Protocol Node.

IBM Storage Scale System has various models and configurations to suit data demands

- Storage media types: NVMe Flash, HL-SAS HDD
- Multiple options for NVMe and HDD disk size based on throughput and total storage capacity needs




Figure: IBM Storage Scale System building blocks and models overview

IBM Storage Scale System models are available in different categories. The x in the models denotes the number of storage enclosures in that model.

Consider the following points:

- IBM Storage Scale System 6000 models are built on NVMe flash storage for high performance. Each IBM Storage Scale System 6000 canister supports up to eight PCIe5 cards. PCI cards are able to be used for either networking or HBA attachment for external storage expansion. A single NIC supports up to two high-speed networking interfaces. A single HBA supports up to two high-capacity external storage enclosures.
- IBM Storage Scale System 3500 models are built on NVMe flash storage for high performance, and optionally high capacity 4U102 HDD storage enclosures. Each IBM Storage Scale System 3500 canister supports up to four PCIe4 cards. A single card either supports up to two high-speed networking interfaces or up to two high-capacity storage enclosures. IBM Storage Scale System components also require the following items:
 - An IBM Storage Scale System Management Server
 - An IBM Storage Scale System Protocol Node

IBM Storage Scale System 6000 models

For the following table, all flash media are SED and you cannot inter-mix NVMe and FCM4 in the same chassis.

Models	All Flash SSS6000 NVMe TLC or FCM	ALL-Flash Tiered SSS6000 NVMe TLC and QLC	Hybrid NVMe + HDD SSS6000 Hx	HDD SSS6000 Cx	All-Flash SSS6000 NVMe and JBOF QLC
Models	Quarter, Half or Full 2U48 12, 24 or 48 drives NVMe TLC or FCM	Full 2U48 24 TLC and 24 QLC or 12 TLC and 36 QLC	12, 24 or 48 drives NVMe TLC plus Full 1–9 4U91 Storage enclosures	Full 1–9 4U91 Storage enclosures	12, 24 or 48 drives NVMe TLC plus Full 1–4 2U26 all-flash QLC storage enclosures (26 drives)
TLC Drive Size options	3.84TB, 7.68TB, 15.36TB, 30.72TB	3.84TB, 7.68TB, 15.36TB, 30.72TB	TLC 3.84TB, 7.68TB, 15.36TB, 30.72TB		3.84TB, 7.68TB, 15.36TB, 30.72TB
FCM Drive Size options	FCM4 38.4TB				
QLC Drive Size options		30.72TB, 61.44TB, 122.8TB		HDD 12TB, 16TB, 20TB, 22TB, 24TB	QLC 30.72TB, 61.44TB, 122.8TB
HDD Drive Size options			12TB, 16TB, 20TB, 22TB, 24TB	12TB, 16TB, 20TB, 22TB, 24TB	

IBM Storage Scale System 3500 models

For the following table, all flash media are SED.

	Performance SSS3500	Hybrid SSS3500	Capacity SSS3500
Models	Half or Full 2U24 12 or 24 drives NVMe TLC	12 or 24 drives NVMe TLC plus Full 1-8 4U102 Storage enclosures (102 drives)	12 or 24 drives NVMe TLC plus Full 1-8 4U102 Storage enclosures (102 drives)
TLC Drive Size	3.84TB, 7.68TB, 15.36TB, 30.72TB	3.84TB, 7.68TB, 15.36TB, 30.72TB	
HDD Drive Size		12TB, 16TB, 20TB, 22TB, 24TB	12TB, 16TB, 20TB, 22TB, 24TB

IBM Storage Scale System usable capacity is a percentage of the raw capacity. The usable capacity varies, depending on the IBM Storage Scale RAID parity that is used. The default 8+2P parity is the most common configuration. Approximately 73% of the raw capacity is the usable capacity when you use the default 8+2P parity in external HDD storage enclosures. Small drive configurations provide a lower usable percentage.

IBM or IBM Business Partners can provide the exact usable capacities of the various models on request by using the IBM Storage Modeler (StorM) to calculate the exact usable capacity. Contact your IBM representative or IBM Business Partner representative for assistance in determining the usable capacity in your IBM Storage Scale System model.

IBM Storage Scale System 6000 models




Figure: IBM Storage Scale System 6000 configuration options

IBM Storage Scale System 3500 models




Figure: IBM Storage Scale System 3500 configuration options

Enclosures are fully redundant. For configurations that require enclosure-level protection, use the following erasure code options:

- 8+3P provides enclosure protection with four or more external enclosures.
- 8+2P provides enclosure protection with five or more external enclosures.
- Configurations with one, two, or three enclosures require three-way or four-way erasure code for enclosure protection.

IBM Storage Scale System 7.0.1 release provides new longer erasure codes, 16+2P and 16+3P, that provide better usable capacity.

For more information about the various IBM Storage Scale System models specifications, see [Appendix A: IBM Storage Scale System Models](#).

IBM Storage Scale System Utility Node and Protocol Node

Released as part of IBM Storage Scale System 6.1.8.1, the IBM Storage Scale System Utility Node (MTM 5149-23E) replaced the existing IBM Power9® based platforms, both EMS and Protocol Nodes.

IBM Storage Scale System 3500 can be managed by an IBM Power9 EMS or an IBM Storage Scale System Utility Node. However, the IBM Storage Scale System 6000, released with IBM Storage Scale System version 6.1.9.1, can be managed by only the Utility Node. When you combine IBM Storage Scale System 3500 and 6000 on the same cluster, use the Utility Node instead of the legacy IBM Power9 EMS. Use the Utility Node to help ensure a unified management approach and to avoid multiple management points.

The following figures show the front and back of the Utility Node.



Figure: IBM Storage Scale System Utility Node front view



Figure: IBM Storage Scale System Utility Node rear view

The same MTM 5149-23E is used for both the Management Server and Protocol Node, but it differs on hardware configuration and adapter placement rules.

The Utility Node was released as a Management Server, so it includes a single socket with only two adapter slots available: one slot for mandatory 4-port management card and one slot for a network adapter.

The release of IBM Storage Scale System version 6.2.3 introduces the option of a dual-socket Utility Node as a Management Server. The dual-socket Utility Node also includes an option for more memory and includes four adapter slots. The four adapter slots consist of a mandatory 4-port management card (always in slot 5) and three high-speed network adapter slots (and).



Figure: IBM Storage Scale System single-socket Utility Node EMS adapter layout




Figure: IBM Storage Scale System dual-socket Utility Node EMS adapter layout

Utility Node for Protocol Node is a dual-socket server with different adapter rules than the Utility Node EMS dual-socket. For example, 4-port 10GbE is not needed on the Protocol Node so slot C5 is used for a high speed network adapter (4).




Figure: IBM Storage Scale System Utility Node Protocol Node adapter layout

IBM Storage Scale System 6.2.3 release also introduces the capability to perform an MES of existing Utility Nodes. For single-socket or dual-socket configurations, you can change or update Utility Node memory. For dual-socket versions only, you can change or update Utility Node adapters.

With IBM Storage Scale System version 7.0.0, you can deploy the Utility Node protocol node by using the `essrun` command. Therefore, all the nodes can be deployed by using a single tool. The tool also adds support for Active File Management (AFM) gateway node deployment.

With IBM Storage Scale System version 7.0.1, a 4-port 64 Gb Fibre Channel HBA is introduced as part of the supported and configurable adapters. Slot placement rules apply; see the IBM Storage Scale System Utility Node documentation.

Nvidia DGX SuperPod

IBM Storage Scale System is an approved storage solution for Nvidia DGX SuperPOD.

The Nvidia DGX SuperPOD with Nvidia DGX H200, Nvidia DGX B200 or Nvidia DGX B300 systems is an artificial intelligence (AI) supercomputing infrastructure. It provides the computational power necessary to train today's state-of-the-art deep learning (DL) models and to fuel future innovation. The DGX SuperPOD delivers groundbreaking performance, deploys as a fully integrated system, and is designed to solve the world's most challenging computational problems.

This DGX SuperPOD storage reference architecture (RA) is the result of collaboration between DL scientists, application performance engineers, and system architects to build a system capable of supporting the widest range of DL workloads. The performance that is delivered by the DGX SuperPOD with DGX systems enables the rapid training of DL models at great scale. The integrated approach of provisioning, management, compute, networking, and fast storage enables a diverse, multi-tenant system that can span data analytics, model development, and AI inference.

IBM Storage Scale is a parallel file system, and its optimal performance is achieved when multiple processes or workloads are running in parallel. To maximize the results of a high-performance computing (HPC) and AI cluster solution with IBM Storage Scale 6000 and Nvidia DGX SuperPOD architecture, see IBM Storage and Nvidia DGX SuperPOD with IBM Storage Scale and IBM Storage Scale System publications.

Networking switches

The IBM Storage Scale System solution requires IP or InfiniBand networking switches to interconnect the IBM Storage Scale System hardware components.

As a best practice, purchase and use Nvidia networking switches. These network switches are tested by IBM to work with IBM Storage Scale System.

Clients can also provide and support their own Ethernet or InfiniBand networking switches.

Types of networks that are used by IBM Storage Scale System

The IBM Storage Scale System features the following types of networks that are used to interconnect the IBM Storage Scale System components:

- Management and service networks, which must have separate 1 Gb or 10 GbE networks for management and service traffic
- Data network, which is a high-speed network for the data that is written into and read out of the IBM Storage Scale IBM Storage Scale System file systems

Regardless of the source and type of network switches between the IBM Storage Scale System, the IBM Storage Scale cluster and the users, the networking must be solid, robust, reliable, and able to provide consistent low-latency responses. Nonblocking, high-speed switches are highly recommended.

Experience shows that IBM Storage Scale and IBM Storage Scale System can easily generate throughput rates that can overwhelm an over-subscribed network. IBM Storage Scale provides a network load generation testing tool (`nsdperf`), which is available to anyone as an open source repository on GitHub. Use this tool to test network performance, network latency, network topology, and network readiness level before IBM Storage Scale System is installed.

For more information about the `nsdperf` tool, see the [nsdperf tool on GitHub](#).

Management network switches

IBM Storage Scale System components require a low-speed Ethernet connection for scale management and separate baseboard management controller (BMC) traffic. As part of an IBM Storage Scale System order, IBM can supply the low-speed management network switches for IBM MTM 5149-S54.

As a best practice, order and use this 1/10 GbE network management switch because it is integrated into the IBM Storage Scale System solution at the manufacturing stage. This integration allows the entire IBM Storage Scale System to be installed and initially tested without external network dependencies.

Alternatively, the client can provide a 1/10 GbE network switch management infrastructure. In this case, the client is responsible for providing the suitable management IP networking port counts and network switch configuration and definitions, including VLANs.

High-speed data network switches

IBM Storage Scale System components also require a high-speed data network IP or InfiniBand connection for the reading and writing of user data. As part of IBM Storage Scale System testing, IBM tests with Nvidia data network switches.

As a best practice, order and use the qualified data network switches where possible. These switches are tested and integrated with the IBM Storage Scale System solution.

Nvidia high-speed network switch drivers and firmware are tested and delivered integrated within the IBM Storage Scale System software solution stack.

Alternatively, the client can provide their own high-speed data network switch infrastructure. In this case, the client is responsible for provisioning networking port counts, firmware, network configuration, and definitions. The client is also responsible for maintaining and troubleshooting any networking issues.

Racking

The IBM Storage Scale System solution can be ordered with or without the building block being integrated into an IBM rack.

An IBM Storage Scale System solution that is ordered with feature codes that specify integration into the rack in IBM manufacturing results in a fully integrated, tested, and pre-cabled solution.

The IBM 5149-R42 is the default for most models of IBM Storage Scale Systems. The rack provides 42 EIA (42U) of usable space for IBM Storage Scale System components and network switches. A choice of various electrical PDUs is available to be specified for inclusion in the rack side areas.

It is advantageous to order IBM Storage Scale System as a racked system because IBM Storage Scale System Manufacturing delivers an IBM Storage Scale System with the rack mostly pre-cabled and pre-tested. This configuration provides faster time to install with higher quality.

If installing the IBM Storage Scale System in a customer-supplied rack, it is the customer's responsibility to correctly specify all physical planning for electrical power, connectors, cables, and racking locations.

Clients also must plan the physical installation and cabling of the IBM Storage Scale System components in the client-supplied rack. An IBM TSS Services contract is available and recommended to perform physical planning, specification, and physical installation of the IBM Storage Scale System components in a client rack on behalf of the client.

For the rackless IBM Storage Scale System order, the solution must be assembled and integrated into the rack at the client site.

For the purposes of cable management and serviceability, each IBM Storage Scale System building block is integrated and deployed within the same rack.

If you need to split the IBM Storage Scale System components between adjacent racks, ask your IBM or IBM Business Partner representative to submit an RPO special bid request. IBM examines the request and responds regarding whether IBM can grant and support the specific requested physical racking.

Software components

This section describes the software components that are used on the IBM Storage Scale System solution.

Software solution stack overview

The IBM Storage Scale System solution provides an integrated and tested stack of these operating system, adapter drivers, firmware, IBM Storage Scale software, management software, and installation scripts into a full IBM Storage Scale System software stack. This software solution stack is supported as an integrated IBM Storage Scale storage building block solution by IBM Service and Support.

The following sections includes descriptions of the following IBM Storage Scale System solution stack software components:

- RHEL operating system.
- IBM Storage Scale high-performance parallel file system.
- IBM Storage Scale RAID software-defined storage (SDS).
- IBM Storage Scale System solution installation and management software, which includes (but is not limited to) the following components:
 - IBM Storage Scale System-specific documentation for installation and upgrade scripts
 - Other tools for the IBM System Services Representative (IBM SSR) to use while installing IBM Storage Scale System, such as `essinstallcheck`, `essnetverify`, and `essutils`
- Third-generation IBM Storage Scale System systems deploy a container-oriented management software stack in the IBM Storage Scale System Management Server that includes Red Hat Ansible playbooks for enhanced ease of use during installation and orchestration. Installation is faster, and skill requirements are significantly lower than previous-generation IBM Storage Scale System models.

The IBM Storage Scale System solution stack levels are released as a version, release, and modification level.

The IBM Storage Scale System solution components are periodically upgraded to newer release levels, tested as an integrated solution, and released as a new level of the IBM Storage Scale System solution software. Additionally, generalized recommendations are made available to help clients implement a code update strategy. It is a full-field perspective, and as such, a customized recommendation that accounts for specifics such as business upgrade windows, length of time since last update, and decommission plans. You might require assistance from local support teams. For more information, see [IBM Storage Scale Software Version Recommendation Preventive Service Planning](#).

Operating system

The IBM Storage Scale System solution runs RHEL as the operating system on the IBM Storage Scale Data Servers.

Each IBM Storage Scale System solution release level integrates and tests a suitable current level of RHEL, including any necessary RHEL fixes and errata that are required for the successful operation of the IBM Storage Scale System solution stack. IBM periodically provides new IBM Storage Scale System solution release levels that incorporate newer levels of RHEL. These releases are provided often enough to assure that a current level of RHEL is always available.

IBM uses an embedded RHEL license. This license is included with every order.

IBM Storage Scale

This section provides a description of IBM Storage Scale, the high-performance parallel clustered file system software that is used in IBM Storage Scale System. IBM Storage Scale software is released several months before the

equivalent IBM Storage Scale System release. Intentionally, the release numbers of the Storage Scale software differ in the first digit from the Storage Scale System software. At the time of publication, the current software version is 6.0.0. The current Storage Scale System version is 7.0.0.

IBM Storage Scale is highly distributed, clustered file system software that provides high-speed concurrent data access to applications that run on multiple nodes and clusters. In addition to providing parallel high-performance file storage capabilities at exabyte scale, IBM Storage Scale provides tools for tiering, management, administration, and archiving of that enterprise-level data. IBM Storage Scale is the IBM strategic SDS for enterprise big data, analytics, and AI applications.

In the most common IBM Storage Scale deployment architecture, IBM Storage Scale data is accessed by IBM Storage Scale clients and users over a LAN network, accessing disk volumes that are known as Network Shared Disk (NSDs). The NSDs are attached to IBM Storage Scale nodes that are known as NSD Data Servers. In this IBM Redpaper publication, these nodes are also referred to as Data Servers.

IBM Storage Scale RAID

IBM Storage Scale RAID is an SDS controller that performs all of the storage controller functions that are normally associated with hardware storage controllers. IBM Storage Scale RAID integrates all HA and features of an advanced storage server into IBM Storage Scale SDS.

IBM Storage Scale RAID runs on the IBM Storage Scale System NSD Data Servers. IBM Storage Scale RAID provides sophisticated data placement and error correction algorithms to deliver high levels of storage reliability, availability, and serviceability, and performance.

IBM Storage Scale RAID implements a declustered erasure code parity schema, distributing data, redundancy information, and spare space across all disks of the IBM Storage Scale System enclosures. With this approach, a significant improvement is realized on the application performance. Also, storage rebuild time overhead is reduced (disk failure recovery process) compared to conventional RAID controllers.

IBM Storage Scale RAID implements large cache for performance by using memory on the IBM Storage Scale System NSD Data Servers. The large cache intelligently improves read and write performance, particularly for small block I/O operations.

IBM Storage Scale RAID mitigates performance effects of storage rebuilds

If storage failures occur, IBM Storage Scale RAID reconstructs lost or erased stripes for I/O operations dynamically. By using the highly distributed erasure coding, IBM Storage Scale RAID mitigates the performance effect of storage media failures.

IBM Storage Scale RAID end-to-end checksums

IBM Storage Scale RAID includes integrated end-to-end checksums that detect data corruption that might otherwise go undetected by a conventional storage controller. Unlike conventional storage controllers, IBM Storage Scale RAID is integrated with the IBM Storage Scale file system and performs end-to-end checksum comparison all the way out to the IBM Storage Scale client code on the workstations. This feature ensures data integrity at a file system level, detecting and automatically correcting data corruption errors that might occur in conventional storage environments.

In an environment where a customer experienced excessive file system checks and suffered downtime to repair file systems, the use of IBM Storage Scale RAID end-to-end checksums mitigates file system check problems. This feature assures availability of data and removes application outages that are caused by file system checks.

IBM Storage Scale RAID disk hospital

One of the key features of IBM Storage Scale RAID is the disk hospital. This powerful function asynchronously diagnoses errors and faults in the IBM Storage Scale System storage media, down to the level of the individual drive and the individual performance of each drive.

IBM Storage Scale RAID is fully aware of and tracks the performance of each individual drive because not all drives perform equally. IBM Storage Scale RAID uses the individual performance history of each drive to make intelligent data allocation and data retrieval decisions.

Extensive health metrics down to the level of the individual drive are maintained by the disk hospital. Performance variation is continually monitored. If or when a disk metric exceeds a threshold, the storage media is marked for replacement according to the disk maintenance replacement policy for the declustered array.

As an example, disk hospital features the following metrics:

- `relativePerformance`, which characterizes response times. Values are compared to the average speed. If the metric falls below a specific threshold, the hospital adds `sLow` to the pdisk state, and the disk is prepared for replacement.
- `dataBadness`, which characterizes media errors (hard errors) and checksum errors.

For more information about IBM Storage Scale RAID implementation and best practices, see the [Administering IBM Storage Scale RAID documentation](#).

IBM Storage Scale System solution installation and management scripts

This section includes an overview of the components of the IBM Storage Scale System solution installation and management scripts. This overview includes the following information:

- IBM Storage Scale System-specific documentation for installation and upgrade scripts
- Ansible playbooks
- IBM Storage Scale System specific tools and utilities for the IBM SSR and administrators to use while installing or maintaining the IBM Storage Scale System

For more information about these IBM Storage Scale System solution components, see the IBM Storage Scale System solution release-specific level information at [IBM Storage Scale System Software documentation](#).

Installation, upgrade, and administering guides

IBM provides manuals and documentation for deploying and administering IBM Storage Scale System, including the following publications:

- Quick Deployment Guide (QDG), which documents IBM Storage Scale System-specific scripts for installing, deploying, and upgrading IBM Storage Scale System for experienced users
- IBM Storage Scale RAID Administration, which focuses on administering IBM Storage Scale RAID on IBM Storage Scale System
- Problem determination guide, which provides more information about monitoring, troubleshooting, and maintenance procedures

For more information, see [IBM Storage Scale System Software documentation](#).

Ansible playbooks and container-based software installation

Third-generation IBM Storage Scale System models include a software update process and methodology that are based on Ansible playbooks.

On this IBM Storage Scale System Management Server node, Ansible is configured to store the container and cluster configuration and definitions for all the managed IBM Storage Scale System Data Servers. Network services (such as Dynamic Host Configuration Protocol (DHCP) and SSH) are enabled to allow IBM Storage Scale System Management Server to deploy IBM Storage Scale System solution stack components to all IBM Storage Scale System nodes in the cluster.

`essutils` are IBM Storage Scale System installation and deployment toolkits that facilitate IBM SSR hardware setup, installation, deployment, and upgrade tasks. Any authorized IBM Storage Scale System administrator can also use these tools.

For more information about `essutils` and `gssutils`, see the Quick Deployment Guide and the IBM Storage Scale System Hardware Planning and Installation Guide and IBM Storage Scale System software deployment preparation,

which are available at [IBM Storage Scale System Software documentation](#). For PDF versions, see [IBM Storage Scale System 6000 PDFs](#).

Footnotes

1. IBM newsroom: IBM Storage Scale System 6000 delivers industry-leading throughput for data-intensive workloads. IBM, 2024. [↩](#)

IBM Storage Scale System planning and integration

This section provides guidelines and considerations for the correct planning, installation, and configuration of the IBM Storage Scale System.

This section also describes configurations and integration considerations for a smooth IBM Storage Scale System deployment into an existing or a new IT environment. This section includes the following elements:

- Roles and responsibilities
- Planning
- Networking
- Installation
- Maintaining hardware
- Maintaining software
- Upgrades

IBM Storage Scale System planning overview

This section provides recommendations and highlights best practices for various roles and responsibilities when an IBM Storage Scale System is deployed.

Best practices were learned from successful IBM Storage Scale System installations regarding the roles and responsibilities for each of the areas of expertise.

Overall planning

Appoint an experienced project leader to coordinate and manage the IBM Storage Scale System installation project. Deploying an IBM Storage Scale System requires coordination across different IT departments and involves various team members (IBM and client).

The project manager must coordinate with various client and IBM teams, such as the following groups:

- Physical planning for the data center, including the networking team
- Networking setup, switch installation, cabling, and verification
- Hardware setup, installation, and verification
- Software planning and integration of the file system into the environment
- Acceptance testing and criteria for moving into production

Networking

IBM Storage Scale System components are connected by using IP or InfiniBand networking.

IBM Storage Scale System data users access IBM Storage Scale System by way of high-speed IP or InfiniBand networking.

The networking planning includes (but is not limited to) the following tasks:

- Order or allocate suitable networking switches and include cable planning, ordering, racking, and proper connectors.
- Plan network addressing for the IBM Storage Scale System and the integration of that network into existing client IP or InfiniBand network.
- After the network is designed and approved, implement and test IP addresses.
- Design, configure, and test multiple VLANs on the management network switches.
- Configure and test suitable networking authentication and firewall rules to ensure that they comply with client network audit and security requirements.
- Test the network before IBM Storage Scale System is installed to identify and verify the network's ability to provide IBM Storage Scale System with the required network bandwidth, latency, and consistent performance. Use tools, such as `nsdperf` or Network Readiness Tools that are provided by IBM Storage Scale as no-charge, open source tools.

These tasks require the client IBM Storage Scale System infrastructure implementation team to complete the following tasks:

- Identify all IP or InfiniBand networking-related requirements.
- Engage early with the infrastructure networking team.
- Document and communicate all required IP network provisioning, implementation, and network implementations, including all network stress testing and acceptance criteria.
- Monitor these tasks closely to ensure that the networking is in place and load-tested before IBM Storage Scale System is installed. Experience shows that networking addressing mismatches, VLAN issues, configuration issues, and a failure to remedy these issues are the major cause of unforeseen delays in IBM Storage Scale System installation progress.

The client IP or InfiniBand networking team is responsible for the following tasks:

- Receive the specified networking requirements.
- Consult and collaborate with the IBM Storage Scale System implementation team to design, optimize, and agree upon a suitable IP or InfiniBand networking design and implementation plan.
- Implement the agreed-upon networking changes.
- Run requested network load tests by using tools that are provided by IBM Storage Scale before IBM Storage Scale System is installed, including agreeing to and implementing any necessary quality assurance and acceptance criteria.
- Adjust and tune the network based on pre-installation tests or during the IBM Storage Scale System implementation.

The following best practices from successful IBM Storage Scale System installations help ensure success:

- Engage early with the client networking team.
- Obtain collaborative buy-in and well-defined ownership of networking-related task completions.
- Run network performance tests by using IBM Storage Scale tools far enough in advance of the IBM Storage Scale System installation allows time for necessary tuning adjustments to be made.
- Ongoing IBM Storage Scale System implementation team and client networking team dialog and interaction during the IBM Storage Scale System installation process quickly remedies any networking-related installation issues.

Installation

In this phase, the following IBM and client teams must be involved and coordinated:

- Physical site administration and engineering
- Networking team and network engineering

- Hardware implementation team, including client IT administrators, IBM System Services Representative (IBM SSR) for IBM Storage Scale System physical installation and setup of physical hardware
- Software implementation team, potentially including IBM System Lab Services consultants, and client operating system and software administrators

These teams collaborate to implement the following phases:

- Physical site planning

This planning includes the receipt and staging of IBM Storage Scale System components after shipment, raised floor space planning and racking; electrical power planning and provisioning, including specifying suitable electrical power, cables, connectors, and power distribution units (PDUs); electromagnetic compatibility; and ensuring proper airflow and air conditioning.

- Software planning and integration

IBM Storage Scale software planning involves selecting the appropriate edition (Data Access or Data Management), determining cluster architecture and node roles, planning Storage Scale file system implementation parameters, integration testing, and acceptance with the client applications and procedures.

- Hardware physical setup, installation, and verification

This planning and scheduling includes the IBM Storage Scale System hardware setup with the IBM SSR and the IT physical installation teams.

- Acceptance testing

Running well-defined criteria for acceptance testing and moving the IBM Storage Scale System into production are featured in this phase.

Maintaining hardware

After the IBM Storage Scale System is installed and running in production, many parties are responsible for various IBM Storage Scale System hardware tasks, including the following examples:

- Client IBM Storage Scale System administrators:
 - Monitoring IBM Storage Scale System physical hardware by using provided tools and the GUI
 - Basic hardware problem determination
 - Optional basic maintenance, such as replacing a failed solid-state drive (SSD) or hard disk drive (HDD)
- IBM SSR: Responsible for physical IBM Storage Scale System hardware maintenance and repair

Maintaining software

After the IBM Storage Scale System is installed and running in production, many parties are responsible for various IBM Storage Scale System software tasks. IBM Storage Scale administrators have the following responsibilities:

- Monitoring and tuning IBM Storage Scale performance
- Data management, which includes monitoring data and storage usage patterns, and writing and running policies that manage data, backups, and snapshots

Learning services and classes

IBM Learning Services or IBM Business Partners can provide education courses about the customer topics for maintaining IBM Storage Scale and IBM Storage Scale System hardware. More worldwide resources might be available.

The following classes and others are available for IBM Storage Scale and IBM Storage Scale System:

- [IBM Storage Scale Basic Administration for Linux and AIX](#)

- [IBM Storage Scale Advanced Administration for Linux](#)

For more information, contact your IBM representative or IBM Business Partner.

Upgrades

The IBM Storage Scale System provides an integrated software solution stack, which is composed of specific software release levels and the following fix levels:

- RHEL kernel levels and components
- Firmware (System/Chassis/Drive)
- Network interface card (NIC) drivers
- IBM Storage Scale and IBM Storage Scale RAID software

IBM recommends upgrading the IBM Storage Scale System software solution stack at least once a year and to keep the IBM Storage Scale System current on levels of RHEL, firmware, and IBM Storage Scale.

IBM System Lab Services is available and it is recommended to engage Lab Services to assist and perform these software upgrade actions on behalf of the client, if required.

Technical and Delivery Assessment

An IBM Technical and Delivery Assessment (TDA) is a technical inspection of a completed solution design. A TDA uses technical subject matter experts (SMEs) who were not involved in the solution participate to ensure the client's satisfaction.

In a pre-sales TDA, which is performed before the IBM Storage Scale System solution is ordered, the IBM IMPACT tool and process is used to examine and approve the IBM Storage Scale System solution design to determine the following information:

- Will the IBM Storage Scale System solution work as expected?
- Is the configuration and implementation plan sound?
- Will the IBM Storage Scale System solution meet client requirements and expectations?

In a pre-installation TDA, SMEs also evaluate the client's environment for readiness to install, implement, and support the proposed solution.

IBM and IBM Business Partner teams are responsible to conduct these pre-sale and pre-installation processes. Contact your IBM representative or IBM Business Partner to schedule and perform these TDA and StorM processes.

The IMPACT and pre-installation TDA tools can help your IBM account team, client technical specialist, and IBM Business Partner select the best IBM Storage Scale System and IBM Storage Scale solutions match your business requirements and needs. For more information, see the following web pages:

- [Storage Modeler \(StorM\)](#)
- [IMPACT/TDA Tool](#)

These resources are available to IBMers or Business Partners after logging in with their IBMid or IBM PartnerWorld ID.

IBM Storage Scale System hardware planning

This section includes a description of the necessary planning and implementation steps for installing the IBM Storage Scale System hardware.

Planning for network switches and networking

Suitable network planning and configuration are one of the most important success factors for an IBM Storage Scale System installation. This planning is crucial because the IBM Storage Scale and IBM Storage Scale System cluster

components are connected by way of an extensive set of IP addresses and associated networking switches and cabling.

The following networks are required to install an IBM Storage Scale System:

- Service network

This network connects the Baseboard Management Controller (BMC)/Flexible Service Processor (FSP) on the management server and I/O server nodes, depending on the IBM Storage Scale System generation.

- Management and provisioning network

This network connects the management server to the I/O data server nodes. The management server runs the Dynamic Host Configuration Protocol (DHCP) on the management and provisioning network. If an IBM Storage Scale System Management Server is not included in the solution order, it is assumed that an IBM Storage Scale System Management Server exists in the same IBM Storage Scale cluster that can be used.

- Clustering network

This high-speed data network is used for IBM Storage Scale System data transfer to and from the rest of the IBM Storage Scale cluster and client node access.

- External and campus management network

This public network is used for external and campus management of the IBM Storage Scale System Management Server, and optionally to other nodes in the scale cluster.

The following figure shows an example network topology for IBM Storage Scale System.




Figure: IBM Storage Scale System network topology example

Nvidia CX-7 cards from the IBM Storage Scale System 6000 are required for connecting to the Nvidia BlueField-3 DPUs in the IBM Storage Scale 5149-F26 enclosure. For more information, see the [Storage Scale System 5149-F26 Storage Enclosure documentation](#).




Figure: IBM Storage Scale System 6000 network topology with JBOF 5149-F26 enclosures

An overview of the many networking steps that are required in the installation of an IBM Storage Scale System are described next. This example is representative (not all inclusive) of the necessary network planning and implementation steps that must be completed before installing the IBM Storage Scale System hardware. For more information about network planning, see the [IBM Storage Scale System Hardware Planning and Installation Guide](#).

Example: Network IP addressing, planning, and implementation for IBM Storage Scale System

This section includes an example of typical networking-related items for installing an IBM Storage Scale System. The following template can be used for the installation process:

- Confirm Private IP range for DHCP server.
- Confirm Private Service network with defined IP addresses and private management network with defined IP addresses. Separate the networks by using switches or VLAN.
- Confirm Public network connections for IBM Storage Scale System Management Server. Multiple IP addresses are needed.
- Select **Host IP mappings** to confirm the mappings for the following items:
 - IBM Storage Scale System Management Server
 - IBM Storage Scale System Data Servers: I/O server 1, I/O server 2, I/O server 3, I/O server 4, and so on
 - 200 Gb Ethernet (GbE), 100 GbE, 40 GbE, 25 GbE, 10 GbE, or InfiniBand switches
- Set up domain names for management private network.
- Set up domain names for high-speed interconnect.
- Set up Data I/O partition and partition profile names.
- Confirm Server names.
- Confirm that the Ethernet or InfiniBand switches are in place and cabled.
- Set up Bonding if it is used.
- Set up Public network in place and cabled to IBM Storage Scale System Management Server.

- Confirm that all building block components are correctly cabled in the rack (Data I/O servers, IBM Storage Scale System Management Server, and switches).
- Set up and confirm dual feed power to racked IBM Storage Scale System components.
- Set up the IBM Storage Scale System Management Server console or terminal, including the network connections.
- Prepare for and install RHEL errata.
- Define and specify IBM Storage Scale installation parameters and requirements. For example, how many file systems, what are the block sizes, how is metadata allocated and split, and what replication factor is used.
- Confirm all physical NVMe, SSD, or HDD disks are in place. IBM Storage Scale System installation checks with scripts.
- Confirm that all physical cabling in place. IBM Storage Scale System installations are checked by scripts.
- Confirm that wifi access exists in lab to set up Webex (for access to and by support teams).

This template shows the processes that IBM Systems Lab Services can perform as part of a contract to install an IBM Storage Scale System. As part of that contract, IBM Systems Lab Services provides and works with the customer on a detailed IP networking planning document.

Planning for servers that are used in IBM Storage Scale System

This section a review of planning information for the physical servers that are used in an IBM Storage Scale System. The section also includes specific recommendations for memory, NICs, and cabling and connectors.

IBM Storage Scale System Management Server

Each IBM Storage Scale cluster that uses an IBM Storage Scale System requires at least one IBM Storage Scale System Management Server. The IBM Storage Scale System Management Server acts as the central management focal point for all IBM Storage Scale System systems that are in the cluster. It runs specific functions, such as the IBM Storage Scale System GUI and stores and distributes the master boot image for all the IBM Storage Scale System systems that are in the cluster by using Red Hat Ansible playbooks.

IBM Storage Scale System Data Servers

Every IBM Storage Scale System is an integrated IBM Storage Scale storage building block that consists of two IBM Storage Scale Data Servers, and each data server is cross-connected to all the storage for high availability (HA). In normal operation, half of the IBM Storage Scale System storage is running through the first data server, and the other half of the IBM Storage Scale System storage is running through the second data server. If a planned or unplanned outage occurs on one of the data servers, the storage is failed over to the remaining data server in the IBM Storage Scale System, which provides HA and resiliency.

Note: Both data servers must have the same amount of memory in each IBM Storage Scale System building block. When possible, as a best practice, ensure that all IBM Storage Scale System systems in a specific IBM Storage Scale storage pool have the same size and speed. If this configuration is not possible, keep the IBM Storage Scale System similar in size and speed. IBM Storage Scale parallelism means that in a specific storage pool, the speed of data reads/writes is gated by the slowest storage building block.

Connect up to four IBM Storage Scale System 5149-F26 Storage Enclosures to an existing IBM Storage Scale System 6000 performance model using Nvidia CX-7 adapters.

For more information, see the [Storage Scale System 5149-F26 Storage Enclosure documentation](#).

Network interface cards

The IBM Storage Scale System features the following types of networks that are used to interconnect the IBM Storage Scale System components:

- Management and service

Mandatory, separate 1/10 GbE network for management and service traffic.

- Data

A high-speed network for the user data being written into and read out of the IBM Storage Scale and IBM Storage Scale System file system.

Standard Nvidia server NICs are used to provide high-speed network connectivity in an IBM Storage Scale System. The network cards that are available vary depending on the IBM Storage Scale System generation.

It is important to ensure that the network bandwidth requirements are clearly studied, documented, and communicated. Specify a suitable number of NICs to meet expected IBM Storage Scale System performance requirements. Plan and specify the correct cables and connector types.

For more information, see the [IBM Storage Scale System Hardware Planning and Installation Guide](#).

Cabling and connectors

In addition to specifying the correct number of NICs to provide bandwidth requirements, make sure that physical planning for cable lengths, adapter connector types, and switch connector types, are clearly studied, documented, and communicated.

Physical machine room placement of IBM Storage Scale System components relative to the physical location of the IP network switches must be clearly understood and documented.

Physical cable orders, including wanted colors, connector types, and required cable routing must be studied, specified, and verified with network teams, server teams, physical site teams, and infrastructure teams.

Many network connector types and cable types are available. Requirements must be identified for the required networking components, including the following examples:

- Transceivers
- Optical versus copper
- Connector form factor
- Cable lengths

Planning for storage configuration

IBM Storage Scale System storage configuration for ordering purposes is relatively straightforward.

First, determine your IBM Storage Scale storage requirements. Request that your IBM representative or IBM Business Partner use the StorM tool to help identify usable capacity requirements and to estimate wanted performance.

After storage capacity and performance requirements are known, identify which models of IBM Storage Scale System feature the suitable type and quantity of storage media available. Within each model, you choose the suitable drive size.

Each IBM Storage Scale System model includes predefined storage configurations and requirements, which are pre-tested and supported as a solution by IBM. Cable lengths, installation parameters, adapters, drive preparation, and location, all are pre-specified within the IBM Storage Scale System architecture.

Planning for IBM Storage Scale protocol nodes

First announced in 2014, IBM Storage Scale protocol nodes provide external Server Message Block (SMB), Network File System (NFS), object, and Hadoop access to IBM Storage Scale data.

IBM Storage Scale requires that all protocol nodes in an IBM Storage Scale clusters are alike in processor architecture, performance, and memory. For more information about protocol nodes, see the [IBM Storage Scale System Utility Node and Protocol Node](#)

Planning for physical, electrical, and cooling installation

Preparation for installation of your IBM Storage Scale System is as important as the physical installation of the system. To have a smooth and successful installation, preparation and planning processes are needed.

TDA and pre-installation meetings with IBM TSS and IBM Service personnel must be completed. Also, IBM Storage Scale Server (IBM Storage Scale System) environmental specifications provide detailed information about your system, including dimensions, electrical, power, temperature, environmental requirements, and noise emissions. By following the preparation process, IBM service personnel can efficiently install the physical hardware and deploy the software.

For the specifications for IBM Storage Scale System, see [IBM Storage Scale System Models](#).

IBM Storage Scale System software planning

This section highlights key items for planning the installation of the IBM Storage Scale System software stack and in particular, IBM Storage Scale and IBM Storage Scale RAID.

Planning for IBM Storage Scale file system

The IBM Storage Scale file systems that are placed into the IBM Storage Scale System must be thoroughly planned, documented, and communicated. Make sure that user application, authentication, firewall, and backup and restore requirements are understood and included.

Before the arrival and installation, the IBM Documentation manual for IBM Storage Scale and IBM Storage Scale RAID Administration should be consulted. These manuals contain detailed instructions and best practices recommendations for implementing the requirements of the client's environment.

Planning for IBM Storage Scale RAID

The IBM Storage Scale RAID Administration manual provides best practices and recommendations for setting up IBM Storage Scale RAID. Consult this manual and follow the recommendations.

The IBM Storage Scale System and IBM Storage Scale Server implementation of IBM Storage Scale RAID uses just a bunch of disks (JBOD) arrays, provides the required redundancy protection, and delivers usable disk capacity, required spare capacity, and maintenance strategy.

In IBM Storage Scale System, the IBM Storage Scale RAID implementation best practices are enforced as de facto standards by the IBM Storage Scale System installation scripts, which provide configuration parameters that can achieve the best storage performance.

Each JBOD array is connected to each of the two Network Shared Disk (NSD) Data Servers in the IBM Storage Scale System to protect against server failure. Each NSD Data Server has two independent paths to each physical disk to protect against path failure and provide higher throughput to the individual disks.

Ordering IBM Storage Scale System

IBM Storage Scale System is always part of a larger IBM Storage Scale system. Because IBM Storage Scale and IBM Storage Scale System are sophisticated technologies to ensure high client satisfaction, IBM requires the use of TDA and the StorM tools to elicit requirements and ensure that designs are reviewed, assessed, and approved.

The IBM Storage Scale System solution process recommends a pre-sale and a pre-delivery TDA. The purpose of these two assessments is to match the client's requirements and needs. It also evaluates the client's current IT environment to create the best proposed IBM Storage Scale System solution that matches the client's needs.

The following figure provides an example of capacity modeling for an IBM Storage Scale System 6000 that is configured with four JBOFs in StoreM.

Figure: StoreM storage configuration capacity view

As part of the order, various decisions regarding hardware must be made. Be prepared to supply requirement information to the IBM Business Partner. Provide the following requirements:

- IBM or the Business Partner then configures the IBM Storage Scale System. IBM Storage Scale System components can be installed in an IBM rack (preferred) or installed in customer-supplied racks.

51

IBM Storage Scale System Physical Racking Configuration Examples

Storage Scale 6000 + 3500 H1
Configuration Example

42u	Top of Rack Switch
...	
18	
17	Protocol Node
16	
15	Protocol Node
14	
13	Scale Management Server
12	
11	SSS 6000
10	
9	
8	
7	SSS3500
6	
5	5147-102 Storage Enclosure
4	
3	
2	
1	Fans/Power

Note: Racking examples follow the suggested method of having the heaviest equipment on the bottom and preserving as much space as possible for the future expansion.

Storage Scale 3500H8
Configuration Example

42u	Top of Rack Switch
41	Protocol Node
40	
39	Protocol Node
38	
37	Scale Mgmt. Server
36	
35	SSS3500
34	
33	5147-102 Storage Enclosure
32	
31	
30	5147-102 Storage Enclosure
29	
28	
27	5147-102 Storage Enclosure
26	
25	
24	5147-102 Storage Enclosure
23	
22	
21	5147-102 Storage Enclosure
20	
19	
18	5147-102 Storage Enclosure
17	
16	
15	5147-102 Storage Enclosure
14	
13	
12	5147-102 Storage Enclosure
11	
10	
9	5147-102 Storage Enclosure
8	
7	
6	5147-102 Storage Enclosure
5	
4	
3	5147-102 Storage Enclosure
2	
1	Fans/Power

Storage Scale 6000 with
MTM 5149-F26

42u	Top of Rack Switch
...	
38	PDU
37	PDU
36	Model F26
35	
34	Model F26
33	
32	Model F26
31	
30	Model F26
29	
28	IBM Scale System 6000
27	
26	
25	Model F26
24	
23	
22	Model F26
21	
20	
19	Model F26
18	
17	
16	IBM Scale System 6000
15	
14	
13	Model F26
12	
11	
10	Model F26
9	
8	
7	Model F26
6	
5	
4	IBM Scale System 6000
3	
2	
1	

Figure: IBM Storage Scale System server physical racking configuration examples

The configuration on the left in the following figure is an example of an IBM Storage Scale System single rack configuration. The configuration on the right illustrates an example configuration of four MTM 5149-F26 connected to each Storage Scale System 6000. These configurations are examples only. Configurations are tailored to your requirements. However, when you plan the rack layout, plan to install the heaviest components on the bottom and the lightest components on the top.

IBM Storage Scale System installation

The installation process starts with the client preparing their environment to be ready for the IBM Storage Scale System. After the client's environment is prepared, the IBM expert service personnel install the hardware. Depending on the solution, the deployment of IBM Storage Scale System also is performed by IBM specialized service personnel.

Physical installation, racking, and cabling

IBM Technical Support Services (TSS) should be engaged to help you plan for the physical installation and racking of the IBM Storage Scale System components.

IBM TSS works with you to help you plan for the many items to consider for racking, installation, and servicing of the IBM Storage Scale System.

Your TSS representative can help you plan for many factors, including the following examples:

- Electrical power, power connectors, PDU, and amperage
- Network adapter and connectors
- Electrical power cables, network cables, and cable lengths
- Shipping location of the system
- Lift tools
- Ladders
- IBM SSR physical installation instructions (WCII) and IBM SSR scheduling and confirmation

If you have an IBM Storage Scale System HDD model, allow time for the physical installation of the many HDDs into the storage enclosures. HDD IBM Storage Scale System models are shipped with the storage enclosures empty to mitigate shipping damage. Each HDD is shipped in a sealed protected package. In a large IBM Storage Scale System, many HDDs are used and it takes time for the IBM SSR to install each HDD into its slot in the storage enclosures.

When planning the physical installation, remember to consider future expansion needs. Some factors include the following examples:

- When purchasing new Storage Scale System hardware from IBM, is a new rack also needed?
- Is there sufficient space in the current or desired rack location which can be reserved to allow for expansion in the future?
- Is the location of the rack in the data center suitable for future cabling and electrical power supply expansion?
- Is the location of network switches and the IBM Storage Scale System Management Server in a location that has suitable cable lengths and access for current or future networking needs?

Integrating the IBM Storage Scale System into the clients' environments consists of many actions. The following sections include an overview of key best practices, practical advice, and experiences and lessons learned for a successful IBM Storage Scale System deployment.

Integrating into the customer networking environment

In any IBM Storage Scale or IBM Storage Scale System solution, the network interconnect always plays a central role. The goal of this section is to guide you in lessons that are learned from an IBM Storage Scale or IBM Storage Scale System network readiness discussion.

Networking is a vast expertise area, and you should not expect to need to become a networking expert. The following questions guide you through initial discussions you must have with your networking infrastructure team to ensure that your IP network is ready for your IBM Storage Scale or IBM Storage Scale System solution to integrate IBM Storage Scale System into your environment.

Determining which network to use

Document if you are installing IBM Storage Scale or IBM Storage Scale System on an existing network or if this installation is on a new network that is designed specifically for IBM Storage Scale or IBM Storage Scale System.

IBM Storage Scale or IBM Storage Scale System is a high-performance, high-data rate big data solution. Best practices include providing the IBM Storage Scale or IBM Storage Scale System cluster with a dedicated network infrastructure that is co-designed with the solution.

However, especially in a commercial environment, many (if not most) circumstances are such that you cannot implement a dedicated network for IBM Storage Scale or IBM Storage Scale System. Instead, the IBM Storage Scale or IBM Storage Scale System traffic must coexist on an IP network. Especially in this case, consider the following points to validate if the network can provide the high bandwidth and consistent low latency IP networking that IBM Storage Scale or IBM Storage Scale System solution requires:

- IP Network vendors and switch models.

Document and provide a list of the IP network switch vendors and switch models that are to be used to provide the network infrastructure for this IBM Storage Scale or IBM Storage Scale System cluster. IBM Storage Scale or IBM Storage Scale System depends upon IP networking for all of the IBM Storage Scale nodes and users to connect and communicate with each other.

Nvidia is the preferred high-performance IP networking switch vendor for IBM.

- Provide a network topology diagram.

Provide a copy of the network topology and document where the IBM Storage Scale System will be placed and how they will be connected.

- Total number of nodes in the IBM Storage Scale or IBM Storage Scale System cluster.

Document the number and type of IBM Storage Scale and IBM Storage Scale System nodes that are planned for your solution. Provide the number of protocol nodes and the number of IBM Storage Scale NSD Data Servers.

Verifying sufficient port counts and port speeds that are available for the network switches

Document the number of ports and port speeds that are required by your IBM Storage Scale System and ensure that the number of ports and speeds are available within the networking infrastructure.

Determining the interconnects between switches

Document the inter-switch links (ISLs), where they are located, what traffic traverses them and if existing networking links and switches be used.

IBM Storage Scale or IBM Storage Scale System can generate unusually high workload on a network. If IBM Storage Scale or IBM Storage Scale System is to be placed on an existing network, document how heavily used the network is, and what the anticipated usage will be after IBM Storage Scale System is added.

Best practice is to validate whether the client's network can sustain the intended IBM Storage Scale or IBM Storage Scale System workload, use the IBM network performance tool `nsdperf` (written by IBM Storage Scale Development). The `nsdperf` tool mimics IBM Storage Scale workloads and provides an open-source, no-charge method to test the network before installation.

For more information, see the [nsdperf tool on GitHub](#).

Determining what other network traffic types and amounts must coexist on these network links

Document what other traffic is to compete for bandwidth on the IP network. IBM Storage Scale depends on low latency, consistent delivery of many small metadata messages to work properly. If any congestion exists in the IP network, IBM Storage Scale can encounter problems. Therefore, the best practice is to have a dedicated network infrastructure for the IBM Storage Scale or IBM Storage Scale System, if possible.

Communicate this need for low latency and consistent, congestion-free delivery of IBM Storage Scale metadata messages to the network infrastructure team.

Determining the basic workload types and sustained bandwidth requirements

Document what other traffic IBM Storage Scale is to compete with for network bandwidth. Identify usage cases that are to be used, looking to check the workload mix (small block IOPs versus large block sequential, for example). Engage a qualified IBM Storage Scale architect with networking skills to project how the IBM Storage Scale or IBM Storage Scale System network infrastructure performs in your specific environment.

Determine performance expectations. In addition to documenting GBps sequential streaming workload requirements, what other specific workloads are to be running? What tools or criteria are to be used to quantify that performance? What specific criteria determine “good” or “not so good” performance?

Determining if validation is required to confirm that the existing network can provide necessary GBps and IOPS

If it is deemed necessary to examine the network further and validate that it is ready for an IBM Storage Scale or IBM Storage Scale System solution after reviewing these questions and answers, use the `nsdperf` tool.

For more information about this tool, see the [nsdperf tool on GitHub](#).

Integrating into a customer software environment

IBM Storage Scale and IBM Storage Scale System provide various file or object interfaces to the user applications. Determine, inventory, and document the user applications and their required interfaces that the applications require. From that information, design the file systems, directory structures, and data architecture that must be supported by the IBM Storage Scale or IBM Storage Scale System solution.

Design the IBM Storage Scale or IBM Storage Scale System cluster to deploy the wanted POSIX, Linux, or UNIX file systems. If necessary, design and deploy the IBM Storage Scale protocol nodes and their associated shares, mount points, and IP addresses that the software applications require.

Integrating into a customer security and authentication environment

IBM Storage Scale System is a part of an IBM Storage Scale cluster. IBM Storage Scale System and IBM Storage Scale support security and authentication by using various interfaces, including LDAP and Active Directory. Determine, document, and inventory the required security, authentication, and access control requirements. Then, use that information to design the IBM Storage Scale or IBM Storage Scale System solution.

For more information about IBM Storage Scale authentication and security, see *IBM Storage Scale Security*.

Reliability, availability, and serviceability

IBM Storage Scale System uses servers and design that have a rich tradition of reliability, availability, and serviceability (RAS).

IBM Storage Scale System’s reliability (and the availability of the applications it supports) is a function of much more than just the reliability of the processors, or even of the entire system hardware. A full description of a system design for RAS must include all of the hardware, firmware, operating system, applications, operating environment, duty cycle, and so on.

Reliability in hardware is all about how often a hardware fault requires a system to be serviced, the less frequent the failures, the greater the reliability.

Availability is how infrequently such a failure affects the operation of the system or application. For high levels of availability, correct system operation must not be adversely affected by hardware faults. That is, a highly available (HA) system design ensures that most hardware failures do not result in an application outage.

Serviceability is concerned with identifying what fails and ensuring an efficient repair (of that component, firmware, or software).

IBM Call Home: IBM Storage Scale System hardware and software

The IBM Storage Scale System Management Server and I/O server hardware includes IBM Call Home capability. When a serviceable event occurs on one of the monitored components (such as servers or storage), the Call Home generates a Call Home event.

IBM Storage Scale software also has its own separate Call Home capability for purposes of automatically collecting and securely uploading data that is related to IBM Storage Scale to IBM Service and Support.

Monitoring IBM Storage Scale System

Monitoring the IBM Storage Scale System system includes system health, performance, and capacity monitoring. You can monitor the system through IBM Storage Scale System GUI, RESTful API, or with the help of the CLI.

For more information, search for the topic "Monitoring IBM Storage Scale RAID" in the [IBM Storage Scale System Software documentation](#).

For Storage Scale System v 6.2.3.0, the IBM Storage Scale System GUI provides robust facilities to monitor, manage, and upgrade the system. For more information, see the [IBM Storage Scale System Software documentation](#).

IBM Documentation provides more information about how you can configure the IBM Storage Scale System GUI to monitor the performance of the following functional areas in the system:

- Network
- System resources
- IBM Storage Scale RAID
- NSD server
- IBM Storage Scale client
- NFS
- SMB
- Object
- Waiters
- Active File Management (AFM)

For more information, see *Monitoring and Managing the IBM Elastic Storage Server Using the GUI*, REDP-5471.

Maintaining and repairing IBM Storage Scale System

For IBM Storage Scale System features, a full library of documentation for maintaining the hardware and software of the IBM Storage Scale System solution is available at the [IBM Storage Scale System Software documentation](#). At the website, select the IBM Storage Scale System solution release level that you are using.

IBM provides service and support for IBM Storage Scale System as a solution. IBM Storage Scale System-specific keywords and procedures are in place so that you, as a client, when calling or electronically requesting IBM Storage Scale System support, are routed to an IBM Storage Scale System cross-component solution support team that understands the interaction of all of the IBM Storage Scale System components. They can also diagnose and respond to your query in a holistic fashion, and where necessary, direct your query to suitable Level 2 or Level 3 resources.

For more information, see the [IBM Storage Scale System Support Reference Guide](#).

Upgrading IBM Storage Scale System

Maintaining and upgrading your IBM Storage Scale System software solution stack is a customer's responsibility.

You can also choose to engage the IBM Technology Expert Lab or an IBM Business Partner to perform the upgrades for you.

The IBM Storage Scale System solution is an integrated package of hardware, firmware, and software. All components of an IBM Storage Scale System are integrated, tested, and supported by IBM as a solution. IBM periodically provides a refreshed or up-leveled IBM Storage Scale System software solution stack.

Existing IBM Storage Scale System customers are entitled to download the newer IBM Storage Scale System solution stack from IBM Fix Central. IBM recommends that you plan to upgrade your IBM Storage Scale System solution stack at least once a year.

For hardware upgrades, you should contact and work with your IBM or Business Partner specialist to plan and run the IBM Storage Scale System upgrades when necessary.

IBM Storage Scale System hardware and firmware

IBM Storage Scale System hardware and firmware updates are included in the IBM Storage Scale System solution stack. Upgrading to a new level of IBM Storage Scale System solution stack also upgrades the server and storage firmware levels, if necessary.

Typically, the IBM Storage Scale System upgrades are done non-disruptively. During a planned maintenance window, fail over the storage to one of the I/O data servers, upgrade the I/O data server, and then fail over the storage to this newly upgraded server and upgrade the other server.

Storage enclosure firmware can typically be applied non-disruptively.

IBM Storage Scale System also supports nondisruptive addition of more storage enclosures, which allows an IBM Storage Scale System to be non-disruptively upgraded to a larger capacity.

The client is responsible for keeping the IBM Storage Scale System operating system, software, firmware, and GUI current with the recommended levels of software. For more information about upgrade instructions, see the [IBM Storage Scale System Software documentation](#).

IBM Storage Scale System software

IBM Storage Scale System provides a fully integrated and tested solution stack of software that consists of the following components at a specific release and fix level (including all prerequisites) that were fully integrated and tested:

- RHEL
- IBM Storage Scale
- Cluster Export Services code (CES) (SMB, NFS, and S3 Object Storage)
- IBM x86 server firmware
- RHEL kernel at specific level
- RHEL system errata at specific level
- RHEL Network Manager at specific level
- RHEL device drivers
- Nvidia OpenFabrics Enterprise Distribution (OFED) network adapter drivers
- Electronic Service Agent (ESA)

Keeping the IBM Storage Scale System software solution stack current is a customer responsibility. IBM recommends that clients plan to upgrade their IBM Storage Scale System software solution stack at least once a year. Typically, the IBM Storage Scale System hardware and software solution stack upgrades are performed at the same time.

For more information about the levels of the software components for each IBM Storage Scale System software solution release level, see the [IBM Storage Scale System Software documentation](#).

Upgrading the IBM Storage Scale System solution software is done by a full replacement of the software solution stack with a newer level.

New levels of the IBM Storage Scale System solution software stack are available at IBM Fix Central. For more information about these newer levels of IBM Storage Scale System software, log on to [IBM Fix Central](#) by using your IBMid. At the website, in the Product Selector field, enter IBM Storage Scale System. Then, select your IBM Storage Scale System version or release level.

A list of available IBM Storage Scale System solution software fix packs for download is shown. When you start the download, IBM FixCentral checks your IBMid authentication profile to ensure that you set up your IBMid with the suitable authentication and client customer number information. This information authenticates that you are authorized to download IBM Storage Scale System solution software.

For more information about applying the IBM Storage Scale System Fix Pack, see the [IBM Storage Scale System Software documentation](#).

IBM Technology Expert Lab is available to perform the IBM Storage Scale System software upgrade for you. As a best practice, take advantage of these services. For more information, contact your IBM representative.

IBM Storage Scale System models

This section includes the differences between the three generations of IBM Storage Scale Systems (IBM Storage Scale System) that are used today. Understanding the nomenclature and the differences between these generations is helpful when you are designing your IBM Storage Scale System environment, especially if different generations of IBM Storage Scale System coexist in the same IBM Storage Scale cluster.

IBM Storage Scale System model specifications

IBM Storage Scale System 6000 specifications

The specifications of the IBM Storage Scale System 6000 model are listed in the following table.

IBM Storage Scale System Model	Enclosure U/Number drives IBM machine type - model	Drive Count	Drive size
IBM Storage Scale System 6000	4U48 Machine type 5149-F48	24 or 48	3.84 TB 2.5 Inch PCIe Gen4 NVMe Flash Drive 7.68 TB 2.5 Inch PCIe Gen4 NVMe Flash Drive 15.36 TB 2.5 Inch PCIe Gen4 NVMe Flash Drive 30.72 TB 2.5 Inch PCIe Gen4 NVMe Flash Drive 19.2TBu NVMe FlashCore Module 38.4TBu NVMe FlashCore Module 3.84 TB U.3 PCIe 5.0 NVMe SSD 7.68 TB U.3 PCIe 5.0 NVMe SSD 15.36 TB U.3 PCIe 5.0 NVMe SSD 30.72 TB U.3 PCIe 5.0 NVMe SSD 30.72 TB 2.5 Inch QLC NVMe SSD 61.44 TB 2.5 Inch QLC NVMe SSD 122 TB 2.5 Inch QLC NVMe SSD



Figure: IBM Storage Scale System 6000 fully populated with 24 TLC and 24 QLC drives



Figure: IBM Storage Scale System 6000 fully populated with 12 TLC and 36 QLC drives




Figure: IBM Storage Scale System 6000 back panel view and connections

IBM Storage Scale System 3500 specifications

The IBM Storage Scale System 3500 is one of the third-generation IBM Storage Scale System offerings.

The IBM Storage Scale System 3500 model includes NVMe flash storage.

The specifications of the IBM Storage Scale System 3500 model are listed in the following table.

IBM Storage Scale System model	Enclosure U/Number drives IBM Machine Type - Model	NVMe drives	Drive size
IBM Storage Scale System 3500 or IBM Storage Scale System 3500 H0	2U24 Machine type 5141-FN2	NVMe 12 or 24	3.84 TB 7.68 TB 15.36 TB 30.72 TB




Figure: IBM Storage Scale System 3500 rear panel view and connections

The IBM Storage Scale System 3500Hx, shown in the following table, includes machine type 5141-FN2 and 1-4 model 5147-102.

IBM Storage Scale System Model	Enclosure U/ Number drives IBM machine type - model	NVMe	HDD enclosure	HDD drives	Drive size
IBM Storage Scale System 3500H1	2U24 Machine type 5141-FN2 + 1@4U external storage Model 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	1	52 or 102	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H2	2U 5141-FN2 + 2@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	2	204	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H3	2U 5141-FN2 + 3@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	3	306	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H4	2U 5141-FN2 + 4@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	4	408	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H5	2U 5141-FN2 + 5@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	5	510	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H6	2U 5141-FN2 + 6@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	6	612	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H7	2U 5141-FN2 + 7@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	7	714	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500H8	2U 5141-FN2 + 8@4U 5147-102	Quantity: 12 or 24 Size: 3.84 TB 7.68 TB 15.36 TB 30.72 TB	8	816	HDD: 10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB

The IBM Storage Scale System 3500Cx, shown in the following table, includes machine type 5141-FN2 logtip only NVMe drives and 1—8 model 5147-102 enclosures.

IBM Storage Scale System Model	Enclosure U/ Number drives IBM machine type - model	NVMe	HDD enclosure	HDD drives	HDD Drive size
IBM Storage Scale System 3500C1	2U24 Machine type 5141-FN2 + 1@4U external storage Model 5147-102	Quantity: 4 (used for logtip)	1	52 or 102	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C2	2U 5141-FN2 + 2@4U 5147-102	Quantity: 4 (used for logtip)	2	204	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C3	2U 5141-FN2 + 3@4U 5147-102	Quantity: 4 (used for logtip)	3	306	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C4	2U 5141-FN2 + 4@4U 5147-102	Quantity: 4 (used for logtip)	4	408	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C5	2U 5141-FN2 + 5@4U 5147-102	Quantity: 4 (used for logtip)	5	510	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C6	2U 5141-FN2 + 6@4U 5147-102	Quantity: 4 (used for logtip)	6	612	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C7	2U 5141-FN2 + 7@4U 5147-102	Quantity: 4 (used for logtip)	7	714	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB
3500C8	2U 5141-FN2 + 8@4U 5147-102	Quantity: 4 (used for logtip)	8	816	10 TB, 14 TB, 18 TB, 20 TB, 22 TB 24 TB

IBM Storage Scale RAID technical overview

The following section provides a high-level, technical overview of the IBM Storage Scale RAID that is used in all IBM Storage Scale System models.

The IBM Storage Scale RAID software that is used in the IBM Storage Scale System solution runs on SAS disks in just a bunch of disks (JBOD) arrays. IBM Storage Scale RAID on IBM Storage Scale System provides JBOD cost reduction while simultaneously providing enterprise class reliability.

Different IBM Storage Scale System models provide solid-state drives (SSDs or NVMe) when more performance is needed. The IBM Storage Scale System solution does not require or use any external RAID controller or acceleration.

IBM Storage Scale RAID supports multiple RAID codes and distributes client data, redundancy information, and spare space across the disks in such a way that if a physical disk loss or even a group of physical disk loss occurs, it does not affect data availability.

Instead of relying on the disks or conventional RAID controller only, IBM Storage Scale RAID implements erasure coding, which is embedded within the IBM Storage Scale file system to provide reliability and high performance. This configuration detects and reports storage media faults and read or write errors, and resolves other data integrity problems.

IBM Storage Scale RAID also implements an end-to-end checksum from the storage media all the way out to the IBM Storage Scale client.

IBM Storage Scale RAID erasure codes

IBM Storage Scale RAID in the IBM Storage Scale System supports different data protection algorithms and can detect and correct up to two or three concurrent storage media faults per erasure code array.




Figure: IBM Storage Scale RAID tracks

End-to-end checksum

If you use the IBM Storage Scale software client to access data on the IBM Storage Scale System, the IBM Storage Scale client is aware if the IBM Storage Scale file system is based on IBM Storage Scale RAID network shared disks. If this configuration is used, an 8-byte checksum is calculated during a write operation, which is appended to the data, and sent over the network to the IBM Storage Scale RAID server. The checksum is verified and then IBM Storage Scale RAID writes the data along with its checksum on the disks and logs the version number on its metadata.

When a read operation is requested, IBM Storage Scale RAID verifies the checksum and version on its metadata. If verification is successful, it sends the data to the client. If verification fails, the data is rebuilt based on parity or replication and then sent to the client along with a newly generated checksum.

The end-to-end checksum feature provides a robust means to prevent and correct silent disk errors or missing disk writes. This feature provides an essential level of reliability to ensure data integrity and prevent file system checks, which ensure reliability and uptime for ever larger petabyte-scale file systems.

Declustered RAID arrays

IBM Storage Scale RAID implements its own data and spare disk layout scheme that reduces overhead and mitigates the performance impact to users when recovering from disk failures. IBM Storage Scale RAID spreads or declusters user data, redundancy information, and spare space across all the disks of the array instead of leaving all spare space in a single disk. A conventional 1+1 RAID layout is compared to a declustered array in the following figure.




Figure: Declustered array versus a 1+1 array

Consider an example where seven stripes of data are on each disk. The declustered array diagram above shows the left three arrays of two disks in a replicated 1+1 configuration and a spare. On the left, you can see the data stripes spreading all over the seven disks of the declustered array.

In a failure of one disk, all data from the remaining disks of the array must be replicated to the spare disk on the traditional 1+1 array. On the declustered array, the replication occurs on spare space of all the remaining disks, which can decrease the rebuild impact from three to four times.

The IBM Storage Scale RAID overhead reduction is shown in the following figure.




Figure: Array rebuild operation

On the IBM Storage Scale System, using RAID 8+2 or 2-way replication (1+2), if one disk loss occurs, the rebuild operation starts with low priority with even lower impact for the clients.

With this array configuration, the rebuild is considered critical and run on high priority only if two concurrent disk losses occur.

By using 8+3 RAID or 3-way replication (1+3), the rebuild operation becomes critical only if three concurrent disk losses occur in the same declustered array.

Abbreviations

Abbreviation	Full Name
AFM	Active File Management
AI	Artificial Intelligence
API	Application Programming Interface
BMC	Baseboard Management Controller
CES	Cluster Export Services
CNSA	Container Native Storage Access
CPU	Central Processing Unit
CSI	Container Storage Interface
CUDA	Compute Unified Device Architecture
DAT	Data Acceleration Tier
DHCP	Dynamic Host Configuration Protocol
DL	Deep Learning
DPU	Data Processing Unit
EIA	Electronic Industries Alliance
EMS	Enterprise Management Server
ESA	Electronic Service Agent
ESS	Elastic Storage System
FCM	FlashCore Module
FCM4	FlashCore Module 4
FSP	Flexible Service Processor
GbE	Gigabit Ethernet
GBps	Gigabytes per Second
GNR	GPFS Native RAID

Abbreviation	Full Name
GPFS	General Parallel File System
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HA	High Availability
HBA	Host Bus Adapter
HDD	Hard Disk Drive
HDR	High Data Rate
HPC	High-Performance Computing
I/O	Input/Output
IBM	International Business Machines
IOPS	Input/Output Operations Per Second
IP	Internet Protocol
IT	Information Technology
JBOD	Just a Bunch of Disks
JBOF	Just a Bunch of Flash
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LED	Light-Emitting Diode
MES	Miscellaneous Equipment Specification
MTM	Machine Type Model
NDR	Next Data Rate
NFS	Network File System
NIC	Network Interface Card

Abbreviation	Full Name
NL-SAS	Nearline Serial Attached SCSI
NSD	Network Shared Disk
NVMe	Non-Volatile Memory Express
NVMe-oF	NVMe over Fabrics
OFED	OpenFabrics Enterprise Distribution
PB	Petabyte
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PDU	Power Distribution Unit
POSIX	Portable Operating System Interface
QLC	Quad-Level Cell
RAID	Redundant Array of Independent Disks
RAS	Reliability, Availability, and Serviceability
RDMA	Remote Direct Memory Access
RHEL	Red Hat Enterprise Linux
RJ45	Registered Jack 45
RoCE	RDMA over Converged Ethernet
S3	Simple Storage Service
SAS	Serial Attached SCSI
SATA	Serial Advanced Technology Attachment
SDS	Software-Defined Storage
SED	Self-Encrypting Drive
SMB	Server Message Block

Abbreviation	Full Name
SSH	Secure Shell
SSD	Solid-State Drive
TB	Terabyte
TDA	Technical and Delivery Assessment
TiB	Tebibyte
TLC	Triple-Level Cell
U	Rack Unit
USB	Universal Serial Bus
VGA	Video Graphics Array
VLAN	Virtual Local Area Network

Additional Resources

This document complements but does not replace IBM product documentation. For detailed API references, complete configuration options, and the latest updates, consult the following additional resources:

IBM Redbooks Publications

- [IBM Storage Scale System Introduction Guide, REDP-5729](#) — Previous editions of this publication.
- [Implementation Guide for IBM Elastic Storage System 3500, SG24-8538](#)
- [Implementation Guide for IBM Storage Scale System 6000, SG24-8566](#)
- [Monitoring and Managing the IBM Elastic Storage Server Using the GUI, REDP-5471](#)

IBM Documentation

- [IBM Storage Scale System Software documentation](#) — Full documentation library for IBM Storage Scale System software, covering installation, administration, deployment, and maintenance.
- [IBM Storage Scale System product page](#) — Top-level IBM Storage Scale System documentation hub.
- [IBM Storage Scale System 3500 documentation](#)
- [IBM Storage Scale System 5149-F26 Storage Enclosure documentation](#)
- [IBM Storage Scale System 6000 PDF files](#) — PDF versions of the IBM Storage Scale System 6000 Hardware Planning and Installation Guide and Quick Deployment Guide.

Other Related Documentation

Notices

This information was developed for products and services offered in the US. This material might be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 United States of America

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Trademarks

The following terms are trademarks of the International Business Machines Corp.

AI Attribution

This work was primarily human-created. AI was used to make stylistic edits, such as changes to structure, wording, and clarity. AI was used to make content edits, such as changes to scope, information, and ideas. AI was prompted for its contributions, or AI assistance was enabled. AI-generated content was reviewed and approved. The following model(s) or application(s) were used: IBM Bob.

Version History

Version 3.0 - Third Edition (March 2026)

Document Number: MD260040

Previous Document Number: REDP-5729-02 (converted from IBM Redpaper to IBM Redbooks format)

Title: IBM Storage Scale System Introduction Guide

Authors:

- Phillip Gerrard
- Sagrika Arora
- Nilton Carlos Dos-Santos
- Jordi Caubet
- Mike Kieran
- Stieg Klein
- Chris Maestas
- Bis Purkayastha

Publication Date: March 2026

Content:

- Chapter 1: Introducing the IBM Storage Scale System
- Chapter 2: IBM Storage Scale System Architecture
- Chapter 3: IBM Storage Scale System Planning and Integration
- Appendix A: IBM Storage Scale System Models

Summary: Third edition of the IBM Storage Scale System Introduction Guide. Provides an overview of the IBM Storage Scale System family of products, including architecture, hardware and software components, planning and integration guidance, and detailed model specifications.

Summary of changes:

- Updated document to include information related to IBM Storage Scale 7.0.0 and IBM Storage Scale System 6.0.0
- Added information related to support for Nvidia BlueField-3 Data Processing Units (DPUs) in flash storage enclosures
- Added information related to Just a Bunch of Flash (JBOF) support
- Added information related to IBM Storage Scale native REST API
- Added information related to deployment updates: GUI-based deployment, Fix Central deployment of protocol nodes, AFM deployment by using automation and essrun
- Updated information related to supported disk and flash storage devices
- Updated information related to IBM Storage Scale 6000 and 3500 configurations, including IBM Storage Scale System 7.0.1 features
- Updated information related to EMS and Utility node devices

Version 2.0 - Second Edition (July 2025)

Document Number: REDP-5729-01

Title: IBM Storage Scale System Introduction Guide

Authors:

- Phillip Gerrard
- Sagrika Arora
- Nilton Carlos Dos-Santos
- Jordi Caubet
- Stieg Klein
- Chris Maestas
- Bis Purkayastha

Publication Date: July 2025

Summary of changes:

- Added information related to IBM® Storage Scale System 6000 and its features
- Added information related to Just a Bunch of Disk (JBOD) storage expansion
- Removed information related to IBM Storage Scale System models that are no longer offered

Version 1.0 - First Edition (2024)

Document Number: REDP-5729-00

Title: IBM Storage Scale System Introduction Guide

Authors:

- Dino Quintero
- Puneet Chaudhary
- Lee Nee Helgeson
- Brian Herr
- Steven Normann
- Marcelo Ramos
- Richard Rosenthal
- Robert Simon
- Marian Tomescu
- Richard Wale

Publication Date: 2024

Summary: Initial publication providing an overview of the IBM Storage Scale System, its architecture, hardware and software components, and planning and integration guidance.