Redbooks

ibm.com/redbooks

IBM Elastic Storage System Introduction Guide

Stieg Klein

Chris Maestas

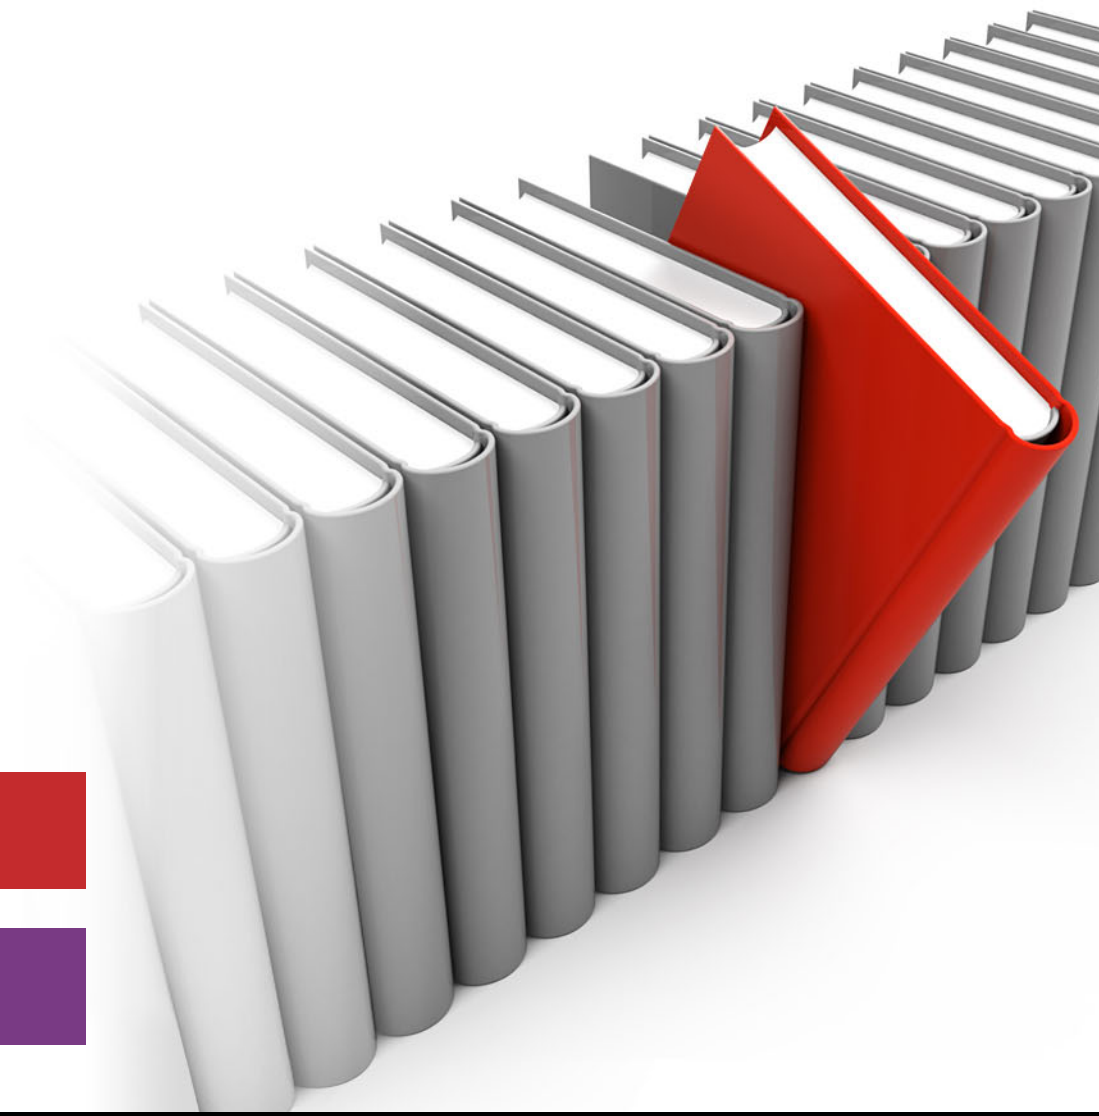

Cloud

Storage

IBM

Redpaper

International Technical Support Organization

IBM Elastic Storage System Introduction Guide

November 2022

REDP-5253-04

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Fifth Edition (November 2022)

This edition applies to IBM Spectrum Scale and IBM Elastic Storage System (IBM ESS), and IBM Elastic Storage Server (also IBM ESS).

Contents

Notices

This information was developed for products and services offered in the U.S.A.

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, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINIBM ESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRIBM ESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNIBM ESS FOR A PARTICULAR PURPOSE. Some states 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 supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®	IBM Spectrum®	POWER9™
IBM Elastic Storage®	PartnerWorld®	Redbooks®
IBM FlashCore®	Passport Advantage®	Redbooks (logo) ®
IBM FlashSystem®	POWER®	
IBM Services®	POWER8®	

The following terms are trademarks of other companies:

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

Red Hat, Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redpaper publication provides an overview of the IBM Elastic Storage® Server (IBM ESS) and IBM Elastic Storage System (also IBM ESS). These scalable, high-performance data and file management solution, are built on IBM Spectrum® Scale technology. Providing reliability, performance, and scalability, IBM ESS can be implemented for a range of diverse requirements.

The latest IBM ESS 3500 is the most innovative system that provides investment protection to expand or build a new Global Data Platform and use current storage. The system allows enhanced, non-disruptive upgrades to grow from flash to hybrid or from hard disk drives (HDDs) to hybrid. IBM ESS can scale up or out with two different storage mediums in the environment, and it is ready for technologies like 200 Gb Ethernet or InfiniBand NDR-200 connectivity.

This publication helps you to understand the solution and its architecture. It describes ordering the best solution for your environment, planning the installation and 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 Elastic Storage Server and IBM Elastic Storage System components is key for planning an environment.

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 big data solutions. The content of this paper can help you to uncover insights among client's data so that you can take appropriate actions to optimize business results, product development, and scientific discoveries.

Authors

This paper was produced by a team of specialists from around the world working with International Technical Support Organization (ITSO), Tucson Center.

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 Elastic Storage System (IBM ESS) and IBM Spectrum 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 also has led global enablement sessions online and face to face, where he described how best to position mature technologies like IBM Spectrum Scale with emerging technologies in cloud, object, container, or AI spaces.

Thanks to the following people for their contributions to this project:

Larry Coyne and Phillip Gerrard
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, John Sing, Marian Tomescu, Richard Wale

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks® publications in one of the following ways:

- Use the online Contact us review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Look for us on LinkedIn:

  `http://www.linkedin.com/groups?home=&gid=2130806`

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  `https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm`

- Stay current on recent Redbooks publications with RSS Feeds:

  `http://www.redbooks.ibm.com/rss.html`

Summary of changes

This section describes the technical changes that were made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for IBM Elastic Storage System Introduction Guide
as created or updated on December 1, 2022.

November 2022, Fifth Edition

This revision includes the following new and changed information.

New information

Introduction of the IBM Elastic Storage System 3500 (IBM ESS 3500). The IBM ESS 3500 provides a high-performance tier of IBM Spectrum Scale file storage for a broad variety of artificial intelligence (AI), analytics, and big data applications. The IBM ESS 3500 also can be configured for high-capacity hard disk drive (HDD)-based storage, or a blend of both high performance and capacity. For more information, see 2.2.3, “IBM Elastic Storage System 3500” on page 20.

Changed information

Added IBM ESS 3500 to functions and specifications like IBM ESS 3000, IBM ESS 3200, and IBM ESS 5000. Reorganized sections and lists with newer models first. Reduced some content for older first-generation models.

July 2021, Fourth Edition

This revision includes the following new and changed information.

New information

Introduction of the IBM Elastic Storage System 3200 (IBM ESS 3200). The IBM ESS 3200 provides a high-performance tier of IBM Spectrum Scale file storage for a broad variety of AI, analytics, and big data applications. For more information, see “IBM Elastic Storage System 3200” on page 81.

Changed information

Added IBM ESS 3200 to functions and specifications similar to IBM ESS 3000 and IBM ESS 5000.

March 2021, Third Edition (minor updates)

This revision includes the following changed information.

Changed information

Updated to reflect announcement of additional IBM ESS 5000 models SL7 and SC9, enhanced IBM POWER9™ processor-based IBM ESS Management Server (IBM EMS) support for second-generation IBM ESS, IBM ESS 5000 mode and capacity upgrades, and High Dynamic Range (HDR) 100-Gb network cards.

February 2021, Third Edition

This revision includes the following changed information.

Changed information

Added IBM ESS 5000 model with POWER9 processor-based data servers with 5147-092 or 5147-106 storage enclosures.

July 2020, Second Edition

This revision includes the following changed information.

Changed information

Updated correct 5148-21L model number on pages 22, 29, and 47.

1

Introducing the IBM Elastic Storage System

This chapter introduces the IBM Elastic Storage Server (IBM ESS) and IBM Elastic Storage System (also IBM ESS) solutions. In this chapter, we describe their characteristics, value, how they integrate with IBM Spectrum Scale, and where they fit in today's business environments.

The latest IBM Elastic Storage Server 3500 is the most innovative system that provides investment protection to expand or build a new Global Data Platform and use current storage. The system allows enhanced, non-disruptive upgrades to grow from flash to hybrid or from hard disk drives (HDDs) to hybrid. IBM ESS can scale up or out with two different storage mediums in the environment, and it is ready for technologies like 200 Gb Ethernet or InfiniBand NDR-200 connectivity.

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

This chapter includes the following topics:

Note: In this IBM Redpaper publication, *IBM ESS* can refer to any of the following products:

- IBM Elastic Storage Server (the first and second generations of this solution).
- IBM Elastic Storage System (the third generation of this solution includes the IBM Elastic Storage System 5000 (IBM ESS 5000), IBM Elastic Storage System 3500 (IBM ESS 3500), IBM Elastic Storage System 3200 (IBM ESS 3200), and IBM Elastic Storage System 3000 (IBM ESS 3000).
- For more information about all the current IBM ESS products, see IBM Elastic Storage 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

The general capabilities of these solutions are the same. Where suitable, we differentiate between the models and generations.

1.1 IBM Elastic Storage System

The IBM Elastic Storage System and IBM Elastic Storage Server (both referred to as IBM ESS) are integrated IBM Spectrum Scale storage solutions for high-performance, reliable file, and object storage. Scalable to petabyte levels, they are suitable for analytics, big data, artificial intelligence, and enterprise data ocean requirements. IBM ESS provides scalable storage for IBM Spectrum Scale (see Figure 1-1), which is the IBM parallel clustered file system. The IBM ESS 3000 was released in December 2019. The IBM ESS 3000 is the first of the generation three IBM ESS models. Generation three models starting with the most recent one includes the IBM ESS 3500, IBM ESS 3200, IBM ESS 5000, and IBM ESS 3000. The prior generation of IBM ESS models were released under the product name IBM Elastic Storage Server.

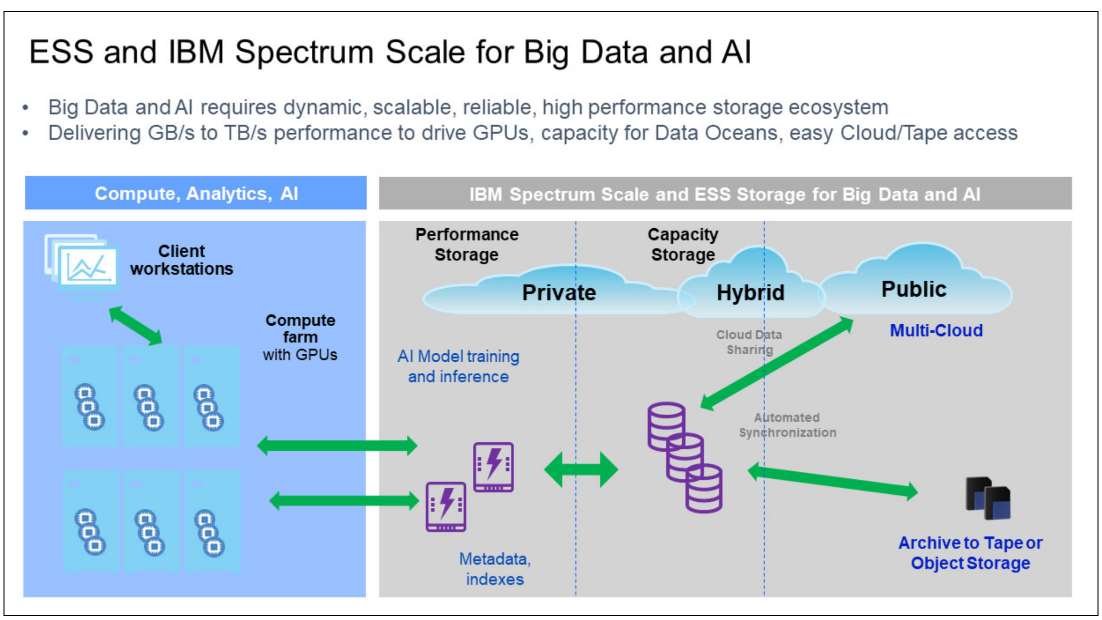

Figure 1-1 IBM Elastic Storage System integrated with IBM Spectrum Scale solutions

IBM ESS is a pair of cross-connected IBM Spectrum Scale Data Servers, which are deployed in a fully integrated, tested storage building block solution (see Figure 1-2). Instead of hardware-based disk RAID controllers, IBM ESS uses IBM Spectrum Scale RAID, which uses declustered erasure code technology to deliver consistent high performance. This feature mitigates storage drive failures and RAID rebuild overhead.

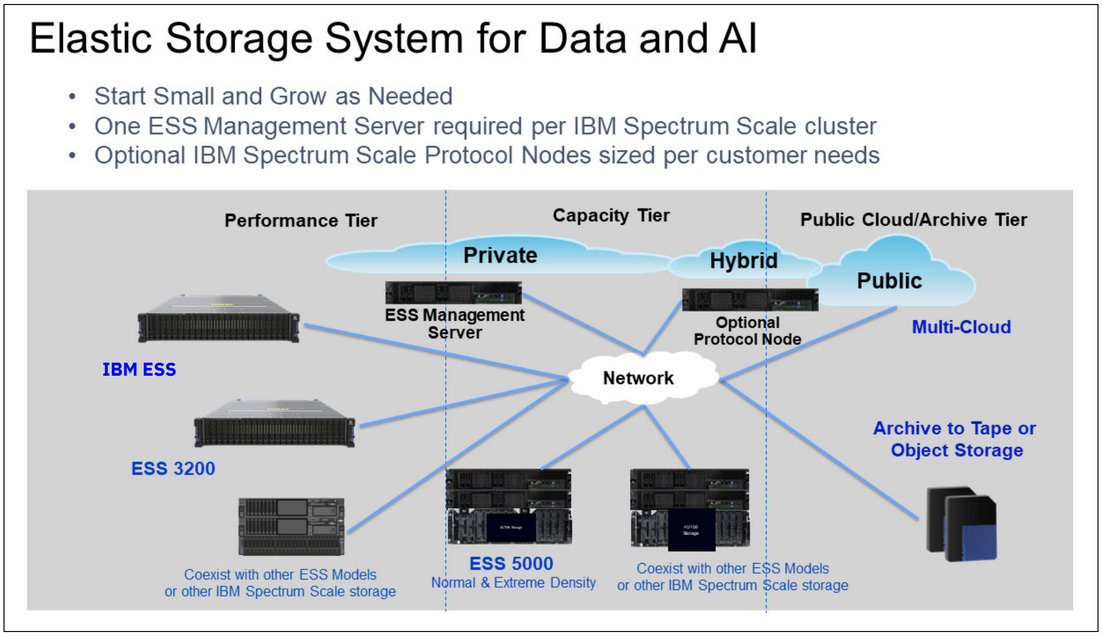

Figure 1-2 Third-generation IBM ESS integrated with IBM Spectrum Scale

IBM ESS with IBM Spectrum Scale RAID also provides superior data protection by reducing rebuild times to a fraction of the time that is needed with conventional hardware-based RAID controllers.

Data access to IBM ESS storage is available by using one of the following methods:

- IBM Spectrum Scale clients on user workstations, which provide a POSIX file system interface to data on IBM ESS.
- IBM Spectrum Scale protocol nodes, which allow IBM ESS data access through Server Message Block (SMB), Network File System (NFS) file protocols, or object data access through the IBM Spectrum Scale implementation of OpenStack Swift. In addition, Hadoop clusters can access IBM ESS data by using the IBM Spectrum Scale Hadoop Transparency Connector.

The IBM ESS is a software-defined storage (SDS) building block that uses the IBM Spectrum Scale high-performance parallel clustered file system (suitable for petabyte scale data storage).

IBM ESS integrates IBM Spectrum Scale, IBM Storage enclosures, IBM Power servers, or x86 servers (depending on the IBM ESS model) into a complete management and installation software stack. It also provides an integrated tested IBM Spectrum Scale storage building block solution. IBM Service supports IBM ESS as an integrated solution.

An IBM ESS base building block consists of a 1 Gb Ethernet (GbE) network switch, servers for IBM ESS Management nodes and I/O nodes or data server nodes, and I/O storage enclosures. Multiple IBM ESS building blocks can be integrated into an IBM Spectrum Scale cluster.

1.2 IBM Elastic Storage System and IBM Spectrum Scale

In this section, we describe how IBM ESS is part of an IBM Spectrum Scale cluster.

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

In the following sections, we provide a brief overview of the following topics:

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

Note: An IBM ESS must be installed and configured, including integration into an IBM Spectrum Scale cluster, before it can be used to hold user data. In isolation, an IBM ESS is not sufficient for an IBM Spectrum Scale cluster, so you must use an IBM ESS Management Server (IBM EMS) too.

Here, we provide a high-level IBM Spectrum Scale and IBM ESS solution overview. For more information, see Chapter 2, "IBM Elastic Storage System architecture" on page 15.

1.2.1 IBM Spectrum Scale overview

IBM Spectrum Scale is IBM's strategic high-performance parallel file system shared storage platform for end-to-end collaborative common enterprise, data platform, big data analytics, and AI workflows (see Figure 1-3).

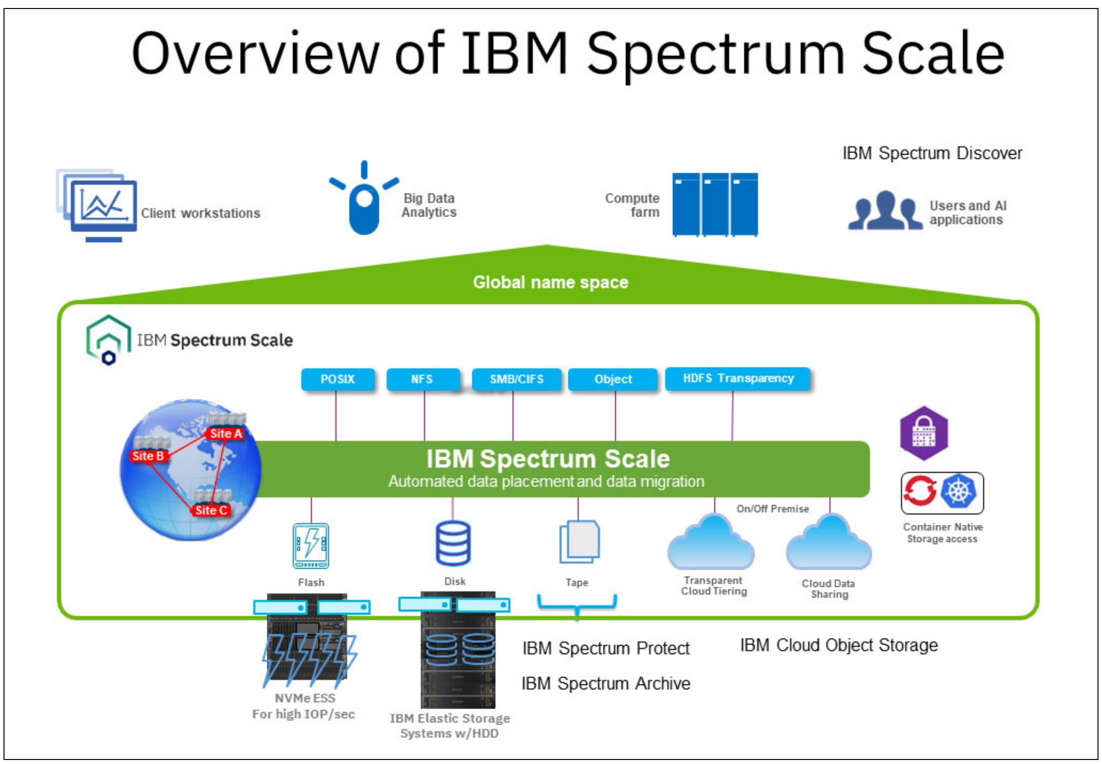

Figure 1-3 IBM Spectrum Scale overview

IBM Spectrum Scale

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

- Simplifies data management by supporting enterprise workflows on a single common enterprise data platform
- Supplies a single global namespace that supports enterprise-level data over high-performance networks
- Enables intelligent automatic tiering of data between storage pools, and externally to tape, object and cloud resources, which deliver cost-effective storage economics by automatically managing and tiering data to different classes of storage

Although multiple types of IBM Spectrum Scale cluster configurations are available, the configuration into which IBM ESS is commonly deployed is the IBM Spectrum Scale Network Shared Disk (NSD) configuration, as shown in Figure 1-4.

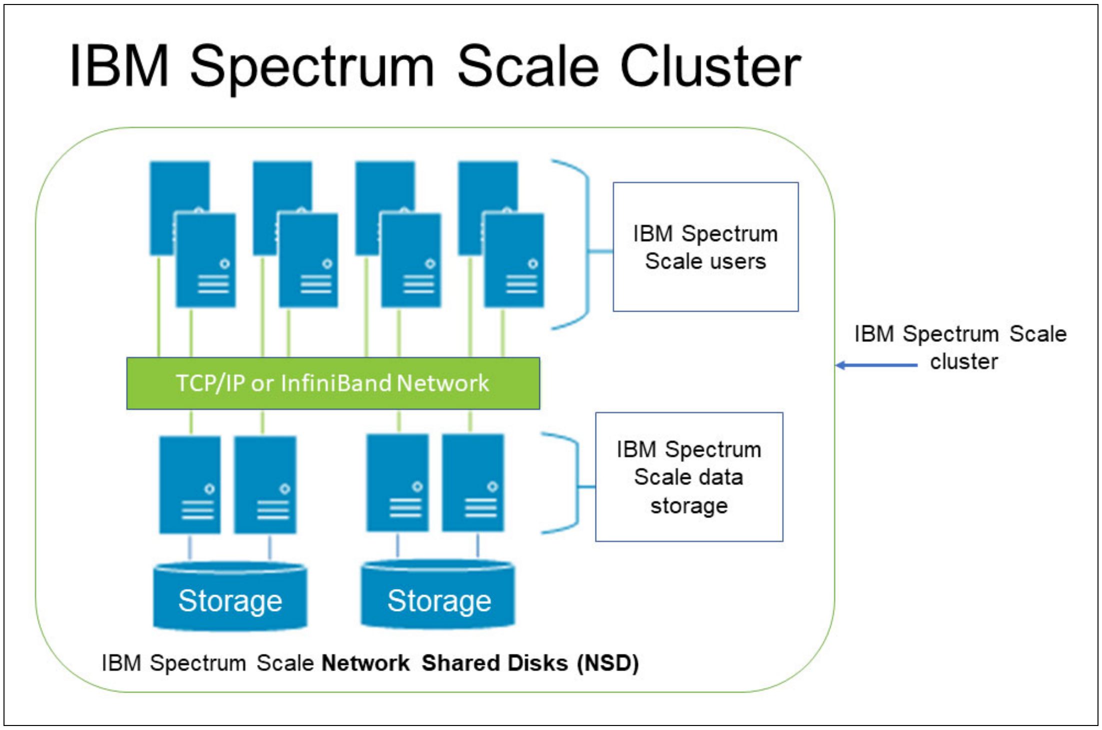

Figure 1-4 IBM Spectrum Scale cluster

IBM ESS is a pair of IBM Spectrum Scale NSD Data Servers, which are configured together as a tested, integrated, highly available (HA), and reliable IBM Spectrum Scale storage building block solution.

As shown in Figure 1-4 on page 5, eight IBM Spectrum Scale nodes are application workstations, servers, or users. Four nodes also are IBM Spectrum Scale data server nodes. The user workstations are running the IBM Spectrum Scale client in this IBM Spectrum Scale configuration. The IBM Spectrum Scale client provides multi-threaded, highly parallel access data on the IBM Spectrum Scale Data Servers, as shown in Figure 1-5.

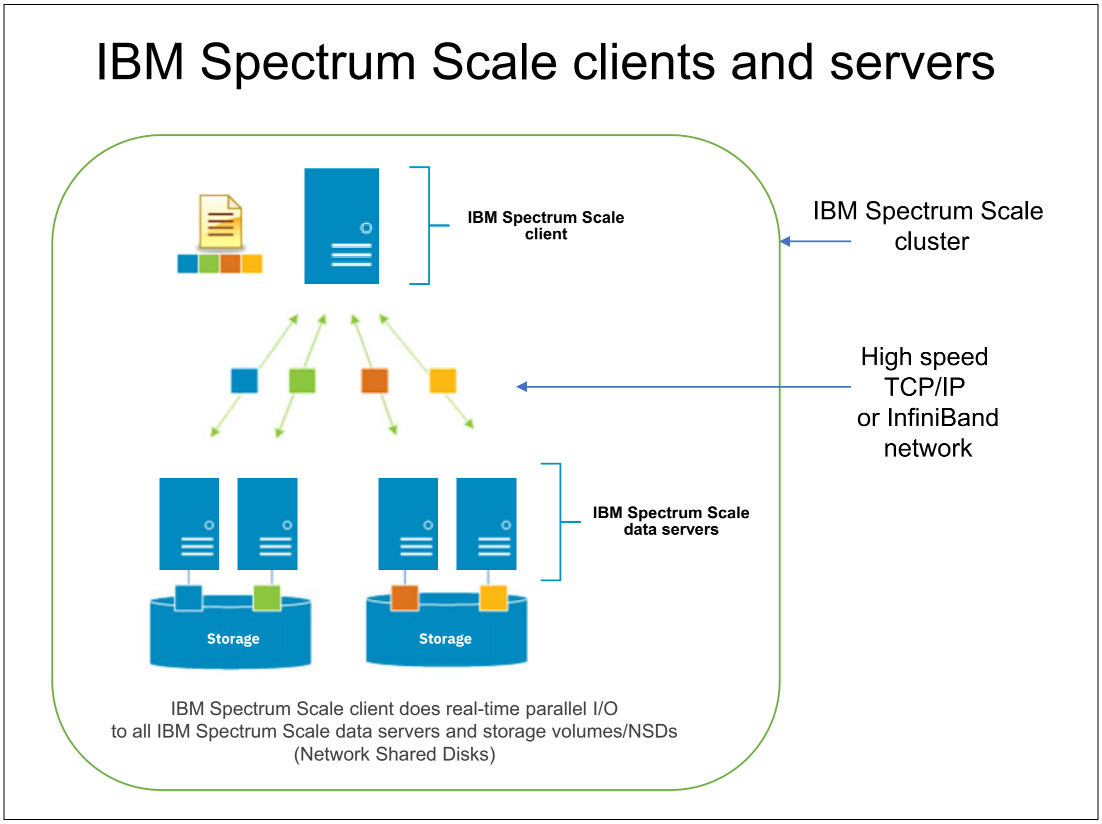

Figure 1-5 IBM Spectrum Scale clients and servers

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

An IBM Spectrum Scale cluster can grow by adding nodes, whether they are IBM Spectrum Scale clients or IBM Spectrum Scale data servers, as shown in Figure 1-6.

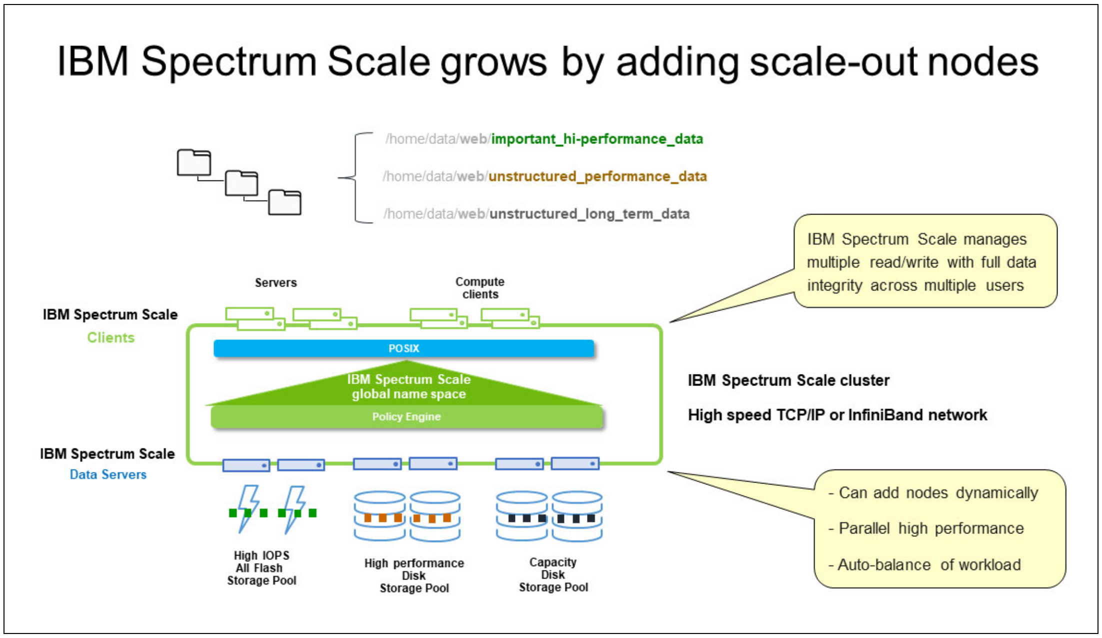

Figure 1-6 IBM Spectrum Scale grows by adding scale-out nodes

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

IBM Spectrum Scale users are unaware of the physical distribution of data in the IBM Spectrum Scale data server physical storage pools. The automatically balanced data distribution is seamlessly determined by the IBM Spectrum 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 Spectrum Scale parallel file system provides an enterprise capability for data management of large amounts of data, while also performing 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 Spectrum Scale data is to install the IBM Spectrum Scale client on every workstation or server that accesses IBM Spectrum Scale data. The IBM Spectrum Scale client provides the multiple threads and communication with multiple data servers to provide enterprise high-performance parallel throughput. While doing so, IBM Spectrum Scale also manages full read/write data integrity between multiple users who are working with the data in the file system.

1.2.2 Protocol nodes overview

Protocol nodes are part of the IBM Spectrum Scale solution. You can access IBM Spectrum Scale data by using SMB, NFS, or object protocols without installing IBM Spectrum Scale software on a node. To provide this access, IBM Spectrum Scale provides the protocol node functions.

The value of protocol nodes is that applications, workstations, and users that do not have the IBM Spectrum Scale client can still access IBM Spectrum Scale data through an SMB, NFS, or object protocol. IBM Spectrum Scale often is used as an enterprise data lake or central enterprise data repository.

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

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

Figure 1-7 shows protocol nodes that were added to the IBM Spectrum Scale cluster.

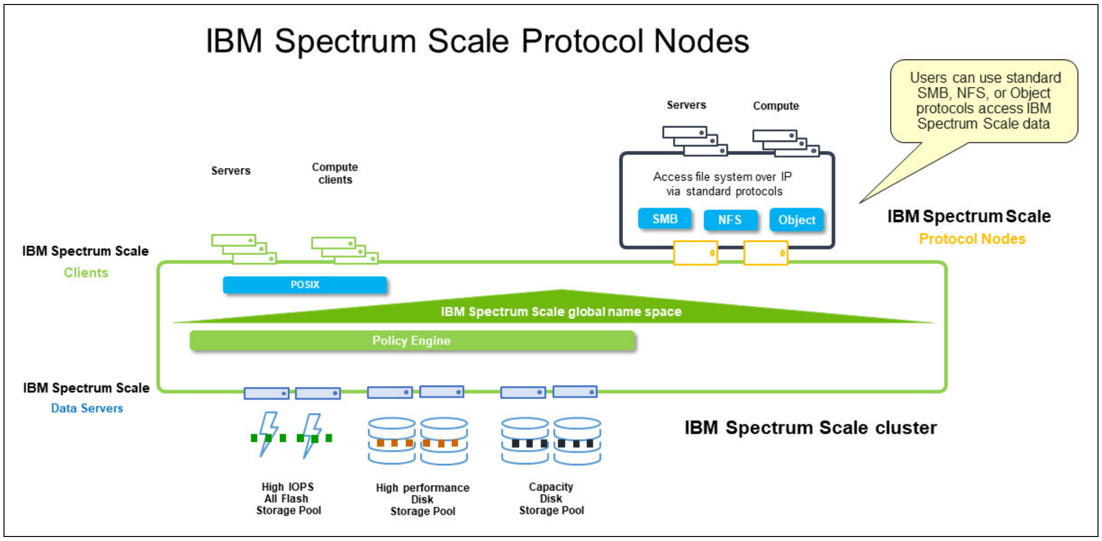

Figure 1-7 IBM Spectrum Scale protocol nodes

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

1.2.3 IBM ESS storage building block solution for IBM Spectrum Scale

IBM ESS with IBM Spectrum Scale is a pair of IBM Spectrum Scale Data Servers that are cross-configured together into a tested, integrated, HA, and reliable IBM Spectrum Scale storage building block. IBM ESS provides a fully tested, integrated, and supported IBM solution to deploy, manage, and maintain IBM Spectrum Scale data and storage by using a building block approach.

Two IBM ESS models are deployed to provide most of the IBM Spectrum Scale storage, as shown in Figure 1-8.

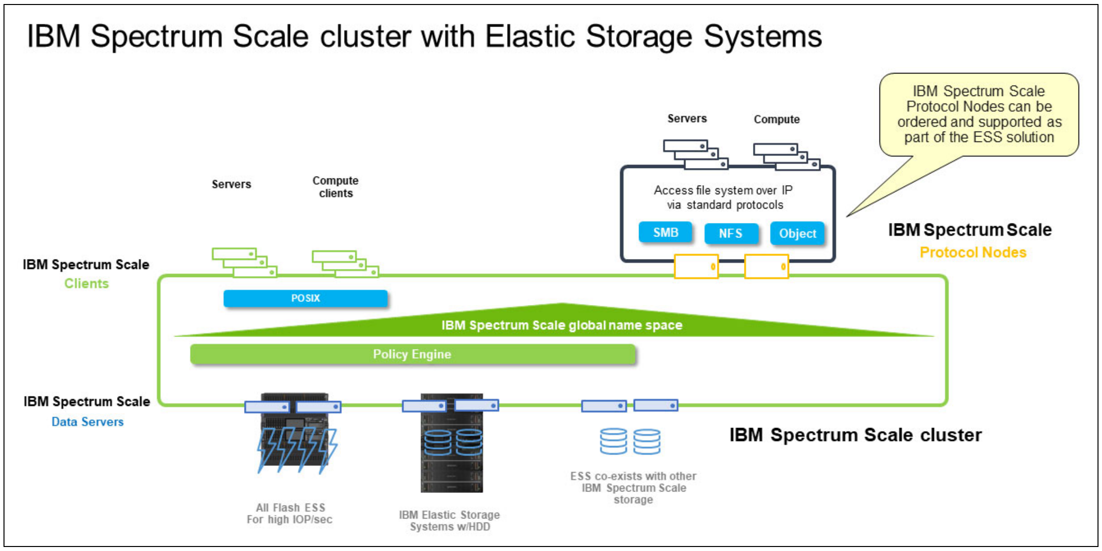

Figure 1-8 IBM Elastic Storage Systems are IBM Spectrum Scale storage building blocks

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

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

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

As part of an IBM ESS order, you also can order more IBM Power servers to use as IBM Spectrum Scale protocol nodes. These protocol nodes are managed by the IBM ESS solution software stack and the IBM ESS GUI. They also provide a complete, integrated IBM solution for an IBM Spectrum Scale cluster that includes IBM ESS and IBM Spectrum Scale protocol nodes.

For more information about planning for the usage of protocol nodes, see 2.2.8, “Protocol nodes” on page 32.

1.3 IBM Elastic Storage System solution value

The following sections describe why the IBM ESS brings added value advantages for deploying IBM Spectrum Scale storage.

Today's AI and big data applications require large, high-performance, manageable, and flexible storage growth.

The IBM ESS is designed to provide an integrated (see Figure 1-9) and tested IBM Spectrum Scale storage building block that provides this storage growth. Incorporating IBM ESS storage servers flexibly adds to the overall IBM Spectrum Scale capacity, bandwidth, and performance, all within a single global namespace.

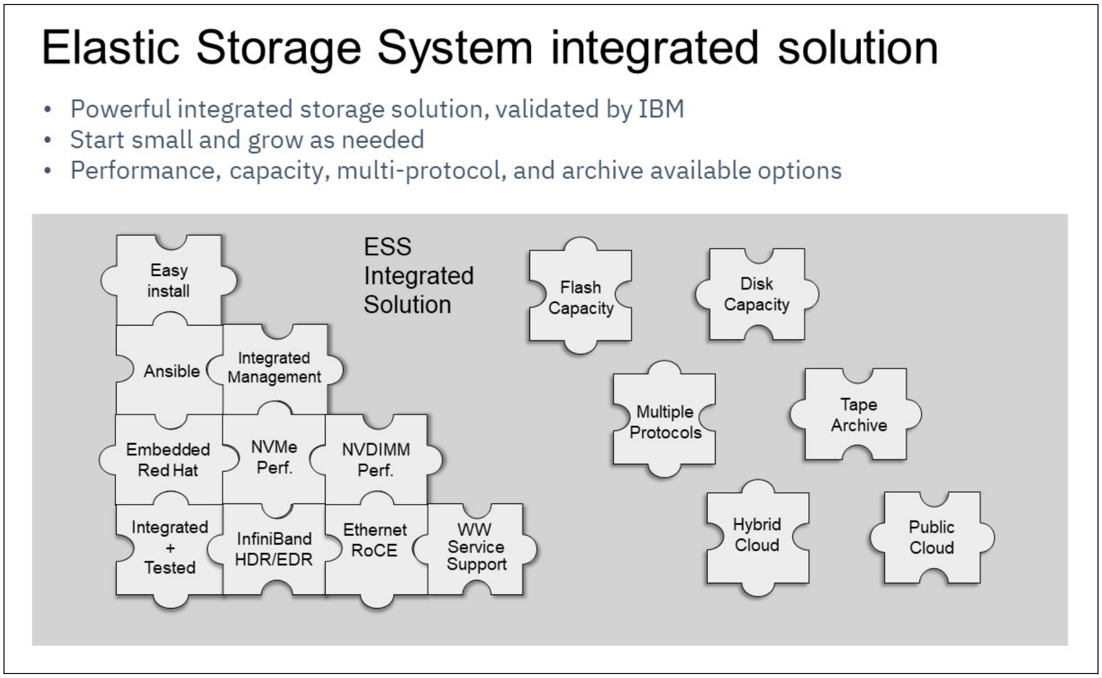

Figure 1-9 IBM Elastic Storage System integrated solution

As shown in Figure 1-9, IBM ESS is an integrated IBM Spectrum Scale solution that significantly accelerates the time-to-value of deploying IBM Spectrum Scale environments. IBM ESS provides the following benefits:

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

For more information about the IBM ESS, see the following web pages:

- IBM Elastic Storage System
- IBM Elastic Storage Server documentation

1.3.1 Integrated and reliable storage building block

Every part of the IBM ESS solution is engineered to provide high performance and enterprise data availability and reliability, from entry system to petabyte scale.

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

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

IBM ESS 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 ESS offers scalability from terabytes to hundreds of petabytes. IBM ESS supports high-speed data networks, including 200 GbE, and Extended Data Rate (EDR) or High Dynamic Range (HDR) InfiniBand. Alternatively or in addition, IBM ESS supports slower speed network access, including 10, 25, and 40 GbE.

1.3.2 Predictable and consistent high performance

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

A high-performance parallel file system environment, such as IBM Spectrum Scale, can run only as fast as its slowest component. IBM ESS that uses IBM Spectrum Scale RAID is designed 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 ESS assures that the larger IBM Spectrum Scale parallel file system always runs at optimum efficiency and consistently delivers expected high performance, even if storage hardware or media failures occur.

1.3.3 Simplified IBM ESS installation

IBM ESS technology deployment methods use Red Hat Ansible driven container orchestration, which results in faster installation, upgrades, and less effort and skills that are required (see Figure 1-10).

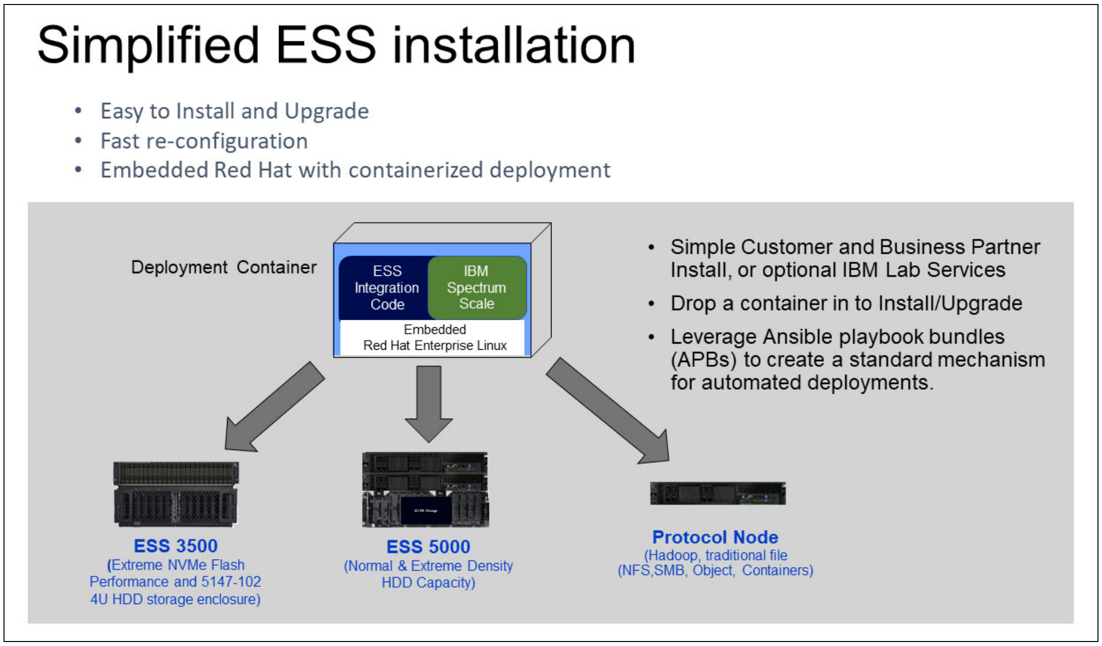

Figure 1-10 Simplified IBM ESS installation

As shown in Figure 1-10, IBM ESS provides a fully tested and integrated solution that is composed of Red Hat Enterprise Linux (RHEL), IBM Spectrum Scale, and IBM ESS integration code. Installation modules are driven by Ansible-driven container orchestration and Ansible playbook bundles (APBs), which create a standard, fast, and manageable mechanism of automated, faster deployment of IBM ESS and IBM Spectrum Scale storage.

1.3.4 Supported as a solution

The IBM ESS is supported as a solution. Clients who call IBM Service and Support for assistance on IBM ESS and IBM Spectrum Scale RAID receive support from a dedicated IBM team that is cross-trained on all IBM ESS components (hardware and software) for the IBM ESS solution. This specialized IBM ESS Service and Support team can debug and advise on all aspects of IBM ESS.

If more detailed support is needed by an IBM ESS component, IBM ESS 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 Spectrum Scale IBM ESS storage solution from IBM.

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

1.4 IBM Elastic Storage System models

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

The general physical arrangement for an IBM ESS is an IBM rack that features the following components:

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

An example of the various models of IBM ESS over time is shown in Figure 1-11.

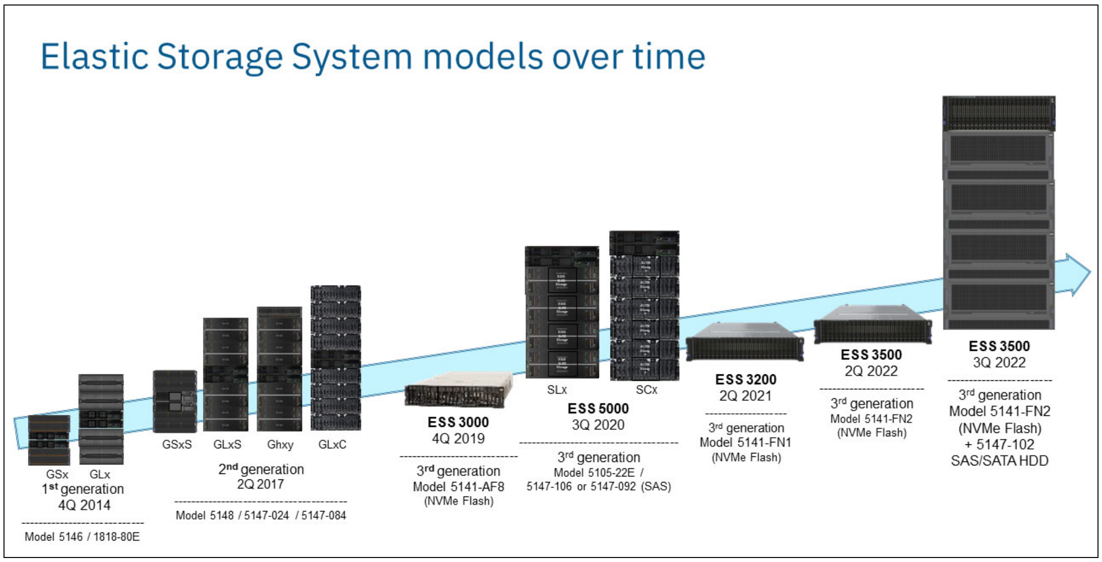

Figure 1-11 IBM Elastic Storage Server generations 1 and 2 and IBM Elastic Storage System (generation 3) models

2

IBM Elastic Storage System architecture

This chapter describes the architecture, hardware, and software of the IBM Elastic Storage System (IBM ESS), including the following components:

- IBM ESS Management Server (IBM EMS)
- I/O/data servers
- Storage enclosures

This chapter also describes the different building block models and software components and features, such as:

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

Note: IBM ESS is an integrated packaged solution. The user should *not* install different kernel levels or drivers into the IBM ESS system. The user also should *not* run any non IBM ESS (client) application or non IBM ESS workloads on the I/O nodes or IBM EMS.

This chapter includes the following topics:

2.1 Architecture overview

In this section, we describe the overall hardware and software architecture of the IBM ESS solution.

An IBM ESS is defined as a combine set of hardware and software that are tightly coupled together and tested as a single unit. The major components of an IBM ESS solution release include the following items:

- Server hardware
- Storage hardware
- IBM Spectrum Scale software
- 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 ESS solution communicate to each other over an IP network. IBM ESS 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 ESS 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 Ethernet (GbE) network switches for management and service networks
- High-speed network switches for data read and write over a high-speed data network

2.2 Hardware components

In this section, we describe the following hardware components of the IBM ESS:

- Solution models
- Server hardware (IBM ESS 5000 uses IBM POWER9 processor-based servers. IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000 use x86 servers.)
- Storage enclosures
- Network interface cards (NICs) and network switches
- Rack

IBM ESS is a rack-mounted IBM Spectrum Scale storage solution. The initial IBM ESS storage building block in an IBM Spectrum Scale cluster includes the following components:

- One IBM EMS (one is required for every IBM Spectrum Scale cluster).
- Two I/O data servers.
- The IBM ESS 3500 supports up to four storage enclosures. The IBM ESS 5000 supports up to seven or nine storage enclosures (depending on the IBM ESS 5000 model).

Third-generation IBM ESS 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 ESS 3500, which is the latest third-generation IBM ESS model, was announced and delivered in 2Q2022. The IBM ESS 3500 design is based on the IBM ESS 3200, but has a faster x86 based processor and various design improvements that improve serviceability. The IBM ESS 3500 was expanded in 3Q2022 to include support for up to four storage enclosures.

The IBM ESS 5000, which is also part of the third-generation IBM ESS models, was announced and delivered in 3Q2020. IBM ESS 5000 is a fully integrated IBM Spectrum Scale HDD storage building block with 1 - 9 storage enclosures (depending on the model) and a pair of POWER9 processor-based IBM Spectrum Scale Data servers.

The IBM ESS 3200, which also is a third-generation IBM ESS model, was announced and delivered in 2Q2021. IBM ESS 3200 is a fully integrated 2U storage building block that has 12 or 24 NVMe drives in the front of the 2U24 enclosure. It features a faster pair of PCI Gen4-based x86 server canisters in the back of the 2U enclosure.

The IBM ESS 3000, which was the first of the third-generation IBM ESS models, was announced and delivered in 4Q2019. IBM ESS 3000 is a fully integrated 2U storage building block that has 12 or 24 NVMe drives in the front of the 2U24 enclosure, and a pair of x86 server canisters in the back of the 2U enclosure.

First- and second-generation IBM ESS storage building blocks contain a pair of IBM POWER8® processor-based IBM Spectrum Scale Data servers.

2.2.1 IBM Elastic Storage System generations

This section describes the different generations and various IBM ESS models.

All IBM ESS generations can coexist in an IBM Spectrum Scale cluster.

Third-generation IBM ESS models

A third generation of IBM ESS was announced in November 2019 with the IBM ESS 3000 NVMe flash storage system. This announcement was followed by the announcement of the IBM ESS 5000 HDD offering in July 2020, the IBM ESS 3200 offering in May 2021, and the IBM ESS 3500 offering in May 2022.

The third-generation IBM ESS models include the following IBM ESS unique IBM machine types:

- IBM ESS 3500: IBM machine type 5141-FN2 for the 2U24 form factor IBM ESS with embedded NVMe flash storage, and optionally 1 - 4 of the 4U105 form factor IBM ESS Storage Enclosures model 5147-102
- IBM ESS 5000:
  - IBM machine type 5105-22E for the IBM ESS IBM POWER9 processor-based data servers, management server, and protocol nodes
  - IBM machine type 5147-092 or 5147-106 for the IBM ESS IBM Storage Enclosures

- IBM ESS 3200: IBM machine type 5141-FN1 for the 2U24 form factor IBM ESS with embedded NVMe flash storage. IBM announced the hardware withdrawal for IBM ESS 3200 on 23 August 2022, which is effective 31 March 31 2023. On or after the effective date of withdrawal, you can no longer order this product directly from IBM. The IBM ESS 3500 is the replacement.
- IBM ESS 3000: IBM machine type 5148-AF8 for the 2U24 form factor IBM ESS with embedded NVMe flash storage. New orders for IBM ESS 3000 were withdrawn from marketing effective 15 July 2022. The IBM ESS 3200 is the replacement.

These machine types uniquely identify these hardware machines as part of an IBM ESS solution.

Second-generation IBM ESS models

The second generation of IBM Elastic Storage Server (IBM ESS) was announced in April 2017.

Second-generation IBM ESS models do not include a solution IBM machine type. The IBM ESS server and storage hardware components feature the following IBM ESS unique machine types:

- IBM machine type 5148 for IBM ESS POWER8 processor-based servers
- IBM machine type 5147 for IBM ESS IBM Storage Enclosures

These machine types uniquely identify these hardware components as part of an IBM ESS solution. The models of second-generation IBM ESS are available in the following categories:

- Solid-state drive (SSD) flash storage IBM ESS models
- HDD IBM ESS models
- Hybrid IBM ESS models that consist of enclosures of SSD flash drives and HDDs

New orders for second-generation IBM ESS models can no longer be miscellaneous equipment specification (MES) upgraded to add more storage enclosures. They were announced to be withdrawn from marketing effective 31 December 2021.

> **Note:** The second-generation IBM ESS models are still in full service and supported by IBM. At the time of writing, no end of service (EoS) date has been announced yet.

For more information about second-generation IBM ESS models, see Appendix A, “IBM Elastic Storage System models” on page 65.

First-generation IBM ESS

The first generation of IBM ESS was announced in October 2014. It includes a solution IBM machine type 5146. GL and GS models are at End of Support.

2.2.2 IBM Elastic Storage System 5000

IBM ESS 5000 is the IBM HDD-based IBM Spectrum Scale storage platform. This storage platform provides high capacity and high-performance IBM Spectrum Scale storage by using HDD storage drives.

IBM ESS 5000 can include 6 TB, 10 TB, 14 TB, 16 TB, and 18 TB HDD drives depending on the model. At the time of writing, the IBM ESS 5000 SL model supports a superset of the drive sizes in the IBM ESS 5000 SC.

For more information about the full collection of manuals and documentation for IBM ESS 5000, see IBM Elastic Storage System documentation.

IBM ESS 5000 is based on IBM Spectrum Scale and IBM Spectrum Scale RAID, as are all IBM ESS models. IBM ESS 5000 includes the following features:

- IBM POWER9 processor-based data servers, which provide the latest advances in IBM POWER® high memory and memory bandwidth, PCI Gen4-based internal bus transfer speeds for higher internal transfer rates, and support of faster NICs.
- Containerized Red Hat Ansible playbooks that provide significantly improved ease of use and orchestration of complex IBM ESS administration tasks, such as cluster configuration, file system creation, and code update.
- Higher density and better HDD performance per rack than previous HDD-based IBM ESS models.
- IBM ESS 5000 supports nondisruptive capacity upgrade. You can add 5147-092 to an SLx model or a 5147-106 to a SCx model. Therefore, an installed SC1 to SC8 model, or SL1 to SL6 model, can be upgraded to any model configuration up to SC9 or SL7, without causing downtime or an interruption in service. (the additional HDD drives must be of the same size as the previous drives).

Figure 2-1 shows the IBM ESS 5000 SL models.

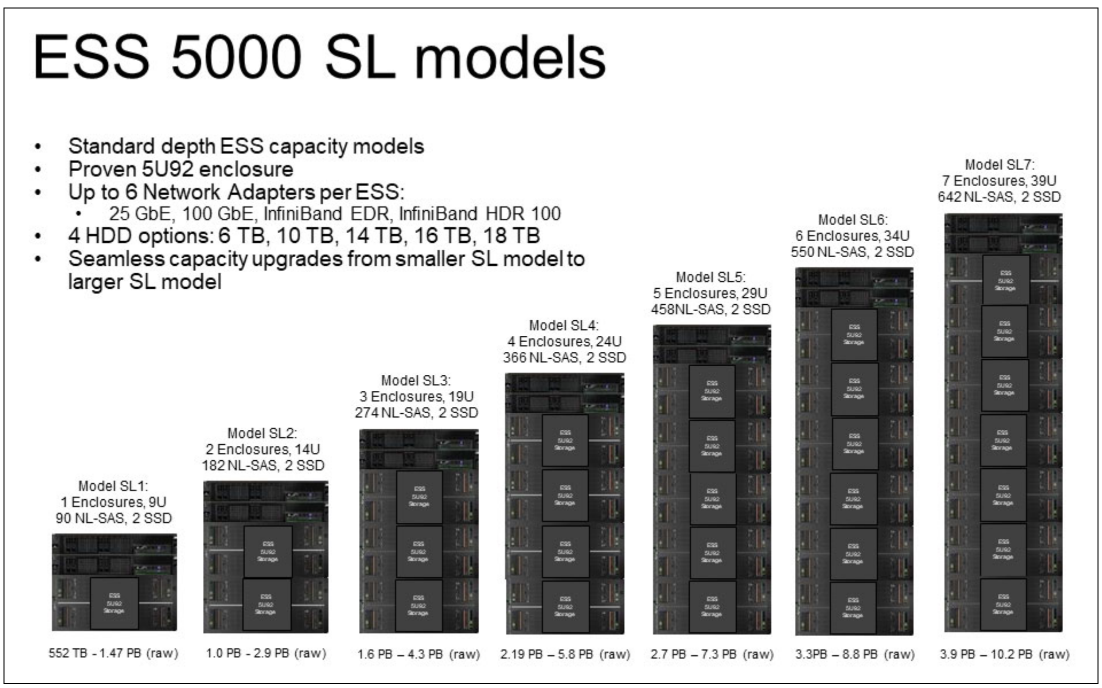

Figure 2-1 IBM ESS 5000 SL models

Figure 2-2 shows the IBM ESS 5000 SC models.

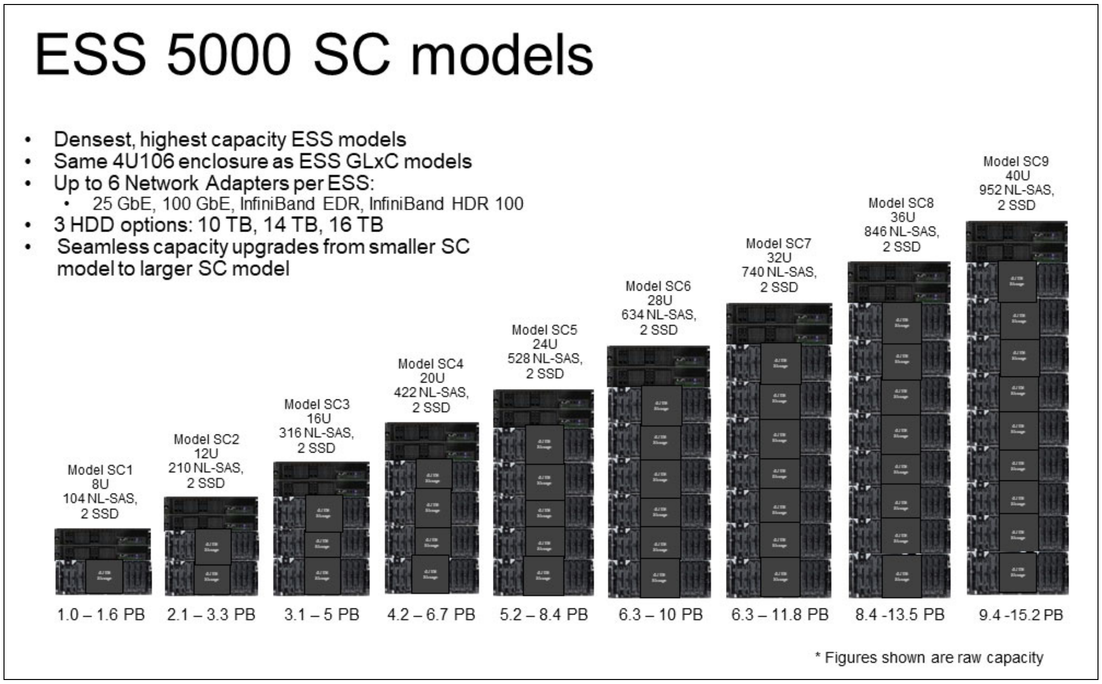

Figure 2-2 IBM ESS 5000 SC models

2.2.3 IBM Elastic Storage System 3500

IBM ESS 3500 is a hybrid model offering high-performance local NVMe storage and capacity storage through SAS or SATA connected storage enclosures.

IBM ESS 3500 can include 12 or 24 drives of 3.84, 7.68, 15.36, or 30.74 TB NVMe drives for user data. Alternatively, four NVMe drives are used exclusively for log tip use. For more information, see IBM Elastic Storage System 3500.

IBM ESS 3500 can include 1 - 4 storage shelves of SAS or SATA drives. The first storage shelves can contain 52 or 102 drives. Extra shelves contain 102 drives. The drives are 10, 14, 18, or 20 TB. For more information, see IBM Elastic Storage System 5147-102 Storage Enclosure.

IBM ESS 3500 is based on IBM Spectrum Scale and IBM Spectrum Scale RAID.

IBM ESS 3500 supports non-disruptive capacity upgrade. You can add 5147-102 storage shelves. IBM ESS 3500 supports up to four storage shelves, but all HDDs must be the same size.

Figure 2-3 on page 21 shows the IBM ESS 3500 models.

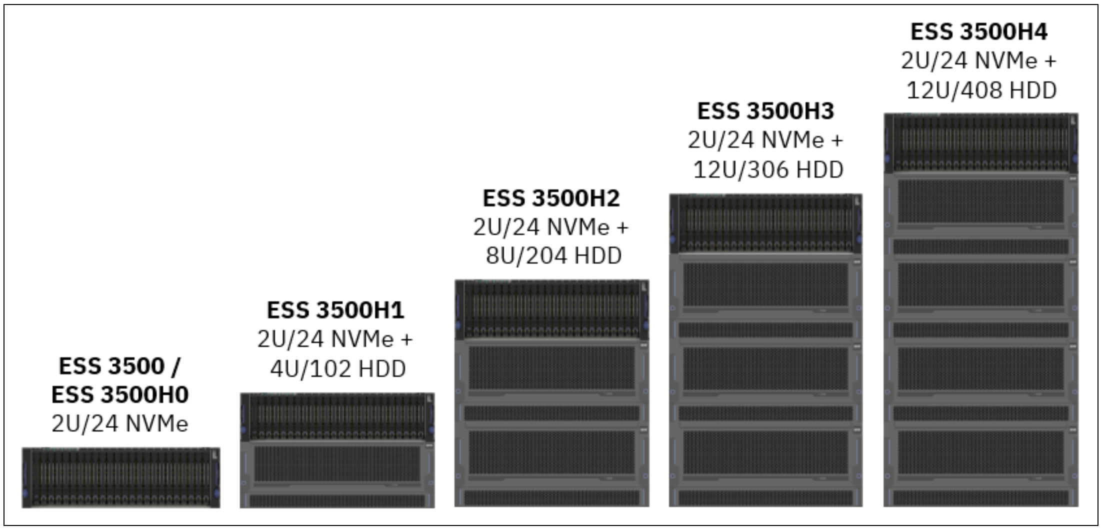

Figure 2-3 IBM ESS 3500 models

2.2.4 Positioning various IBM Elastic Storage System models

The different models of IBM ESS storage building blocks provide varying levels of performance and capacity according to the needs of the IBM Spectrum Scale storage pool in which they are deployed. The various models of third-generation IBM ESS can provide NVMe or SAS or SATA storage in selected disk sizes.

Figure 2-4 shows a general positioning of the available models of IBM ESS.

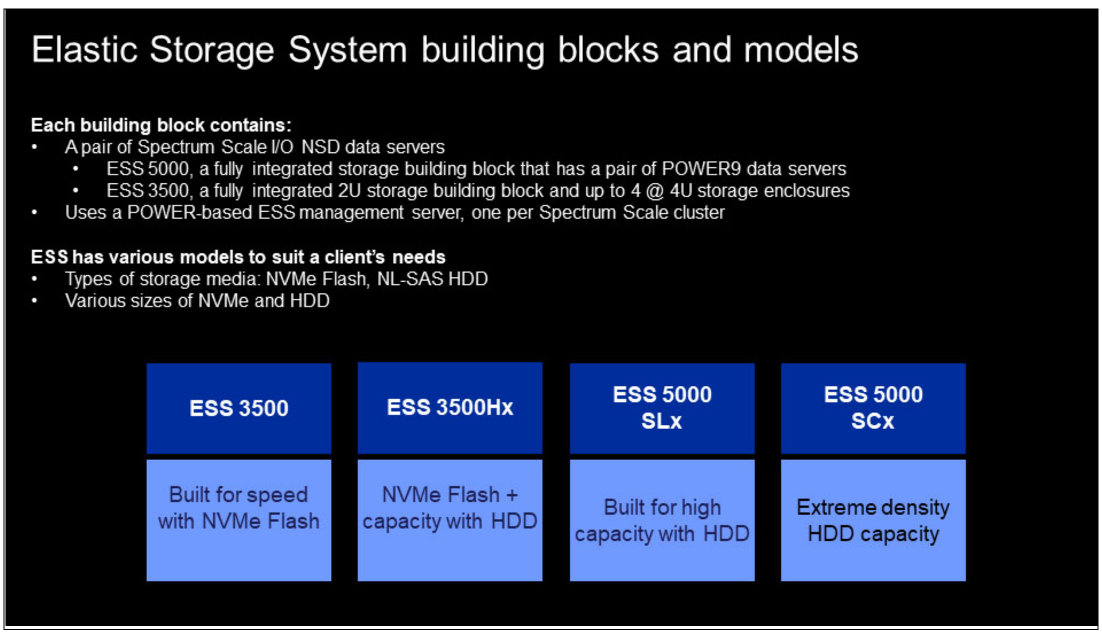

Figure 2-4 General positioning of available IBM ESS models

IBM ESS models are available in different categories. The “x” in the models (see Figure 2-5) denotes the number of storage enclosures in that model.

Consider the following points:

- IBM ESS 3500 models are built on NVMe flash storage for high performance, and optionally high capacity 4U102 HDD storage enclosures. Each IBM ESS 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.
- The IBM ESS 5000 SL model is built on high-capacity 5U92 HDD storage enclosures.
- The IBM ESS 5000 SC model is built on high-capacity 4U106 HDD storage enclosures.
- Supporting IBM ESS components include the following ones:
  - IBM EMS is required (see 2.2.5, “IBM Elastic Storage System Management Server” on page 23).
  - IBM ESS protocol node (see 2.2.8, “Protocol nodes” on page 32).

The available IBM ESS models are shown in Figure 2-5.

Elastic Storage System models at a glance

	ESS 3500	ESS 3500Hx Storage Enclosure	ESS 5000 SLx	ESS 5000 SCx
Models	2U24 With 12 or 24 drives	0 to 4 @ 4U102 Storage enclosure	SL1 SL2 SL3 SL4 SL5 SL6 SL7	SC1 SC2 SC3 SC4 SC5 SC6 SC7 SC8 SC9
Drive sizes	NVMe: 3.84 TB 7.68 TB 15.36 TB 30.74 TB	HDD: 10 TB 14 TB 18 TB 20 TB	HDD: 6 TB 10 TB 14 TB 16 TB 18 TB	HDD: 10 TB 14 TB 16 TB

Figure 2-5 IBM Elastic Storage System models

IBM ESS usable capacity is a percentage of the raw capacity. The usable capacity varies, depending on the IBM Spectrum Scale RAID parity that is selected for use. Generally, approximately 73% of the raw capacity is the usable capacity when the default 8+2P parity is used, which is the most commonly selected parity for most scenarios.

IBM or IBM Business Partners can provide the exact usage capacities of the various models on request by using the IBM File Object Solution Design Engine (FOS DE) or the Storage Modeler (StorM) to calculate the exact usable capacity. See your IBM or IBM Business Partner representative for assistance in determining the usable capacity that your IBM ESS model and implementation provides.

For more information about the various IBM ESS models specifications, see Appendix A, “IBM Elastic Storage System models” on page 65.

2.2.5 IBM Elastic Storage System Management Server

Released with the IBM ESS 3200 was IBM EMS. This management system can be used with the IBM ESS 3000, and it is required on IBM ESS 3500, IBM ESS 3200, and IBM ESS 5000. IBM EMS (IBM machine type 5105-22E) is based on the IBM POWER9 processor.

The POWER9 processor-based IBM EMS server features the following specifications:

- 1x DD2.3 20 Small Cores, 190W/225W, 2.5 GHz or 2.9 GHz
- 128 GB default memory, no NVDIMMs (nonvolatile dual inline memory module)
- No SAS HBAs
- The same network interface options as available on all third-generation IBM ESS models

Every IBM Spectrum Scale cluster with IBM ESS requires a minimum of one IBM EMS. This IBM POWER9 processor-based server is used to manage and deploy the IBM ESS I/O data server nodes. This server performs the following tasks:

- Runs an IBM ESS GUI server for managing IBM ESS in the cluster.
- Can optionally provide other IBM Spectrum Scale functions, such as a quorum node.
- Runs the newer management and ease of installation enhancements of third-generation IBM ESS by using Red Hat Ansible orchestration.

The initial IBM ESS storage building block in an IBM Spectrum Scale cluster requires an IBM EMS. The purpose of the IBM EMS is a central control point to manage multiple IBM ESS storage building blocks in the IBM Spectrum Scale cluster.

The third-generation IBM ESS 3500, IBM ESS 3200, and IBM ESS 5000 models require a POWER9 processor-based IBM EMS (IBM machine type 5105-22E) running RHEL Little Endian. This Little Endian IBM EMS can manage IBM ESS generation three and two.

The second-generation IBM ESS models use an IBM EMS (IBM machine type 5148-21L) running RHEL Little Endian. This Little Endian IBM EMS can manage second-generation IBM ESS and IBM ESS 3000, but cannot manage first-generation IBM ESS or third-generation IBM ESS 3500, IBM ESS 3200, or IBM ESS 5000.

If you have a mix of generations of IBM ESS in the same IBM Spectrum Scale cluster, you must have a suitable number of IBM EMSs to manage the IBM ESS generations.

IBM ESS GUI and Call Home

The IBM EMS runs the IBM ESS GUI and is the central Call Home control point for the IBM ESS hardware. IBM ESS Call Home monitors server and storage hardware. If a hardware failure event occurs, the IBM EMS can be configured to initiate a Call Home service action to IBM Service and Support.

2.2.6 Data Server hardware

The IBM ESS solution integrates a pair of IBM Spectrum Scale Data Servers that are cross-connected to all storage in the IBM ESS, which provides an IBM Spectrum Scale storage building block. Second-generation IBM ESS models use IBM Power servers for the data servers. The IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000 use an x86-based pair of servers for data servers. The IBM ESS 5000 uses a POWER9 processor-based pair of data servers.

The server hardware, firmware, and drivers are all tested, integrated, and supported as a solution by IBM and as part of the overall IBM ESS solution.

In the following sections, we describe the server specifics for third-generation and second-generation IBM ESS models.

Second-generation IBM ESS Data Servers

All second-generation IBM ESS models use a pair of POWER8 processor-based servers for the integrated IBM Spectrum Scale Data Servers. The second-generation IBM ESS models use POWER8 processor-based (machine type 5148-22L) servers as IBM Spectrum Scale Data servers. These 5148-22L servers all run the RHEL Little Endian operating system.

Third-generation IBM ESS Data Servers

With IBM ESS 5000, IBM introduced a new POWER9 processor-based IBM ESS server: IBM Machine Type 5105-22E. This server is used in the IBM ESS 5000 only.

The 5105-22E is configured in three separate ways to provide of one of the following server functions in the IBM ESS 5000 system:

- IBM EMS
- Protocol node
- IBM ESS Data Server

All these servers are based on the IBM POWER9 processor-based architecture and include the same machine type and model (MTM) (5105-22E). In the IBM ESS 5000 configuration, they are differentiated by feature codes that specify which of the three possible roles this 5105-22E is performing.

The 5105-22E POWER9 processor-based servers feature increased memory and backplane bandwidth. In particular, the usage of PCI Gen4 bus allows each of the ports on the dual-port NICs to run at full rated speed.

The IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000 2U24 form factor includes two x86 server canisters in the rear of the 2U24 storage enclosure, which provide a small footprint and a dense and integrated NVMe flash IBM Spectrum Scale storage building block. These two x86 server canisters run IBM Spectrum Scale RAID as a pair of IBM Spectrum Scale Data Servers, which are cross-connected to all storage in the 2U24 enclosure for high availability (HA).

IBM ESS 3500 Storage Enclosure 5147-102 is an enterprise-class, fully redundant SAS or SATA enclosure for IBM ESS 3500 5141-FN2. The enclosure supports up to one hundred and two 3.5-inch or 2.5-inch drives of any density and spindle speed, and up to four 5147-102 Storage Enclosures can be attached to IBM ESS 3500 (5141-FN2) to provide up to 8 PB of HDD capacity.

Note: The IBM ESS 3000 can be supported by a POWER9 processor-based 5105-22E IBM EMS or a POWER8 processor-based 5148-21L IBM EMS. This management server can be an IBM ESS 5148-21L that is used to manage other IBM ESS models in this same IBM Spectrum Scale cluster. If this cluster is a new IBM Spectrum Scale cluster, order a POWER9 processor-based 5105-22E IBM EMS to manage your third-generation IBM ESS IBM Spectrum Scale cluster.

Network interface cards

The IBM ESS solution features the following types of networks:

- Management and service networks that require basic, low-speed 1 GbE.
- All IBM ESS models support high-speed data networks (Extended Data Rate (EDR) InfiniBand, and 100 GbE, 40 GbE, 25 GbE, or 10 GbE) over which user data is read and written.
- Specific IBM ESS models support even higher speed data networks: IBM ESS 3500, IBM ESS 3200, and IBM ESS 5000 support High Dynamic Range (HDR) InfiniBand or 200 GbE.

The NICs on the IBM ESS Data Servers and the IBM EMS service these various network types.

Every IBM ESS node has a 1 GbE NIC for connection to the management and service networks.

Every IBM EMS, in addition to the 1 GbE NIC, must have one of each of the high-speed network NICs being used on the IBM ESS Data Servers. This configuration allows the IBM EMS to monitor, manage, and measure the activity on the high-speed data networks.

Note: Specific IBM ESS models support NICs that provide Virtual Protocol Interconnect (VPI). VPI allows individual ports of the dual-port network card to be configured as InfiniBand or 200 GbE. If you use the VPI feature to reconfigure your NIC, make sure to configure or add suitable NICs to your IBM EMS so that the IBM EMS can attach to that network.

On the IBM ESS Data Server nodes, a sufficient quantity of high-speed data network NICs must be configured to provide the wanted IBM ESS bandwidth.

Various high-speed data network NICs are available:

- 200 GbE
- 100 GbE
- 40 GbE
- 25 GbE
- 10 GbE

A VPI-enabled InfiniBand or 100 GbE network card is available on all third-generation IBM ESS models. This dual-port InfiniBand and Ethernet VPI network adapter is based on the NVIDIA ConnectX technology. Each of the two ports on the card can be independently configured for 100 GbE or InfiniBand.

The HDR InfiniBand or 100 GbE network card is available on IBM ESS 5000. This dual-port InfiniBand-HDR and Ethernet VPI network adapter is based on the NVIDIA ConnectX-7 card with VPI enabled. Each of the two ports on the card can be independently configured for 100 GbE or HDR-100.

The HDR InfiniBand or 200 GbE network card is available on IBM ESS 3500. This dual-port InfiniBand-HDR and Ethernet VPI network adapter is based on the NVIDIA ConnectX-6 card with VPI enabled. Each of the two ports on the card can be independently configured for 100 GbE or HDR-200.

Note: Consider the following points:

- IBM ESS 3500 and IBM ESS 5000 use a PCIe Generation 4 bus, which fully supports both ports of the 200 GbE or InfiniBand NICs to run at full rated speed.

  IBM and IBM Business Partners use the IBM File and Object Design Engine tool to estimate and configure an appropriate amount of NIC bandwidth.

  Experience shows that throughput on IBM ESS is often limited to the amount of network NIC bandwidth that is available. The best performance is achieved over InfiniBand.
- Suitable network cables and connectors must be specified and ordered for your specific network requirements. If you use the VPI feature to reconfigure your high-speed network ports, make sure that your IBM EMS also can access the newly configured network.
- RDMA over Converged Ethernet (RoCE) might be available on third-generation IBM ESS models for specific configurations, networks, and environments only. If you have this requirement, submit a Request for Price Quotation (RPQ) to IBM by contacting your IBM or IBM Business Partner representative.
- For more information about estimating and sizing your performance, or for specifying the suitable network cables and connectors, contact your IBM or IBM Business Partner representative.

You can intermix high-speed data network NIC types if enough data server slots are available. Each of the data servers in an IBM ESS must have same number and type of NICs.

Configure your IBM ESS to have the suitable type and number of NICs for your throughput requirements. It is supported and a common practice on IBM ESS Data Servers to bond multiple ports for increased availability and throughput.

2.2.7 Storage enclosures

A storage enclosure is a specialized storage drawer holds and powers flash or HDD storage, while providing the connections to allow them to communicate to one or more separate firmware service modules that run the entire enclosure.

First-generation IBM Elastic Storage Server

The first generation of IBM ESS was announced in October 2014. It was available in 2U24 storage enclosure (GS models) with SSD or HDD storage, or high capacity GL storage models that used IBM 1818-80E 4U60 HDD Storage Enclosures. New orders and storage enclosure upgrades are withdrawn from marketing and no longer available, effective 19 January 2018.

The announced EoS date for the first-generation IBM ESS models is 31 December 2021. For more information, contact your IBM or IBM Business Partner representative.

For more information about the first-generation IBM ESS Storage Enclosures, see Appendix A, “IBM Elastic Storage System models” on page 65.

Second-generation IBM Elastic Storage Server

The second generation of IBM ESS was announced in April 2017 and was available for order until the last order date of 20 December 2020.

The following components have been withdrawn from marketing:

- IBM 5147-084 HDD Storage Enclosures to add to a second-generation IBM ESS GLxS model until 30 June 2021.
- IBM 5147-024 SSD Storage Enclosures to add to a second-generation IBM ESS GSxS or GH model until 31 December 2021.
- IBM 5147-106 HDD Storage Enclosures to add to a second-generation IBM ESS GLxC model until 31 December 2021.

Note: The second-generation IBM ESS is still in full service and supported by IBM. No EoS date was announced at the time of this writing.

Third-generation IBM Elastic Storage Server

The following third-generation IBM ESS models were announced, and they use the associated storage enclosures:

- IBM ESS 3500 was announced and delivered in 2Q2022. It is a self-contained 2U24 form factor storage enclosure with NVMe flash storage in the front, and PCI Gen4 x86-based IBM Spectrum Scale Data Servers in the back. The IBM ESS 3500 was expanded in 3Q2022 to include support for up to four storage enclosures (IBM machine type 5147-102).
- IBM ESS 3200 was announced and delivered in 2Q2021. It is a self-contained 2U24 form factor storage enclosure with NVMe flash storage in the front, and PCI Gen4 x86-based IBM Spectrum Scale Data Servers in the back.
- IBM ESS 5000 was announced and delivered in 3Q2020. It uses 5U92 (IBM machine type 5147-092) or 4U106 (IBM machine type 5147-106) HDD storage enclosures.
- IBM ESS 3000 was announced and delivered in 4Q2019. It is a self-contained 2U24 form factor storage enclosure with NVMe flash storage in the front, and x86-based IBM Spectrum Scale Data Servers in the back.

In the following sections, we describe storage enclosures that you are likely to encounter in working with second-or third-generation IBM ESS. Machine type 5147 is used for the IBM ESS IBM Storage Enclosures. Machine type 5147 uniquely identifies these storage enclosure hardware components as part of an IBM ESS.

2U24 (5147-024)

The 5147-024 2U Storage Enclosure is used for SSDs in the second-generation IBM ESS GSxS and second-generation GH hybrid models. The following sizes of SSDs are available:

- 3.84 TB
- 15.36 TB

The 5147-024 in the IBM ESS is available in a fully populated 24-drive configuration only.

The IBM 5147-024 Storage Enclosure is shown in Figure 2-6.

Figure 2-6 IBM 5147-024 Storage Enclosure

5U84 (5147-084)

The 5147-084 5U Storage Enclosure is available to hold HDDs in the second-generation IBM ESS GLxS models and the GH hybrid models. The following sizes of HDDs are or were available:

- 4 TB
- 8 TB
- 10 TB
- 14 TB

The IBM 5147-084 Storage Enclosure is shown in Figure 2-7.

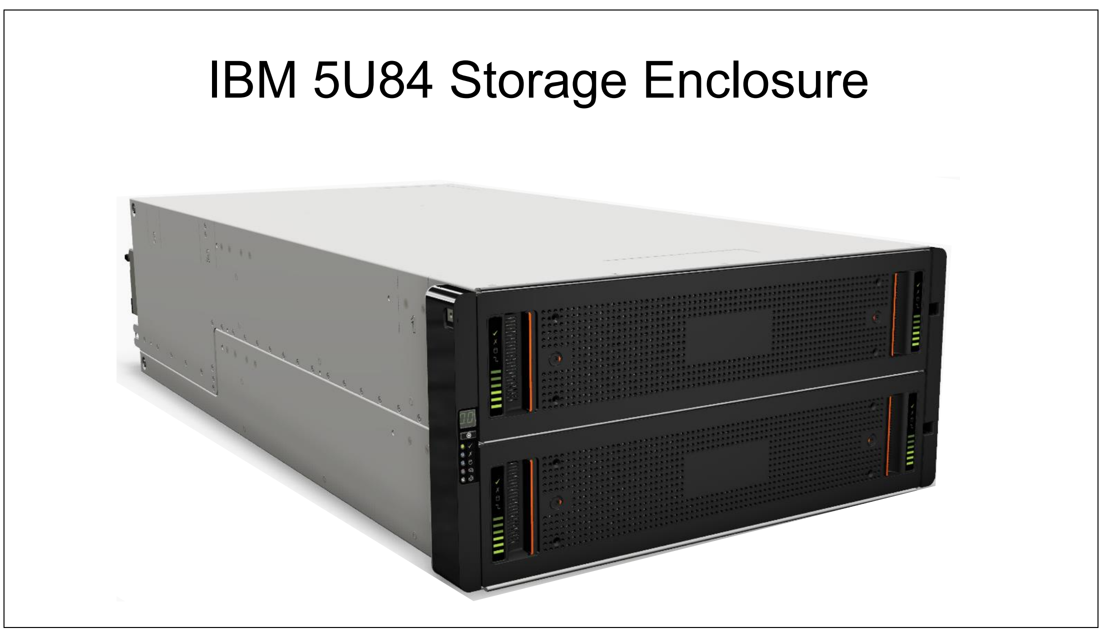

Figure 2-7 IBM 5147-084 Storage Enclosure

The 5147-084 in the second-generation IBM ESS is available in a fully populated configuration only.

The first 5147-084 Storage Enclosure in an IBM ESS also contains two SSDs that hold internal IBM Spectrum Scale RAID metadata. The capacity of these two SSDs is not visible or available as user capacity for the IBM Spectrum Scale or IBM ESS file system.

4U106 (5147-106)

The 5147-106 4U Storage Enclosure is an extreme density storage enclosure that holds HDDs and available in second-generation IBM ESS GLxC and third-generation IBM ESS 5000 SC models. The following sizes of HDDs are or were available:

- 10 TB
- 14 TB
- 16 TB (IBM ESS 5000 only)

The 5147-106 in the IBM ESS is available in a fully populated configuration only.

The first 5147-106 in an IBM ESS also contains two SSDs that hold internal IBM Spectrum Scale RAID metadata. The capacity of these two SSDs is not visible or available as user capacity for the IBM Spectrum Scale IBM ESS file system.

The IBM 5147-106 Storage Enclosure is shown in Figure 2-8.

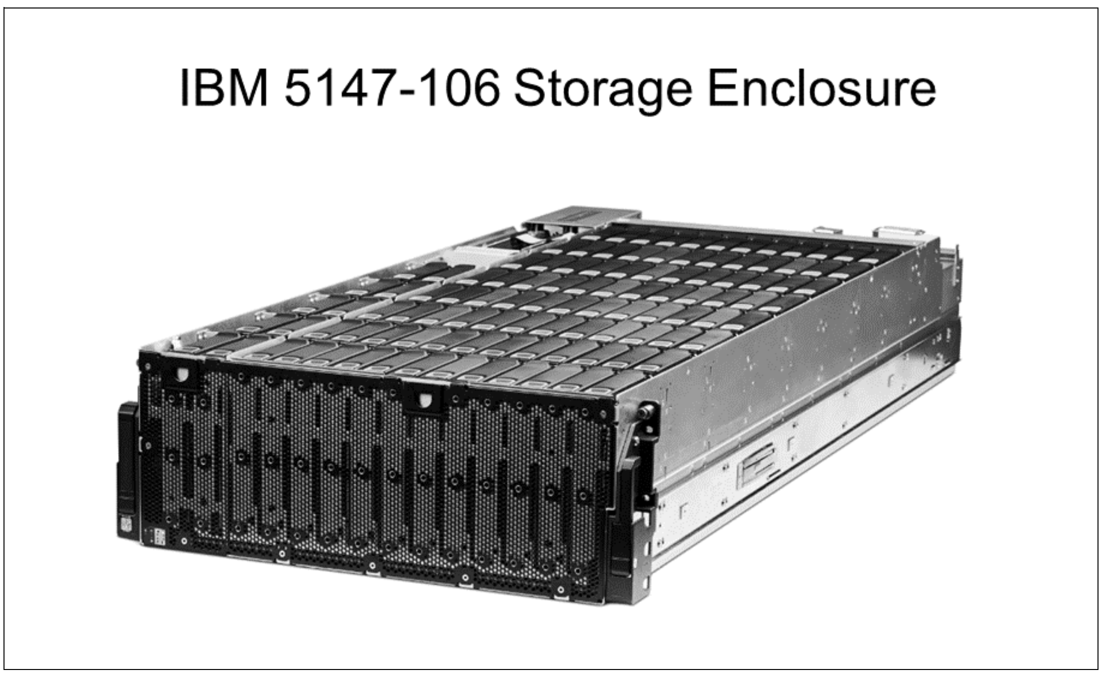

Figure 2-8 IBM 5147-106 Storage Enclosure

The IBM 5147-106 Storage Enclosure is an unusually deep enclosure. Physical planning for this enclosure must be performed. Also, the racking, aisle widths, and rear clearances all must be reviewed and approved by the IBM Installation Planning representative.

Note: Because of the depth of this 5147-106 enclosure, the IBM ESS models that use 5147-106 must be installed in the IBM 7965-S42 rack. These enclosures (and associated IBM ESS SC models) cannot be installed in the IBM 7014-T42 rack.

Second-generation IBM ESS Hybrid models storage enclosures

Second-generation IBM ESS Hybrid models combined 5147-024 Storage Enclosures with flash SSD storage and 5147-084 enclosures with spinning HDDs and in a single IBM ESS building block. These IBM ESS Hybrid models provided a combination of drawers of 2U24 SSD and 5U84 HDDs, which combines the speed and low latency of flash SSD storage with the density and affordability of spinning HDD.

This configuration provides a flexible, smaller footprint than using a separate flash SSD storage building block and disk HDD storage building block. Where the IBM Spectrum Scale IBM ESS configuration indicates the usage of the IBM ESS GH hybrid models, these IBM ESS hybrid storage building blocks were used to give better storage density and a smaller footprint at a lower cost.

Third-generation IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000 Storage Enclosures

The third-generation IBM ESS 3500, IBM ESS 3200, or IBM ESS 3000 with NVMe is available with a 2U24 storage enclosure, with slots for 24 NVMe storage drives in the front of the enclosure. The IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000 use different physical storage enclosures (they are *not* the same). Conversion from one model to another is *not* possible. These different 3x00 storage enclosure models are similar in that they provide front panel hot-swappable drive access. Rear access is provided to the two power supplies and two x86-based IBM Spectrum Scale Data Servers.

Figure 2-9 shows a front view of the IBM ESS 3500 and IBM ESS 3000 Storage Enclosures.

Figure 2-9 Sample IBM ESS 2U24 front panel accessible NVMe Storage Enclosure

> **Note:** IBM ESS 3500 and IBM ESS 3000 use enterprise-class NVMe drives. The IBM ESS 3200 supports either NVMe or IBM FlashCore®, depending on size.

IBM ESS 3x00s use a mirrored set of internal boot disks for the x86-based IBM Spectrum Scale Network Shared Disk (NSD) Data Servers.

5U92 (5147-092)

The 5147-092 5U Storage Enclosure is available to hold HDDs in the IBM ESS 5000 SL models. The following sizes of HDDs are available:

- 6 TB
- 10 TB
- 14 TB
- 16 TB

The IBM 5147-092 Storage Enclosure is shown in Figure 2-10.

Figure 2-10 IBM 5147-092 Storage Enclosure

The 5147-092 in the IBM ESS 5000 SL model is available in a fully populated configuration only. The first 5147-092 in an IBM ESS 5000 also contains two SSDs that hold internal IBM Spectrum Scale RAID metadata. The capacity of these two SSDs is not visible or available as user capacity for the IBM Spectrum Scale or IBM ESS file system.

4U102 (5147-102)

The 5147-102 4U Storage Enclosure is available to hold HDDs in the IBM ESS 3500 models. The following HDD sizes are available:

- 10 TB
- 14 TB
- 18 TB
- 20 TB

The IBM 5147-102 Storage Enclosure is shown in Figure 2-11.

Figure 2-11 IBM 5147-102 Storage Enclosure

IBM ESS 3500 Storage Enclosure 5147-102 is an enterprise-class, fully redundant SAS or SATA enclosure for IBM ESS 3500 5141-FN2. The enclosure supports up to one hundred and two 3.5-inch or 2.5-inch drives of any density and spindle speed. Up to four 5147-102 Storage Enclosures can be attached to IBM ESS 3500 (5141-FN2), which provides up to 8 PB of HDD capacity.

2.2.8 Protocol nodes

The IBM ESS IBM Spectrum Scale protocol node feature code makes it possible to order POWER9 processor-based 5105-22E server nodes with a specific hardware configuration, which are tested and tuned by IBM to provide IBM Spectrum Scale protocol node (also known as Cluster Export Services (CES) services. The MTM for protocol nodes that are ordered with an IBM ESS is IBM machine type 5105-22E.

The default protocol node hardware features the following configuration:

- 5105-22E POWER9 processor-based model with two 10 core POWER9 processors
- Runs RHEL Little Endian
- 192 GB or greater memory in default configuration
- 1 GbE 4-port network adapter that is required for management network
- Up to seven network adapters

All NICs that are available for other 5105-22E Data Servers are for 5105-22E servers that are ordered as protocol nodes.

Note: The 5105-22E protocol nodes are supported for management by the POWER9 processor-based 5105-22E IBM EMS node only.

5105-22E protocol nodes may coexist in an IBM Spectrum Scale cluster with previous generation IBM ESS 5148-22L protocol nodes and with other non-5148-22L IBM Power Little Endian servers that also are protocol nodes.

The IBM ESS 3500 can run embedded virtualized protocol nodes. The limitations for running CES require no more than 1024 Network File System (NFS) active connections or 512 Server Message Block (SMB) or Common Internet File System (CIFS) active connections. Object support is not supported on the embedded virtualized protocol nodes.

Note: IBM Spectrum Scale requires that in any one IBM Spectrum Scale cluster, all protocol nodes must be of the same architecture and of the same Endian. The following types of IBM Spectrum Scale protocol nodes architecture and Endian are supported:

- IBM Power servers running RHEL Little Endian, including the IBM 5105-22E POWER9 processor-based protocol node, and the previous generation POWER8 processor-based 5148-22L protocol node
- IBM Power servers running RHEL Big Endian (for example, IBM 8247-2xL servers)
- x86 servers running supported Linux operating systems

For more information, see the IBM Spectrum Scale Functional Support Matrices in the IBM Spectrum Scale FAQ at IBM Documentation.

If you are ordering IBM Power 5105-22E servers as protocol nodes, you might have other IBM Power RHEL Little Endian protocol nodes in the IBM Spectrum Scale cluster that are not IBM 5105-22E protocol nodes or previous generation 5148-22L protocol nodes. These other protocol nodes are installed separately and deployed, upgraded, and managed by the customer. The IBM EMS tool sets are not supported and cannot be used for the management of operating system, kernel, network manager, systemd, OpenFabrics Enterprise Distribution (OFED), or firmware on the non-5105-22E or non-5148-22L protocol nodes.

2.2.9 Networking switches

The IBM ESS solution requires IP or InfiniBand networking switches to interconnect the IBM ESS hardware components.

As a best practice, purchase and use NVIDIA networking switches. These network switches are tested by IBM to work with IBM ESS.

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

Types of networks that are used by IBM ESS

The IBM ESS features the following types of networks that are used to interconnect the IBM ESS components:

- Management and service networks, which must have separate 1 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 Spectrum Scale IBM ESS file systems

Regardless of the source and type of network switches between the IBM ESS, the IBM Spectrum Scale cluster and the users, the networking must be solid, robust, reliable, and provide consistent latency and response time. Non-blocking, high-speed switches are highly recommended.

Experience shows that IBM Spectrum Scale and IBM ESS can easily generate throughput rates that can overwhelm an over-subscribed network. IBM Spectrum Scale provides a network load generation testing tool (**nsdperf**), which is available to anyone as an open source repository on GitHub. This tool should be used to test network performance, network latency, network topology, and network readiness level before IBM ESS is installed.

For more information about the **nsdperf** tool, see this GitHub repository.

Management network switches

IBM ESS components require a 1 GbE Ethernet IP connection for management and service network traffic. As part of an IBM ESS order, IBM can supply the NVIDIA 1 GbE management network switch for IBM machine type 8831-S52 or 8831-S48.

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

Alternatively, the client can provide a 1 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 ESS components also require a high-speed data network IP or InfiniBand connection for the reading and writing of user data. As part of IBM ESS testing, IBM tests with NVIDIA data network switches.

As a best practice, order and use the NVIDIA data network switches where possible. These switches are tested and integrated with the IBM ESS solution.

NVIDIA high-speed network switch drivers and firmware is tested and delivered integrated within the IBM ESS software solution stack.

Alternatively, the client can provide their own 1 GbE management network and 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 also is responsible for maintaining and troubleshooting any networking issues.

2.2.10 Racking

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

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

The following IBM racks are available for IBM ESS:

- IBM 7014-T42 Enterprise Rack
- IBM 7965-S42 Enterprise Slim Rack

Both racks provide 42 EIA (42U) of usable space for IBM ESS components and network switches. A choice of various electrical PDUs is available to be specified for inclusion in the rack side areas. The difference between the racks is the width and depth.

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

The IBM 7014-T42 is the default for most models of IBM ESS, except for the IBM ESS GLxC models and IBM ESS 5000 SC models.

For the IBM ESS models that use the 5147-106 Storage Enclosure, use the IBM 7965-S42 rack, which is designed to accommodate the deep IBM 5147-106 Storage Enclosure. The IBM 7965-S42 rack is narrower and fits on a 24-inch floor tile. This rack might better fit clients with tight racking requirements and space. It can be ordered for other IBM ESS models where the client wants an alternative to the IBM 7014-T42 rack.

It is also possible to order IBM ESS without a rack. In this case, the IBM ESS components are assembled and tested by IBM Manufacturing. Then, the IBM ESS components (servers, storage, networking switches, and cables) are uninstalled from the manufacturing rack, labeled, and shipped.

If installing the IBM ESS 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 ESS 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 ESS components in a client rack on behalf of the client.

For the rackless IBM ESS order, the solution must be assembled and integrated into the rack at the client site. Extra IBM Technical Support Services (TSS) installation hours for installation might be required. For more information, see your IBM TSS representative.

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

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

2.3 Software components

This section describes the software components that are used on the IBM ESS solution.

2.3.1 Software solution stack overview

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

In the following sections, we briefly describe the following IBM ESS solution stack software components:

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

IBM preinstalls the complete integrated and tested IBM ESS solution stack on the IBM ESS Data Servers in IBM Manufacturing.

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

The IBM ESS solution stack components are periodically upgraded to newer release levels, tested as an integrated solution, and released as a new level of the IBM ESS 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 Spectrum Scale Software Version Recommendation Preventive Service Planning.

For more information about the release levels of the IBM ESS software solution and the levels of the software components for that IBM ESS release level, see question 2.2 in the IBM Spectrum Scale RAID FAQ for IBM ESS in IBM Documentation.

2.3.2 Operating system

The IBM ESS solution runs RHEL as the operating system on the IBM Spectrum Scale Data Servers.

Each IBM ESS 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 ESS solution stack. IBM periodically provides new IBM ESS 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.

Consider the following points:

- All generations of IBM ESS can coexist in the same IBM Spectrum Scale cluster.
- Second-generation and subsequent-generation IBM ESS models use RHEL.

Embedded RHEL licensing on IBM ESS 3500, IBM ESS 3200, IBM ESS 3000, and IBM ESS 5000

Starting with third-generation IBM ESS, IBM uses an embedded RHEL license. This license is ordered and included with every generation three+ order.

First- and second-generation IBM ESS Red Hat Enterprise Linux subscriptions

For all first- and second-generation IBM ESS models, Red Hat requires the following separate subscriptions to be in effect for each IBM Power server:

- RHEL Level 3 support subscription
- Red Hat Linux Extended Update Support (EUS) subscription

IBM includes the following required Red Hat subscriptions (with RHEL and EUS feature codes) in every IBM ESS order:

- First-generation IBM ESS models: IBM Program ID 5639-RH7 - Red Hat Enterprise Linux Big Endian
- Second-generation IBM ESS models: IBM Program ID 5639-RLE - Red Hat Enterprise Linux Little Endian

Client responsibilities for RHEL licensing on IBM ESS

For first- and second-generation IBM ESS models, specific customer responsibilities exist for registering the required IBM ESS RHEL subscriptions. For more information about the IBM ESS customer's responsibilities, see "How do I register the required Red Hat Enterprise subscription keys for my Elastic Storage Server (IBM ESS)?"

2.3.3 IBM Spectrum Scale

In this section, we describe IBM Spectrum Scale, the high-performance parallel clustered file system software that is used in IBM ESS in enterprise-level, high-performance computing environments, university and research environments, and commercial business environments worldwide.

Formerly known as General Parallel File System (GPFS), IBM Spectrum 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 petabyte scale, IBM Spectrum Scale provides tools for tiering, management, administration, and archiving of that enterprise-level data. IBM Spectrum Scale is the IBM strategic SDS for enterprise big data, analytics, and AI applications.

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

IBM Spectrum Scale and IBM ESS

IBM Elastic Storage System and IBM Elastic Storage Server are pre-integrated and pre-tested IBM Spectrum Scale storage building blocks.

IBM ESS 3500 uses a pair of x86-based servers, and IBM ESS 5000 uses a pair of POWER9 processor-based servers. These servers run RHEL and IBM Spectrum Scale, and they are defined as IBM Spectrum Scale NSD Data Servers. In all IBM ESS, the NSD Data Servers are cross-connected to provide failover and redundancy. If one of the NSD Data Servers fails, IBM ESS fails over the storage and data to the other NSD Data Server, which ensures continued availability of the data in the IBM Spectrum Scale cluster.

IBM Spectrum Scale for IBM ESS licensing

All nodes in an IBM Spectrum Scale cluster run a copy of the IBM Spectrum Scale software. IBM ESS typically is licensed for IBM Spectrum Scale by using the following specific IBM Spectrum Scale for IBM ESS IBM program IDs:

- 5765-DAE IBM Spectrum Scale for IBM ESS Data Access Edition
- 5765-DME IBM Spectrum Scale for IBM ESS Data Management Edition

This specific IBM Program ID with a "Per Disk" metric is normally used for licensing IBM Spectrum Scale for IBM ESS software on an IBM ESS. IBM Spectrum Scale for IBM ESS software licenses include IBM Spectrum Scale RAID license entitlement. The license price for IBM Spectrum Scale for IBM ESS is tiered as flash and HDDs have different list prices per TB. You count only the physical number of flash drives or HDDs for the "Per Disk" metric.

An advantage of the IBM Spectrum Scale for IBM ESS "Per Disk" metric is that the size of the SSD or HDD does not affect the license list price. For example, if your IBM ESS model has 550 HDDs, your IBM Spectrum Scale for IBM ESS license list price is the same, regardless of whether you are specifying 6 TB HDDs or 16 TB HDDs.

IBM Spectrum Scale for IBM ESS licensing helps to contribute in building a complete IBM hardware and software solution by integrating the IBM Spectrum Scale and IBM ESS solution.

For more information about IBM Spectrum Scale and IBM Spectrum Scale for IBM ESS software licensing, options, and considerations (such as the use of IBM ESS in a Socket-licenses IBM Spectrum Scale cluster, or for IBM Spectrum Scale capacity licenses, or for using IBM ESS in an IBM Spectrum Scale Enterprise License Agreement environment), see the following resources:

- "Where can I find detailed information about IBM Spectrum Scale and ESS licensing and pricing?"
- IBM Spectrum Scale IBM Elastic Storage System - Licensing Information

2.3.4 IBM Spectrum Scale RAID

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

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

IBM Spectrum Scale RAID implements a declustered erasure code parity schema, distributing data, redundancy information, and spare space across all disks of the IBM ESS 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 Spectrum Scale RAID implements large cache for performance by using memory on the IBM ESS NSD Data Servers. The large cache intelligently improves read and write performance, particularly for small block I/O operations.

IBM Spectrum Scale RAID mitigates performance effects of storage rebuilds

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

IBM Spectrum Scale RAID end-to-end checksums

IBM Spectrum 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 Spectrum Scale RAID is integrated with the IBM Spectrum Scale file system and performs end-to-end checksum comparison all the way out to the IBM Spectrum 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 Spectrum 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 Spectrum Scale RAID disk hospital

One of the key features of IBM Spectrum Scale RAID is th disk hospital. This powerful function asynchronously diagnoses errors and faults in the IBM ESS storage media, down to the level of the individual drive and the individual performance of each drive. IBM Spectrum Scale RAID is fully aware of and tracks the performance of each individual drives because all drives do not perform equally. IBM Spectrum 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 Spectrum Scale RAID implementation and best practices, see the Administering IBM Spectrum Scale RAID manual for your particular IBM ESS model. at IBM Documentation.

2.3.5 IBM Elastic Storage System solution installation and management scripts

In this section, we provide an overview of the components of the IBM ESS solution installation and management scripts. This overview includes the following information:

- IBM ESS-specific documentation for installation and upgrade scripts
- Ansible playbooks and container-based software installation
- IBM ESS specific tools and utilities for the IBM SSR and administrators to use while installing or maintaining IBM ESS

The third-generation IBM ESS models introduce container-based deployment methods, which use Red Hat Ansible playbooks for installation and orchestration.

For more information about these IBM ESS solution components, see the IBM ESS solution release-specific level information at IBM Documentation.

Installation, upgrade, and administering guides

IBM provides manuals and documentation for deploying and administering IBM ESS, including the following publications:

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

The IBM ESS product documentation is available at IBM Documentation.

Ansible playbooks and container-based software installation

Third-generation IBM ESS models include a software update process and methodology that is based on Ansible playbooks. The Ansible playbooks provide powerful, flexible orchestration tools that drive containerized software deployment methodology and speed and simplify the process of installing IBM ESS.

The POWER9 processor-based IBM EMS is the control management node where the Ansible software and containers are installed. The POWER9 processor-based IBM EMS is used as the single point to perform system management for second- and third-generation IBM ESS models.

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

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

For more information about **`essutils`** and **`gssutils`**, see the Quick Deployment Guide, ESS 3500 Hardware Planning and Installation Guide, and ESS software deployment preparation, which are available at IBM Elastic Storage System 3500.

3

IBM Elastic Storage System planning and integration

This chapter provides guidelines and considerations for the correct planning, installation, and configuration of the IBM Elastic Storage System (IBM ESS).

This chapter also describes configurations and integration considerations for a smooth IBM ESS deployment into an existing or a new IT environment. In this chapter, we describe the following elements:

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

This chapter includes the following topics:

3.1 IBM Elastic Storage System planning overview

In this section, we make recommendations and highlight best practices for various roles and responsibilities when an IBM ESS is deployed.

3.1.1 Roles and responsibilities

In this section, we discuss best practices were learned from successful IBM ESS installations regarding who is responsible for each of the areas of expertise.

Overall planning

An experienced project leader should be appointed to coordinate and manage the IBM ESS installation project. Deploying an IBM ESS requires coordination across different IT departments and involves various team members (IBM and client).

The various client and IBM teams that the project manager must coordinate include (but are not limited to):

- 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 ESS components are connected by using IP or InfiniBand networking. IBM ESS data users access IBM ESS by way of high-speed IP or InfiniBand networking.

The networking planning tasks that are required include (but are not limited to) include the following examples:

- Suitable networking switches must be ordered or allocated, including cable planning, ordering, racking, and proper connectors.
- Network addressing must be planned for IBM ESS and the integration of that network into existing client IP or InfiniBand network.
- After the network is designed and approved, IP addresses must be implemented and tested.
- Multiple VLANs must be designed, configured, and tested on the management network switches.
- Suitable networking authentication and firewall rules must be configured and tested to assure that they comply with client network audit and security requirements.
- Testing of the network before IBM ESS is installed (by using tools, such as **nsdperf** or Network Readiness Tools that are provided by IBM Spectrum Scale as no-charge, open-source tools) must be done to identify and verify the network's ability to provide IBM ESS with the required network bandwidth, latency, and consistent performance.

These tasks require the client IBM ESS 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 or InfiniBand network provisioning, implementation, load-testing, and network implementations, including all QA timeframes, implementation lead times, testing windows, and acceptance criteria.
- Monitor these tasks closely to assure that the networking is in place and load-tested before IBM ESS 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 ESS 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 ESS 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 Spectrum Scale before IBM ESS 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 ESS implementation.

The following best practices from successful IBM ESS installations assure success:

- Engage early with the client networking team.
- Obtain collaborative buy-in and well-defined ownership of networking-related task completions.
- Operatively running network performance tests by using IBM Spectrum Scale tools far enough in advance of the IBM ESS installation allows time for necessary tuning adjustments to be made.
- Ongoing IBM ESS implementation team and client networking team dialog and interaction during the IBM ESS 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 ESS 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 ESS components after shipment, physical floor space planning and racking; electrical power planning and provisioning, including specifying suitable electrical power, cables, connectors, and power distribution units (PDUs); and assuring proper airflow and air conditioning.

- Software planning and integration

  This planning includes defining the IBM Spectrum Scale file system implementation parameters, requirements, integration testing, and acceptance with the client applications and procedures.

- Hardware physical setup, installation, and verification

  This planning and scheduling includes the IBM ESS 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 ESS system into production are featured in this phase.

Maintaining hardware

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

- Client IBM ESS administrators:
  - Monitoring IBM ESS 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 ESS hardware maintenance and repair.

Maintaining software

After the IBM ESS is installed and running in production, many parties are responsible for various IBM ESS software tasks, including the following examples:

- Client Red Hat Linux administrators:
  - Assuring Red Hat Enterprise Linux (RHEL) licensing and subscriptions are in place for first- and second-generation IBM ESS
  - Applying required RHEL security patches as recommended by IBM
- Client IBM Spectrum Scale administrators:
  - Monitoring and tuning IBM Spectrum 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 Spectrum Scale and IBM ESS hardware. More worldwide resources might be available.

Classes that are available for IBM Spectrum Scale and IBM ESS include (but are not limited to) the following offerings:

- IBM Spectrum Scale Basic Administration
- IBM Spectrum Scale Advanced Administration
- Monitoring the IBM Elastic Storage System

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

Upgrades

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

- RHEL kernel levels and components
- IBM Power server firmware
- NVIDIA network interface card (NIC) drivers
- IBM Spectrum Scale and IBM Spectrum Scale RAID software

Note: It is a client responsibility to plan and upgrade their IBM ESS solution software.

IBM recommends upgrading the IBM ESS software solution stack at least once a year and to keep the IBM ESS current on levels of RHEL, firmware, and IBM Spectrum 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 and IBM File Object Solution Design Engine

An IBM Technical and Delivery Assessment (TDA) is a technical inspection of a completed solution design. 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 ESS solution is ordered, the IBM File Object Solution Design Engine (FOS DE) tool and process is used to examine and approve the IBM ESS solution design to determine the following information:

- Will the IBM ESS solution work as expected?
- Is the configuration and implementation plan sound?
- Will the IBM ESS 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 FOS DE processes.

For more information about the IMPACT and pre-installation TDA tools, which help your IBM account team, client technical specialist, and IBM Business Partner select the best IBM ESS and IBM Spectrum Scale solutions to match your business requirements and needs, see the following web pages:

- Preinstall TDA
- TDA and FOS DE processes for IBM Elastic Storage Server

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

3.2 IBM Elastic Storage System hardware planning

In this section, we describe the necessary planning and implementation steps that must be followed as part of installing the IBM ESS hardware.

3.2.1 Planning for network switches and networking

Suitable network planning and configuration are one of the most important success factors for an IBM ESS installation. This planning is crucial because the IBM Spectrum Scale and IBM ESS 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 ESS:

- Service network

  This network connects the flexible service processor (FSP) on the management server and I/O server nodes (with or without the Hardware Management Console (HMC), depending on the IBM ESS generation).

- Management and provisioning network

  This network connects the management server to the I/O data server nodes (and HMC if the first-generation IBM ESS machine type 5146 models are used). The management server runs the Dynamic Host Configuration Protocol (DHCP) on the management and provisioning network. If an IBM EMS is not included in the solution order, it is assumed that IBM Power IBM EMS exists in the same IBM Spectrum Scale cluster that can be used.

- Clustering network

  This high-speed data network is used for IBM ESS data transfer to and from the rest of the IBM Spectrum Scale cluster and client node access. High-speed networking is preferred (ideally 200 Gb InfiniBand-High Dynamic Range (HDR) or Ethernet), but other selected speeds are supported, such as 100 Gb InfiniBand-HDR or Ethernet.

  > **Note:** This network must provide consistent, congestion-free network packet delivery with consistent low latency. IBM can supply tools before the installation to test the network. The network might be provided by NVIDIA switches that are ordered through IBM at order time, or the network might be provided by the customer.
  >
  > For more information about for estimating and sizing suitable network performance, contact IBM or your IBM Business Partner.

- External and campus management network

  This public network is used for external and campus management of the IBM EMS, the HMC (if available), or both.

Figure 3-1 shows an example network topology for IBM ESS.

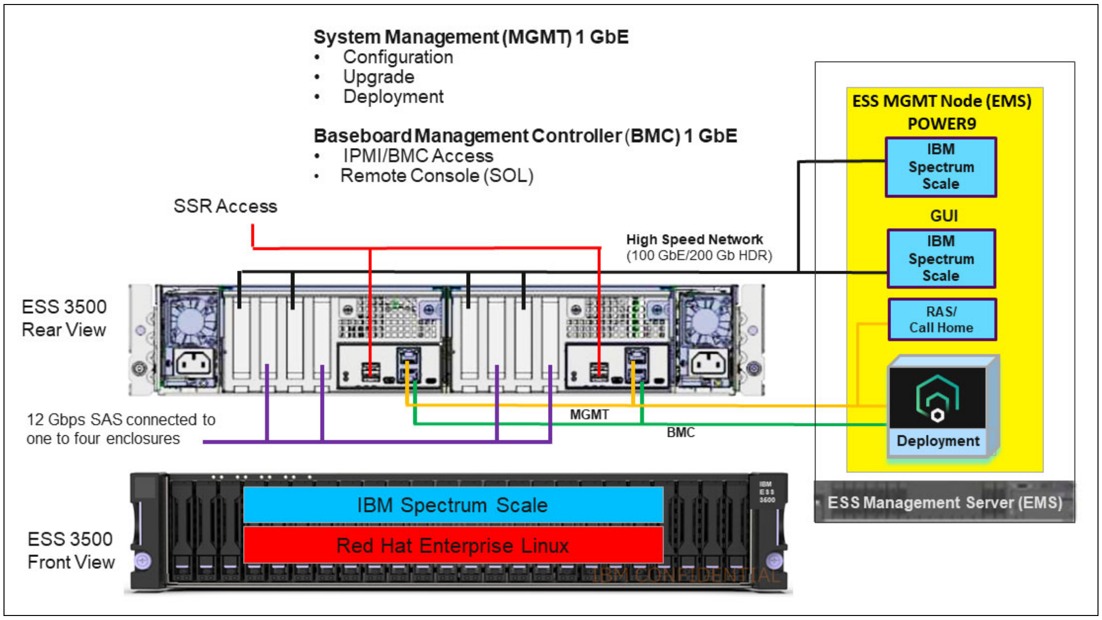

Figure 3-1 IBM ESS 3500 network topology example

An overview of the many networking steps that are required in the installation of an IBM ESS 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 ESS hardware. For more information about network planning, see the IBM ESS Quick Deployment Guide that is available at IBM Documentation.

Example: Network IP addressing, planning, and implementation for IBM ESS

In this section, we present an example of typical networking-related items for installing an IBM ESS. The following template can be used for the installation process:

1. Confirm Private IP range for DHCP server.
2. Confirm Private Service network with defined IP addresses and private management network with defined IP addresses. Separate the networks by using switches or VLAN.
3. Confirm Public network connections for IBM EMS. Multiple IP addresses are needed.
4. Select **Host** → **IP mappings** to confirm the mappings for the following items:
   - IBM EMS
   - IBM ESS 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
5. Set up domain names for management private network.
6. Set up domain names for high-speed interconnect.
7. Set up Data I/O partition and partition profile names.
8. Confirm Server names.
9. Confirm that the 10 GbE, 25 GbE, 40 GbE, 100 GbE, 200 GbE, or InfiniBand switches are in place and cabled.
10. Set up Bonding if it is used.
11. Set up Public network in place and cabled to EMS.

12. Confirm that all building block components are correctly cabled in the rack (Data I/O servers, EMS, and switches).
13. Set up and confirm dual feed power to racked IBM ESS components.
14. Set up the IBM EMS console or terminal, including the network connections.
15. Prepare for and install RHEL errata.
16. Register Red Hat subscription keys for all IBM ESS servers, as needed (if you have first or second-generation IBM ESS).
17. Define and specify IBM Spectrum Scale installation parameters and requirements. For example, how many file systems, what are the block sizes, how do we allocate and split metadata, and what replication factor is used.
18. Confirm all physical NVMe, SSD, or HDD disks are in place. IBM ESS installation checks with scripts.
19. Confirm that all physical cabling in place. IBM ESS installations are checked by scripts.
20. 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 ESS System. As part of that contract, IBM Systems Lab Services provides and works with the customer on a detailed IP networking planning document.

3.2.2 Planning for servers that are used in IBM Elastic Storage System

In this section, we review planning information for the physical servers that are used in an IBM ESS. We also describe specific recommendations for memory, NICs, and cabling and connectors.

IBM EMS

Each IBM Spectrum Scale cluster that uses an IBM ESS requires at least one IBM EMS. The EMS acts as the central management focal point for all IBM ESS systems that are in the cluster; runs specific functions, such as the IBM ESS GUI; and stores and distributes the master boot image for all the IBM ESS systems that are in the cluster by using the xCAT software (first- and second-generation IBM ESS) or by using Red Hat Ansible playbooks (third-generation IBM ESS 3500, IBM ESS 3200, IBM ESS 3000, and IBM ESS 5000).

For first-generation IBM ESS (machine type 5146), an IBM POWER8 processor-based machine type 8247-21L server is used as the EMS. This EMS runs RHEL Big Endian. Only first-generation IBM ESS are supported by the 8247-21L EMS.

For the second-generation IBM ESS, a POWER8 processor-based server with machine type 5148-21L is used as the EMS. This EMS runs RHEL Little Endian. In an IBM Spectrum Scale cluster, a 5148-21L EMS can simultaneously support second-generation IBM ESS and IBM ESS 3000.

For the third-generation IBM ESS (IBM ESS 3500, IBM ESS 5000, IBM ESS 3200, and IBM ESS 3000), an IBM POWER9 processor-based server with machine type 5105-22E is used as the EMS. This EMS runs RHEL Little Endian. If the IBM ESS 3000 is deployed into a cluster that features a 5148-21L EMS, the IBM ESS 3000 can be supported by that 5148-21L EMS.

Third-generation IBM ESS introduces an EMS software deployment mechanism that is based on containers and Red Hat Ansible orchestration playbooks.

An IBM Spectrum Scale cluster can simultaneously support all three generations of IBM ESS.

All first-generation IBM ESS are supported by an POWER8 processor-based 8247-21L IBM EMS running RHEL Big Endian. All second-generation and IBM ESS 3000 are supported by an POWER8 processor-based 5148-21L IBM EMS running RHEL.

All third-generation IBM ESS are supported by IBM POWER9 processor-based 5105-22E IBM EMS running RHEL Little Endian. Announced on 23 February 2021, this POWER9 processor-based 5105-22E IBM EMS can also support second-generation IBM ESS.

IBM ESS Data Servers

Every IBM ESS is an integrated IBM Spectrum Scale storage building block that consists of two IBM Spectrum 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 ESS storage is running through the first data server, and the other half of the IBM ESS 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 ESS, which provides HA and resiliency.

Each first-generation IBM ESS (machine type 5146) storage building block uses two POWER8 processor-based servers (machine type 8247-22L) each acting as an IBM Spectrum Scale Data Server. These 8247-22Ls run RHEL Big Endian. Only first-generation IBM ESS use the 8247-22L RHEL Big Endian-based data servers.

In the second-generation IBM ESS, two POWER8 processor-based servers with machine type 5148-22L are used as the IBM ESS Data Servers. These servers are physically identical to 8247-22L servers, but carry an IBM ESS-specific machine type of 5147 to assure that IBM Service and Support always know that these servers are part of an overall IBM ESS integrated solution. These 5148-22L POWER8 processor-based servers run RHEL Little Endian.

An IBM Spectrum Scale cluster fully supports a mix of all generations of IBM ESS in the same cluster. A 5148-21L EMS can simultaneously support second-generation IBM ESS and IBM ESS 3000. A 5148-21L EMS cannot support a first-generation IBM ESS, IBM ESS 3500, IBM ESS 3200, or IBM ESS 5000.

For the third-generation IBM ESS 3500, IBM ESS 3200, and IBM ESS 3000, the 2U24 form factor uses an integrated pair of x86-based servers as IBM Spectrum Scale Data Servers.

For the third-generation IBM ESS 5000, two POWER9 processor-based servers with machine type 5105-22E are used as the IBM ESS Data Servers. These servers are specific POWER9 processor-based servers for IBM ESS 5000. They carry an IBM ESS specific machine type of 5105-22E to assure that IBM Service and Support always know that these servers are part of an overall IBM ESS-integrated solution. These 5105-22E POWER9 processor-based servers run RHEL Little Endian.

An IBM Spectrum Scale cluster can simultaneously support all three generations of IBM ESS. First-generation IBM ESS in the cluster are supported by an 8247-21L IBM EMS running RHEL Big Endian.

All second-generation and IBM ESS 3000 systems in the cluster are supportable by the 5148-21L IBM EMS or 5105-22E IBM EMS, both of which run RHEL Little Endian.

All third-generation IBM ESS 3500, IBM ESS 3200, and IBM ESS 5000 require a 5105-22E POWER9 processor-based IBM EMS that runs RHEL Little Endian. Announced on 23 February 2021, the POWER9 processor-based 5105-22E IBM EMS also supports second-generation IBM ESS.

Memory

POWER9 processor-based DIMM memory cards (specified with POWER9 processor-based server memory feature codes) are used in the 5105-22E servers. The default memory chip is a 32 GB DIMM.

The default memory size for the 5105-22E when it is used as a data server is 384 GB, which consists of twelve 32 GB DIMMs. When a 5105-22E is used as an IBM EMS, the default memory is 128 GB, which consists of eight 16 GB DIMMs.

When 5105-22E is used as IBM ESS Protocol node with single socket, the default memory is 128 GB, consisting of 16 GB DIMMs with quantity of 8. When 5105-22E is used as IBM ESS Protocol node with two sockets, the default memory is 192 GB, consisting of 16 GB DIMMs with quantity of 12. This memory is enough for most IBM ESS installations.

POWER8 processor-based DIMM memory cards (specified with POWER8 processor-based server memory feature codes) are used in the 8247-21L, 8247-22L, 5148-21L, and 5148-22L servers. The memory size for the 8247-21L and the 5148-21L EMS was originally 32 GB. IBM ESS 5148-21L IBM EMS default memory is now 64 GB.

The IBM EMS memory size is enough for most IBM ESS installations. If many IBM ESS systems are used in your IBM Spectrum Scale IBM ESS configuration, contact your IBM representative to see whether larger IBM EMS memory sizes might be required for your installation. More EMS memory can be specified at order time or added later as a field miscellaneous equipment specification (MES).

The standard memory for the IBM ESS Data Servers is enough for most IBM ESS usage cases. If you have a large workload configuration, or you have a specific workloads type, larger memory might be needed on your IBM ESS Data Servers. Check with your IBM representative to see whether you have specific workloads that might benefit from more main memory.

Standard IBM Power memory upgrades for IBM ESS Data Servers can be ordered and installed without disrupting the access to the IBM Spectrum Scale file system. In the installation process, the storage in one of the two IBM ESS Data Servers that are upgraded is failed over to the other IBM ESS Data Server.

This procedure allows the first IBM ESS Data Server to be gracefully brought down and then upgraded. Then, the same process is repeated for the other IBM ESS Data Server to be upgraded.

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

IBM ESS 3000 uses a pair of x86-based servers that are integrated into the back of the 2U24 storage enclosure. The standard memory for each of these IBM ESS 3000 Data Servers is 384 GB. This memory is enough for most of IBM ESS 3000 usage cases. You can specify another 384 GB of main storage per data server at order time (or later by field upgrade), which results in a total of 768 GB of memory per IBM ESS 3000 Data Server.

IBM ESS 3200 uses a pair of PCI Gen4 x86-based servers that are integrated into the back of the 2U24 storage enclosure. The standard memory for each of these IBM ESS 3200 Data Servers is 512 GB. This memory is enough for most of IBM ESS 3200 usage cases. As of this writing, no option is available to specify any extra main memory per IBM ESS 3200 Data Server.

Network interface cards

The IBM ESS features the following types of networks that are used to interconnect the IBM ESS components:

- Management and service

  Mandatory, separate 1 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 Spectrum Scale and IBM ESS file system.

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

Make sure that the network bandwidth requirements are clearly studied, documented, and communicated. Specify a suitable number of NICs to meet expected IBM ESS performance requirements. Plan and specify the correct cables and connector types.

> **Note:** IBM Power IBM ESS Data Servers do not support every possible IBM Power PCI adapter.

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 ESS 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 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

3.2.3 Planning for storage configuration

IBM ESS storage configuration for ordering purposes is relatively straightforward.

First, determine your IBM Spectrum Scale storage requirements. Request that your IBM representative or IBM Business Partner use the FOS DE 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 ESS feature the suitable type and quantity of storage media available. Within each model, you choose the suitable drive size. Request that your IBM Representative or IBM Business Partner use the FOS DE tool to verify the performance of your chosen IBM ESS model with your requirements.

Part of the value of IBM ESS is that after storage performance and capacity requirements are known, IBM ESS models provide known performance and usable capacity characteristics because these factors were considered when IBM designed the IBM ESS model.

Each IBM ESS 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 ESS architecture. This situation demonstrates the value of the use of IBM ESS as an optimum and tested method to deploy IBM Spectrum Scale storage by using an efficient building block approach.

3.2.4 Planning for IBM Spectrum Scale protocol nodes

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

From a planning perspective, IBM Spectrum Scale requires that all protocol nodes in an IBM Spectrum Scale clusters are all x86, IBM Power Little Endian, or IBM Power Big Endian.

No protocol node intermixing of these nodes is allowed in the same IBM Spectrum Scale cluster.

IBM ESS coexists in an IBM Spectrum Scale cluster with protocol nodes. IBM ESS also has no effect on which type of protocol node can be used in an IBM Spectrum Scale cluster.

With the IBM ESS 3000, IBM ESS 3200, and IBM ESS 5000, IBM provided the option of ordering IBM Spectrum Scale protocol nodes based on the POWER9 processor-based 5105-22E server. These protocol nodes are supported by the IBM EMS and can be managed by the IBM ESS GUI and by IBM ESS installation tools, which provides greater ease of use and management capability for ordering a complete IBM Spectrum Scale or IBM ESS or protocol node solution from IBM.

The 5105-22E IBM Spectrum Scale protocol nodes can coexist with any other IBM Power Little Endian protocol nodes in the same IBM Spectrum Scale cluster.

For more information about 5105-22E protocol nodes, see *Quick Deployment Guide*, which is available at IBM Documentation.

3.2.5 Planning for physical, electrical, and cooling installation

Preparation for installation of your IBM ESS system is as important as the physical installation of the system. To have a smooth and successful installation, preparation processes, and TDAs must be followed.

FOS DE and pre-installation meetings with IBM TSS and IBM Service personnel must be completed. Also, IBM Elastic Storage Server (IBM ESS) 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 also deployment of the software.

The specifications for IBM ESS can be found in IBM Documentation.

3.3 IBM Elastic Storage System software planning

In this section, we highlight key items to be aware of for planning the installation of the IBM ESS software stack; in particular, IBM Spectrum Scale and IBM Spectrum Scale RAID.

3.3.1 Planning for IBM Spectrum Scale file system

The IBM Spectrum Scale file systems that are placed into the IBM ESS 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 Spectrum Scale and IBM Spectrum Scale RAID Administration should be consulted. These manuals contain detailed instructions and best practices recommendations for implementing the requirements of the client's environment.

3.3.2 Planning for IBM Spectrum Scale RAID

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

The IBM Elastic Storage System and IBM Elastic Storage Server implementation of IBM Spectrum 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 ESS, the IBM Spectrum Scale RAID implementation best practices are enforced as de facto standards by the IBM ESS 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 ESS 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.

3.4 Ordering IBM Elastic Storage System

IBM ESS is always part of a larger IBM Spectrum Scale system. Because IBM Spectrum Scale and IBM ESS are sophisticated technologies to assure high client satisfaction, IBM requires that a TDA and the FOS DE tools are used to elicit requirements, assure that designs are reviewed, assessed, and approved.

The IBM ESS solution process engages 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 ESS solution that matches the client's needs.

Before confirming the IBM ESS order, the IBM team or the Business Partner uses the IBM FOS DE tool. FOS DE is a question and answer tool that elicits and documents the various technical requirements for a solution that uses IBM ESS. Contact your IBM representative to request that a pre-order IBM ESS FOS DE design evaluation is performed.

IBM eConfig is used by the IBM or Business Partner team to configure the IBM ESS for ordering. eConfig is a tool that helps the IBM or Business Partner team to configure the correct combination of IBM Storage Enclosures, servers, NICs, and associated feature codes to make up the wanted IBM ESS.

Note: An approved pre-installation TDA and an approved FOS DE run *must* be performed by the IBM account team or IBM Business Partner to have an IBM ESS order to be entered and shipped.

As part of the order, various decisions regarding hardware must be made. Be prepared to supply requirement information to the IBM Business Partner. The following requirements often are included in the information that is sent:

- Storage requirements (total capacity, HDD or SSD sizes, and wanted performance)
- Server memory and NIC requirements
- Wanted network cables, cable lengths, connector types, and network switches
- Physical electrical power requirements (phase and voltage)
- Any other more requirements (Your IBM or Business Partner representative works with you to determine these requirements.)

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

Figure 3-2 on page 55 shows an example of some possible IBM ESS configurations in one rack. Shown on the left in Figure 3-2 on page 55 are the size of various IBM ESS components in rack units. Shown on the right of Figure 3-2 on page 55 are examples of four IBM ESS SL models. These configurations are examples only; your configuration is tailored to your requirements.

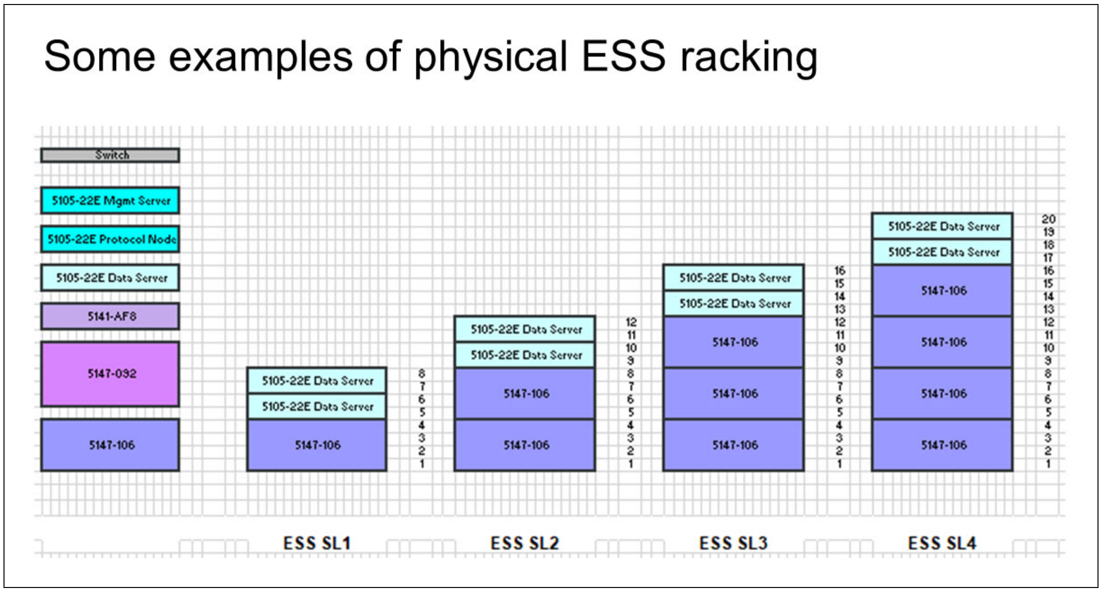

Figure 3-2 IBM Elastic Storage Server physical racking examples

If an IBM ESS 5000 is upgraded with more storage enclosures, it is not necessary to reposition the POWER9 processor-based 5105-22E Data Servers. IBM installs the additional storage enclosures above the POWER9 processor-based data servers in the rack.

3.5 IBM Elastic Storage System installation

In this section, we describe the process that is used to physically install the IBM ESS and deploy it into the client's IBM Spectrum Scale environment.

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

3.5.1 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 ESS components.

IBM TSS works with you to help you plan for the many items to consider for racking, installation, and servicing of the IBM ESS. 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
- Airflow considerations (especially IBM ESS SCx models), which use the deep IBM 5147-106 Storage Enclosures
- Shipping location of the system

- Location of systems that are installed in customer racks (if rackless; that is, not installed in IBM rack)
- IP addresses, networking, and switch configuration that is needed to install and connect the IBM ESS components
- Lift tools
- Ladders
- IBM SSR physical installation instructions (WCII) and IBM SSR scheduling and confirmation

If you have an IBM ESS HDD model, allow time for the physical installation of the many HDDs into the storage enclosures. HDD IBM ESS 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 ESS, many HDDs are used and it takes time for the IBM SSR to install each HDD into its slot in the storage enclosures.

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

- Whether to reserve space in a rack for future expansion
- Location of rack in data center for purposes of future cable lengths and electrical power supply expansion
- Location of network switches and the IBM EMS for purposes of cable lengths for networking

3.5.2 Integrating into client environment

Integrating the IBM ESS into the clients' environment consists of many actions. In this section, we provide an overview of key best practices, practical advice, and experiences and lessons learned for a successful IBM ESS deployment.

Integrating into customer networking environment

In any IBM Spectrum Scale or IBM ESS 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 Spectrum Scale or IBM ESS 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 assure that your IP network is ready for your IBM Spectrum Scale or IBM ESS solution to integrate IBM ESS into your environment.

Will a new or existing network be used

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

IBM Spectrum Scale or IBM ESS is a high-performance, high-data rate big data solution. Best practices include providing the IBM Spectrum Scale or IBM ESS 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 Spectrum Scale or IBM ESS. Instead, the IBM Spectrum Scale or IBM ESS 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 Spectrum Scale or IBM ESS 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 Spectrum Scale or IBM ESS cluster. IBM Spectrum Scale or IBM ESS depends upon IP networking for all of the IBM Spectrum 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 ESS will be placed and how they will be connected.

- Total number of nodes in IBM Spectrum Scale or IBM ESS cluster.

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

Are there 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 ESS and assure that number of ports and speeds are available within the networking infrastructure.

What are the interconnects between switches

Document the inter-switch links, where they are located, and what traffic traverses them.

Can existing networking links and switches be used

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

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

For more information, see this GitHub repository.

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

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

What are the basic workload types

Document what other traffic IBM Spectrum 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 Spectrum Scale architect with networking skills to project how the IBM Spectrum Scale or IBM ESS network infrastructure performs in your specific environment.

Sustained bandwidth required in GBps

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?

Is validation 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 Spectrum Scale or IBM ESS solution after reviewing these questions and answers, use the **nsdperf** tool. For more information about this tool, see this GitHub repository.

Integrating into a customer software environment

IBM Spectrum Scale and IBM ESS 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 Spectrum Scale or IBM ESS solution.

Design the IBM Spectrum Scale or IBM ESS cluster to deploy the wanted POSIX, Linux, or UNIX file systems. If necessary, design and deploy the IBM Spectrum 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 ESS is a part of an IBM Spectrum Scale cluster. IBM ESS and IBM Spectrum 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 Spectrum Scale or IBM ESS solution.

For more information about IBM Spectrum Scale authentication and security, see *IBM Spectrum Scale Security*, REDP-5426.

3.5.3 Reliability, availability, and serviceability

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

IBM ESS’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 ESS hardware

The IBM EMS and I/O server hardware includes IBM Call Home capability.

In the first-generation IBM ESS, the Hardware Management Console (HMC) provided the Call Home focal point for service receiving, logging, tracking system errors, and, if enabled, forwarding problem reports to IBM Service for hardware events.

When a serviceable event occurs on one of the monitored components (such as servers or storage), the Call Home generates a Call Home event.

In the second and subsequent generations of IBM ESS, the HMC is no longer required.

IBM Call Home: IBM Spectrum Scale software

IBM Spectrum Scale software also has its own separate Call Home capability for purposes of automatically collecting and securely uploading data that is related to IBM Spectrum Scale to IBM Service and Support. These functions are separate from the IBM ESS Call Home functions.

3.6 Monitoring IBM ESS

Monitoring the IBM ESS system includes system health, performance, and capacity monitoring. You can monitor the system through IBM ESS GUI or with the help of CLI.

For more information, search for the topic “Monitoring IBM Spectrum Scale RAID” at IBM Documentation.

3.6.1 GUI

The IBM ESS GUI provides robust facilities to monitor and manage the system. IBM Documentation provides more information about how you can configure the IBM ESS GUI to monitor the performance of the following functional areas in the system:

- Network
- System resources
- IBM Spectrum Scale RAID
- NSD server
- IBM Spectrum Scale client
- NFS
- SMB
- Object
- Clustered Trivial Data Base (CTDB), which is used for SMB
- Transparent cloud tiering
- Waiters
- Active File Management (AFM)

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

3.7 Maintaining and repairing IBM ESS

For IBM ESS features, a full library of documentation for maintaining the hardware and software of the IBM ESS solution is available at IBM Documentation. At the website, select the IBM ESS solution release level that you are using.

3.7.1 IBM ESS IBM Solution service and support

IBM provides service and support for IBM ESS as a solution. IBM ESS-specific keywords and procedures are in place so that you, as a client, when calling or electronically requesting IBM ESS support, are routed to an IBM ESS cross-component solution support team that understands the interaction of all of the IBM ESS components. They also can 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 Elastic Storage System (ESS) Support Reference Guide*.

3.8 Upgrading IBM ESS

Maintaining and upgrading your IBM ESS software solution stack is a customer's responsibility.

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

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

Existing IBM ESS customers are entitled to download the newer IBM ESS solution stack from IBM Fix Central. IBM recommends that you plan to upgrade your IBM ESS 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 ESS upgrades when necessary.

3.8.1 IBM ESS hardware and firmware

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

Typically, the IBM ESS 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 ESS also supports nondisruptive addition of more storage enclosures, which allows an IBM ESS to be non-disruptively upgraded to a larger capacity. For example, with IBM ESS 5000, this process supports the addition of a 5147-092 to an SLx model or a 5147-106 to a SCx model. Therefore, an installed SC1 to SC8 model, or SL1 to SL6 model, can be upgraded to any model configuration up to SC9 or SL7 respectively, without causing downtime or an interruption in service.

An IBM ESS 3000 or IBM ESS 3200 that was ordered half-populated with 12 NVM drives, can be upgraded to fully populated 24-drive IBM ESS 3000 in a similar manner, by adding the NVMe drives in the front of the IBM ESS 3000 Storage Enclosure. No downtime or interruption in service is required.

The client is responsible for keeping the IBM ESS operating system, software, firmware, and GUI current with the recommended levels of software. For more information about upgrade instructions, see IBM Documentation.

Note: IBM and IBM Business Partners might offer IBM Services® to assist or perform these upgrades on behalf of the client.

For all hardware upgrades, clients work with the IBM or their Business Partner to match their needs with the IBM ESS upgrades and services offerings that are available.

3.8.2 IBM ESS software

IBM ESS 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 Spectrum Scale
- Cluster export services code (CTDB, Samba for SMB, Ganesha for NFS, and OpenStack Object Storage (Swift))
- xCAT (eXtreme Cloud Administration Tool) (if you use first- or second-generation IBM ESS)
- IBM Power or x86 server firmware
- RHEL kernel at specific level (for example: 3.10.0-957.21.3)
- RHEL system errata at specific level (for example: 219-62.el7_6.6)
- RHEL Network Manager at specific level (for example: 1.12.0-10.el7_6)
- RHEL device drivers (IPR)
- NVIDIA OpenFabrics Enterprise Distribution (OFED) network adapter drivers
- Electronic Service Agent (ESA)
- Hardware Management Console (HMC) if this installation is first-generation IBM ESS

For IBM ESS 3000, IBM ESS 3200, and IBM ESS 5000, Red Hat Ansible playbooks and software is used to manage and provide significantly improved ease of use for upgrading IBM ESS software components

Keeping the IBM ESS software solution stack current is a customer responsibility. IBM recommends that clients plan to upgrade their IBM ESS software solution stack at least once a year. Typically, the IBM ESS 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 ESS software solution release level, see IBM Documentation.

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

New levels of the IBM ESS solution software stack are available at IBM Fix Central. For more information about these newer levels of IBM ESS software, log on to IBM Fix Central by using your IBMid. At the website, in the Product Selector field, enter `IBM Elastic Storage System`. Then, select your IBM ESS version or release level, your platform (for the latest IBM ESS, this platform often is Linux on IBM Power PC 64 Little Endian).

A list of available IBM ESS solution software fix packs for download is shown. When you start the download, IBM FixCentral checks your IBMid authentication profile to assure 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 ESS solution software.

> **Note:** If you confirmed that you are authorized to download fix packs but receive a message that says, “No applicable support agreement found”, this message indicates that your IBMid profile does not match the IBM entitlement database. In this case, click **Request Help** in the download window. An IBM Help ticket is opened and IBM contacts you to help resolve the matter.

For more information about applying the IBM ESS Fix Pack, see the IBM ESS Quick Deployment Guide.

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

3.9 IBM ESS hints and tips

Consider the following hints and tips:

- Physically plan ahead and think through the following issues:
  - Where the IBM ESS servers and components are to go in the rack.
  - Where to place the network switch and the server locations into the rack’s position.
  - Cabling and type of adapter connectors. Always follow the recommendations from IBM.
- IBM recommends that you use defaulted tuning parameters, unless your IBM or IBM Business Partner SME reviewed your environment and recommended other suitable choices.
- Always test and qualify your IP network before installing IBM ESS. IBM Spectrum Scale and IBM ESS can achieve high performance and use all available network bandwidth. Tune your network specifically for optimum IBM Spectrum Scale and IBM ESS performance.
- Have all networks planned, implemented, and tested before IBM ESS hardware is installed. For example, identify and implement network configuration, VLANs, IP addresses, and authentication and security passwords. Make sure that your IP network environment meets IBM’s IBM ESS network criteria and is verified by an IBM SME.
- Define clear IBM ESS performance criteria expectations and test plans to verify performance.

- Use only IBM ESS and IBM Spectrum Scale certified machines and components in your IBM ESS hardware and software environments.
- Ensure that every IBM ESS cluster includes at least one IBM EMS and one management network switch.
- Plan your education and skill sets for IBM Spectrum Scale and IBM ESS so that you can understand and use this technology to its fullest capabilities.
- IBM recommends that each IBM ESS building block has the same IBM ESS code level. For migration or upgrade purposes, adhere to requirements for IBM ESS code levels that are allowed to coexist during the migration or upgrade within each cluster. For example, ensure that a code level must be no more than two different OFED levels apart.
- Ensure that each IBM ESS building block is integrated into the same rack; it cannot be split into multiple racks.

A

IBM Elastic Storage System models

In this appendix, we describe information and differences between the three generations of IBM Elastic Storage Systems (IBM ESS) that are used today. Understanding the nomenclature and the differences between these generations is helpful when you are designing your IBM ESS environment, especially if different generations of IBM ESS coexist in the same IBM Spectrum Scale cluster.

This appendix includes the following topics:

IBM Elastic Storage System model specifications

This section describes the specifications of the various IBM ESS models.

First-generation IBM Elastic Storage System specifications

The first-generation IBM ESS models were first announced and shipped in 4Q2014. These models were withdrawn from marketing in January 2018.

The end-of-service date for the first-generation IBM ESS was announced in July 2020, with an effective date of 31 December 2021. First-generation IBM ESS models are no longer serviced or supported by IBM.

The first-generation IBM ESS consisted of the models that are listed in Table A-1.

Table A-1 First-generation IBM ESS building block models

IBM ESS model	Enclosure type /U/ IBM machine type - model	Enclosures	Number of drives[a]
GS1	EXP24S 2U No machine/type. Specified as feature code on POWER processor-based IBM ESS servers.	1	24 solid-state drives (SSDs)
GS2		2	46 serial-attached SCSI (SAS) + 2 SSDs or 48 SSDs
GS4		4	94 SAS + 2 SSDs or 96 SSDs
GS6		6	142 SAS + 2 SSDs
GL2	DCS3700 4U Machine type 1818-80E	2	116 NL-SAS + 2 SSDs
GL4		4	232 NL-SAS + 2 SSDs
GL6		6	348 NL-SAS + 2 SSDs

a. On the hard disk drive (HDD) configurations, the two SSDs are not visible user capacity. These SSDs hold data that is used internally for IBM Spectrum Scale RAID metadata only.

Consider the following points regarding network interface cards (NICs):

- The IBM ESS GSx and GLx models support three high-speed data NICs per IBM POWER8 processor-based data server.
- Because two POWER8 processor-based data servers are used per IBM ESS building block, a total of six high-speed data NICs are used per IBM ESS GSx or GLx building block.

Second-generation IBM Elastic Storage System specifications

This section describes the second-generation IBM ESS specifications.

The second-generation IBM ESS models were announced in April 2017. At the time of writing, these models are still actively supported and serviced.

The second-generation IBM ESS does not include a solution IBM machine type. The second-generation IBM ESS server and storage hardware components feature the following IBM ESS-unique machine types:

- IBM machine type 5148 for the IBM ESS POWER8 processor-based servers
- IBM machine type 5147 for the IBM ESS IBM Storage Enclosures

These machine types uniquely identify these hardware components as part of an IBM ESS solution.

The second-generation IBM ESS models are available in the following types:

- GSxS models for all SSD storage
- GHxx models hybrids of some SSD storage and some HDD storage
- GLxS models for high-capacity HDD storage
- GLxC models for ultra high-capacity and density HDD storage

These GLxS and GSxS models were first announced and available in the second quarter of 2017. They are fully supported by IBM Service and Support for hardware, firmware, software updates, and maintenance.

The SSD flash storage second-generation IBM ESS models GSxS specifications are listed in Table A-2.

Table A-2 Second-generation IBM ESS SSD flash building block models

IBM ESS model	Enclosure U/Number drives IBM machine type - model	Enclosures	Number of drives
GS1S	2U24 Machine type 5147-024	1	24 SSDs
GS2S		2	48 SSDs
GS4S		4	96 SSDs

The HDD disk drive storage second-generation IBM ESS models “GLxS” specifications are shown in Table A-3.

Table A-3 Second-generation IBM ESS HDD hard disk building block models

IBM ESS model	Enclosure U /Number drives IBM machine type - model	Enclosures	Number of drives[a]
GL1S	5U84 Machine type 5147-084	1	82 NL-SAS + 2 SSDs
GL2S		2	166 NL-SAS + 2 SSDs
GL3S		3	250 NL-SAS + 2 SSDs
GL4S		4	334 NL-SAS + 2 SSDs
GL5S		5	418 NL-SAS + 2 SSDs
GL6S		6	502 NL-SAS + 2 SSDs

a. On the HDD configurations, the two SSDs are not visible user capacity. These SSDs hold data that is used internally for IBM Spectrum Scale RAID metadata only.

Consider the following points regarding NICs:

- The IBM ESS GSxS and GLxS models support three high-speed data NICs per POWER8 processor-based data server.
- Because two POWER8 processor-based data servers are used per IBM ESS building block, a total of six high-speed data NICs are used per IBM ESS GSxS or GLxS building block.

The ultra-high capacity HDD storage second-generation IBM ESS models GLxC specifications are as listed in Table A-4.

Table A-4 Second-generation IBM ESS HDD hard disk building block models

IBM ESS model	Enclosure U /Number drives IBM machine type - model	Enclosures	Number of drives[a]
GL1C	4U106 Machine type 5147-106	1	104 NL-SAS + 2 SSDs
GL2C		2	210 NL-SAS + 2 SSDs
GL3C		3	316 NL-SAS + 2 SSDs
GL4C		4	422 NL-SAS + 2 SSDs
GL5C		5	528 NL-SAS + 2 SSDs
GL6C		6	634 NL-SAS + 2 SSDs
GL8C		8	846 NL-SAS + 2 SSDs

a. On the HDD configurations, the two SSDs are not visible user capacity. These SSDs hold data that is used internally only for IBM Spectrum Scale RAID metadata.

Consider the following points regarding NICs:

- The IBM ESS GLxC models include another SAS adapter per POWER8 processor-based server, compared to the GSxS and GLxS models. This addition provides enough SAS connections for the GL8C model to be built.
- The GLxC models support two high-speed data NICs per POWER8 processor-based data server.
- Because two POWER8 processor-based data servers are used per IBM ESS building block, a total of four high-speed data NICs are used per IBM ESS GLxC building block.

When the ultra-high-capacity IBM ESS GLxC models are configured, ensure that the NIC capacity and bandwidth is sufficient for your workload. If you need higher network bandwidth capability than four high-speed data NICs per POWER8 processor-based data server can provide, choose one of the other IBM ESS models that has three high-speed data NICs per POWER8 processor-based data server.

Second-generation IBM Elastic Storage System Hybrid models

A primary use of the GHxy hybrid models is for a less expensive option to combine an IBM Spectrum Scale SSD storage tier with an IBM Spectrum Scale HDD capacity storage tier.

The IBM ESS GHxy Hybrid models combined one or two drawers of SSD with two or four drawers of 5U84 HDD storage. You can think of these models as a combination of the GSxS models and the GLxS models.

IBM ESS Hybrid models are IBM ESS storage building blocks that consist of SSD flash 5147-024 Storage Enclosures and HDD 5147-084 Storage Enclosures. The IBM ESS Hybrid models (model GHxy) are a blend of 2U24 SSD enclosures and 5U84 HDD enclosures, where “x” is the number of 2U24 SSD enclosures, and “y” is the number of 5U84 HDD enclosures. The following models are available:

- GH12
- GH22
- GH14
- GH24

The IBM ESS GHxy models use the same POWER8 processor-based data server. They also feature the same NIC capacity as the IBM ESS GSxS and GLxS models.

Third-generation IBM Elastic Storage System 3000 specifications

The IBM Elastic Storage System 3000 (IBM ESS 3000) is one of the third-generation IBM ESS models.

The specifications of the IBM ESS 3000 model are listed in Table A-5.

Table A-5 Third-generation IBM ESS NVMe flash building block model

IBM ESS model	Enclosure U/Number drives IBM machine type - model	Enclosures	Number of drives
IBM ESS 3000	2U24 Machine type 5141-AF8	1	12 or 24 NVMe. Available NVMe drive sizes are 3.84 TB, 7.68 TB, or 15.36 TB.

Third-generation IBM Elastic Storage System 3200 specifications

The IBM Elastic Storage System 3200 (IBM ESS 3200) is one of the third-generation IBM ESS models.

The specifications of the IBM ESS 3200 model are listed in Table A-6.

Table A-6 Third-generation IBM ESS NVMe flash building block model

IBM ESS model	Enclosure U/Number drives IBM machine type - model	Enclosures	Number of drives
IBM ESS 3200	2U24 Machine type 5141-FN1	1	12 or 24 NVMe. Available NVMe drive sizes are 3.84 TB, 7.68 TB, or 15.36 TB.

Third-generation IBM Elastic Storage System 3500 specifications

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

The IBM ESS 3500 model includes NVMe flash storage.

The specifications of the IBM ESS 3500 model are listed in Table A-7.

Table A-7 Third-generation IBM ESS NVMe flash building block model

IBM ESS model	Enclosure U/Number drives IBM machine type - model	NVMe drives	Drive size
IBM ESS 3500 or IBM ESS 3500 H0	2U24 Machine type 5141-FN2	NVMe 12 or 24	3.84 TB, 7.68 TB, or 15.36 TB

The IBM ESS 3500Hx, shown in Table A-8, includes machine type 5141-FN2 and 1-4 model 5147-102.

Table A-8 IBM ESS 3500Hx enclosure specifications

IBM ESS Model	Enclosure U/Number drives IBM machine type - model	NVMe	HDD enclosure	HDD drives	Drive size
IBM ESS 3500H1	2U24 Machine type 5141-FN2 + 1@4U external storage Model 5147-102	Quantity: 12 or 24 Size: 3.84 TB, 7.68 TB, or 15.36 TB	1	52 or 102	HDD: 10 TB, 14 TB, 18 TB, or 20 TB
IBM ESS 3500H2	2U24 Machine type 5141-FN2 + 2@4U external storage Model 5147-102		2	204	
IBM ESS 3500H3	2U24 Machine type 5141-FN2 + 3@4U external storage Model 5147-102		3	306	
IBM ESS 3500H4	2U24 Machine type 5141-FN2 + 4@4U external storage Model 5147-102		4	408	

The IBM ESS 3500Cx, shown in Table A-9, includes machine type 5141-FN2 logtip only NVMe drives and 1 - 4 model 5147-102.

Table A-9 IBM ESS 3500Cx enclosure specifications

IBM ESS Model	Enclosure U/Number drives IBM machine type - model	NVMe	HDD enclosure	HDD drives	Drive size
IBM ESS 3500C1	2U24 Machine type 5141-FN2 + 1@4U external storage Model 5147-102	Quantity: 4 (used for logtip)	1	52 or 102	HDD: 10 TB, 14 TB, 18 TB, or 20 TB
IBM ESS 3500C2	2U24 Machine type 5141-FN2 + 2@4U external storage Model 5147-102		2	204	
IBM ESS 3500C3	2U24 Machine type 5141-FN2 + 3@4U external storage Model 5147-102		3	306	
IBM ESS 3500C4	2U24 Machine type 5141-FN2 + 4@4U external storage Model 5147-102		4	408	

Third-generation IBM Elastic Storage System 5000 specifications

The IBM Elastic Storage System 5000 is part of the third generation of IBM ESS offerings.

The third-generation IBM ESS 5000 model with IBM POWER9 processor-based data servers and 5147-092 or 5147-106 Storage Enclosures was announced in July 2020 and general availability was August 2020.

The specifications of the IBM ESS 5000 SL model are listed in Table A-10.

Table A-10 Third-generation IBM ESS 5000 SL model - HDD hard disk building block

IBM ESS model	Enclosure U/Number drives IBM machine type - model	Enclosures	Number of drives[a]
SL1	5U92 Machine type 5147-092 HDD drive sizes that are available: 6 TB, 10 TB, 14 TB, or 16 TB	1	90 NL-SAS + 2 SSDs
SL2		2	182 NL-SAS + 2 SSDs
SL3		3	274 NL-SAS + 2 SSDs
SL4		4	366 NL-SAS + 2 SSDs
SL5		5	458 NL-SAS + 2 SSDs
SL6		6	550 NL-SAS + 2 SSDs
SL7		7	642 NL-SAS + 2 SSDs

a. On the HDD configurations, the first storage enclosure has two slots that are taken up by SSDs, which are not visible user capacity. These SSDs hold data that is used internally by IBM Spectrum Scale RAID metadata only.

The specifications of the IBM ESS 5000 SC model are listed in Table A-11.

Table A-11 Third-generation IBM ESS 5000 SC model - HDD hard disk building block

IBM ESS model	Enclosure U/Number drives IBM machine type - model	Enclosures	Number of drives[a]
SC1	4U106 Machine type 5147-106 HDD drive sizes that are available: 10 TB, 14 TB, or 16 TB	1	104 NL-SAS + 2 SSDs
SC2		2	210 NL-SAS + 2 SSDs
SC3		3	316 NL-SAS + 2 SSDs
SC4		4	422 NL-SAS + 2 SSDs
SC5		5	528 NL-SAS + 2 SSDs
SC6		6	634 NL-SAS + 2 SSDs
SC7		7	740 NL-SAS + 2 SSDs
SC8		8	846 NL-SAS + 2 SSDs
SC9		9	952 NL-SAS + 2 SSDs

a. On the HDD configurations, the first storage enclosure has two slots that are taken up by SSDs, which are not visible user capacity. These SSDs hold data that is used internally by IBM Spectrum Scale RAID metadata only.

First-generation IBM Elastic Storage System models overview

You might encounter IBM Spectrum Scale clusters that include the first-generation IBM ESS models. Although these models continue to operate, IBM support for these models ceased on 31 December 2021.

First-generation IBM Elastic Storage System models

The first-generation IBM ESS models were first announced and shipped in 4Q2014. These models, including upgrades, were withdrawn from marketing in January 2018. As of 31 December 2021 the first-generation IBM ESS models are no longer supported by IBM Service and Support for hardware and software maintenance. First-generation IBM ESS models also no longer receive IBM ESS solution stack hardware, firmware, and software updates.

The first-generation IBM ESS models were available in the following types:

- GSx models
- GLx models

Depending on the first-generation IBM ESS model, the following configurations were available:

- 2U24 or storage enclosures hold SSDs or 10K RPM HDD storage drives.
- 4U60 of storage enclosures to hold high capacity Nearline SAS HDD storage drives.

The first-generation IBM ESS has a solution IBM machine type: 5146. This machine type applies to the first-generation IBM ESS only. Each first-generation IBM ESS building block consists of two POWER8 processor-based servers and a specific number of storage enclosures (1, 2, 4, or 6), depending on the model.

The POWER8 processor-based data servers were model S822L (IBM machine type 8247-22L). The storage enclosures can be one of the following configurations:

- IBM EXP24S SFF (small form factor) Gen2-bay drawer (2U 24 drives), which are used in models GS1, GS2, GS4, or GS6.
- IBM DCS3700 (IBM machine type 1818-80E) is 4U and contain 58 HDD drives each, which is used in models GL2, GL4, and GL6.

The first-generation IBM ESS GS models can use 2.5-inch 10 K rpm HDD (1. 2 TB) or 2.5-inch SSD (400 GB or 800 GB).

The first-generation IBM ESS GL models can use 3.5-inch NL-SAS HDDs (2 TB, 4 TB, or 6 TB).

All IBM ESS models can be mixed in an IBM Spectrum Scale cluster to meet the specific needs of client applications that use IBM ESS storage. For example, a common best practice is to use SSDs for IBM Spectrum Scale metadata and NL-SAS for IBM Spectrum Scale data storage.

First-generation IBM Elastic Storage System GSx models

The GSx models that used 2U24 storage enclosures were available with all SSD or all 10K RPM, high-speed SAS HDD storage.

The first-generation IBM ESS GSx models are shown in Figure A-1.

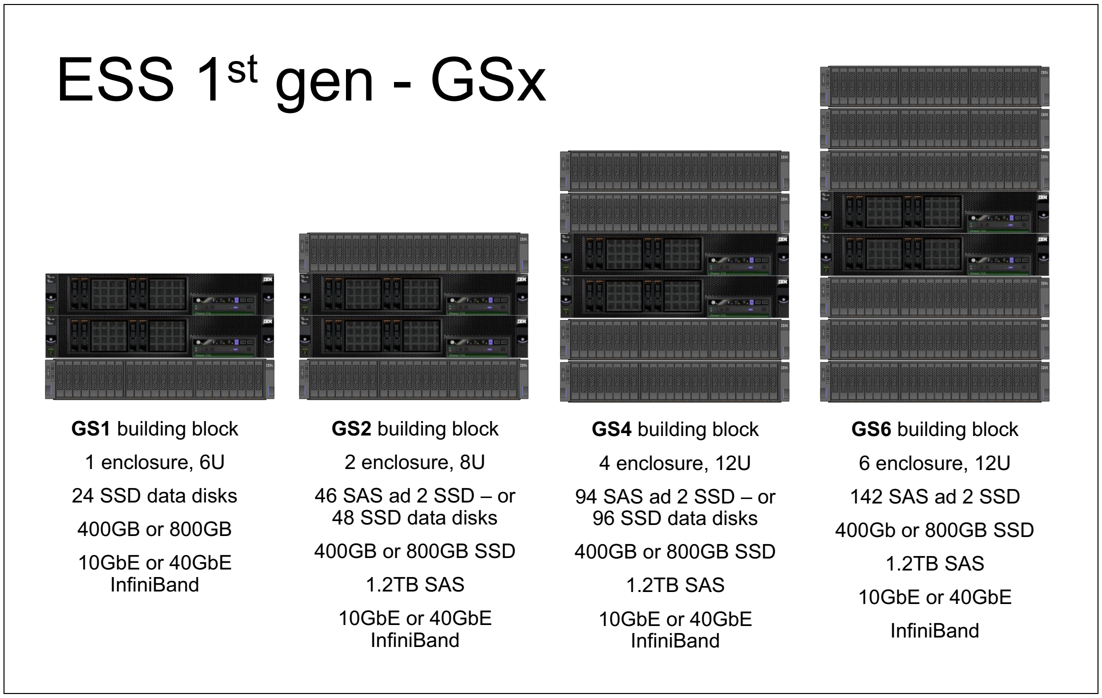

Figure A-1 First-generation IBM ESS GSx models

A primary use of the GSx models with SSD was for IBM Spectrum Scale metadata, or for high-speed HDD storage tier.

First-generation IBM Elastic Storage System GLx models

The HDD models of the first-generation IBM ESS (see Figure A-2) used 4U60 Storage enclosures (IBM 1818-80E).

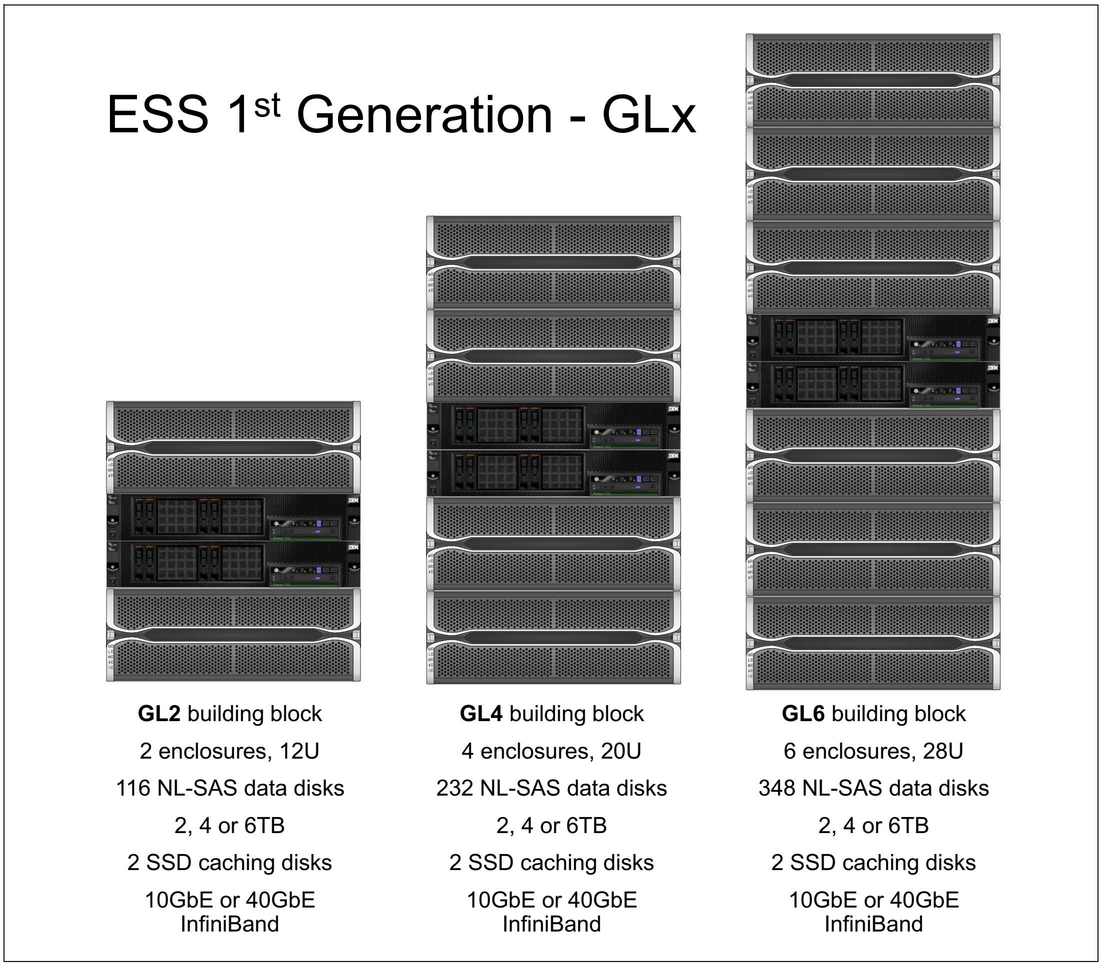

Figure A-2 First-generation IBM ESS GLx models

The primary usage of the GSx models with NL-SAS HDDs was for IBM Spectrum Scale capacity tiers.

First-generation IBM Elastic Storage System Data Servers

The first-generation IBM ESS uses a pair of POWER8 processor-based (machine type 8247-22L) servers as IBM Spectrum Scale Data servers, which are cross-connected to all storage for high availability (HA). These 8247-22L servers run Red Hat Enterprise Linux (RHEL) Big Endian operating system.

The first-generation IBM ESS Management Server (IBM EMS) is also an POWER8 processor-based (machine type 8247-21L) server. This 8247-21L server also runs RHEL Big Endian operating system.

The first-generation POWER processor-based IBM ESS servers were managed by an IBM Hardware Management Console (HMC), which is a hardware component that is commonly used to manage IBM Power server environments.

HMC for the first-generation IBM Elastic Storage System

The first-generation IBM ESS used an IBM HMC to manage the IBM Power server portion of the first-generation IBM Elastic Storage Server (IBM ESS) environment.

> **Note:** IBM ESS required an HMC for the first-generation only. In later IBM ESS generations, the function of the HMC was integrated into the IBM Electronic Service Agent (ESA) running on the second-generation IBM EMS.

First-generation IBM Elastic Storage System Storage Enclosures

In this section, we describe the storage enclosures that were available on the first-generation IBM ESS models.

2U24 (feature code on 8247-22L)

The first-generation IBM ESS 2U24 Storage Enclosures can be ordered with SSDs or HDDs.

These storage enclosures do not have a separate IBM machine type. They were feature codes of POWER8 processor-based 8247-22L Data Servers in the first-generation IBM ESS.

The IBM 2U24 Storage Enclosure is shown in Figure A-3.

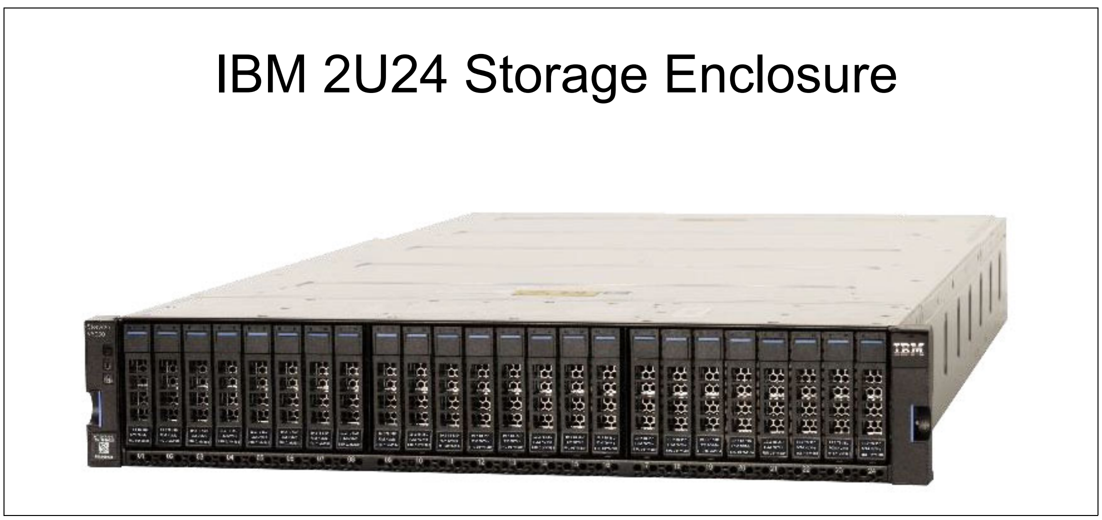

Figure A-3 IBM 2U24 Storage Enclosure

The following storage drives are available in these 2U24 enclosures:

- 2.5-inch 10 K rpm HDD (1.2 TB)
- 2.5-inch SSD (400 GB or 800 GB)

All 2U24 enclosures within one first-generation IBM ESS building block must have the same SSD or HDD size. Different IBM ESS building blocks can have different storage enclosures or drive sizes.

4U60 (1818-80E)

The first-generation IBM ESS HDD models (GL2, GL4, and GL6) used IBM machine type 1818-80E Storage Enclosures for HDD storage. All 1818-80E Storage Enclosures are 4U60 in size (4 rack units, 60 HDD slots available) and were populated with 58 HDDs (no partial population was available).

The first 1818-80E in a first-generation IBM ESS also contained two SSDs that hold internal IBM Spectrum Scale RAID metadata. The capacity of these two SSDs is not visible nor available as user capacity for the IBM Spectrum Scale and IBM ESS file system. The IBM 1818-80 4U60 Storage Enclosure is shown in Figure A-4.

Figure A-4 IBM 1818-80E 4U60 Storage Enclosure

Each IBM 1818-80E used in the first-generation IBM ESS, which held 58 HDDs.

All HDDs in a first-generation IBM ESS building block with 4U60 must have same HDD size. In the 1818-80E, 2 TB, 4 TB, and 6 TB HDDs were available.

Second-generation IBM Elastic Storage System overview

The second generation of IBM ESS was announced in April 2017 and was withdrawn from marketing in December 2021. IBM no longer provides upgrades for second-generation IBM ESS. At the time of writing, IBM continues to provide full hardware and software service and support for second-generation IBM ESS.

The second-generation IBM ESS does not include a solution IBM machine type. The IBM ESS server and storage hardware components feature the following IBM ESS-unique machine types:

- IBM machine type 5148 for the IBM ESS POWER8 processor-based servers
- IBM machine type 5147 for the IBM ESS IBM Storage Enclosures

These machine types uniquely identify these hardware components as part of an IBM ESS solution.

The available models of second-generation IBM ESS were available in the following categories:

- SSD flash storage IBM ESS models
- HDD IBM ESS models
- Hybrid IBM ESS models that consist of enclosures of SSD flash drives and HDDs.

IBM Elastic Storage System GSxS models

The second-generation IBM ESS GSxS models are all-SSD models. GSxS models are *not* available with HDD storage.

Figure A-5 shows the three IBM ESS GSxS models. A primary usage of the GSxS models with SSD is for IBM Spectrum Scale metadata, or to provide a high-speed SSD storage tier.

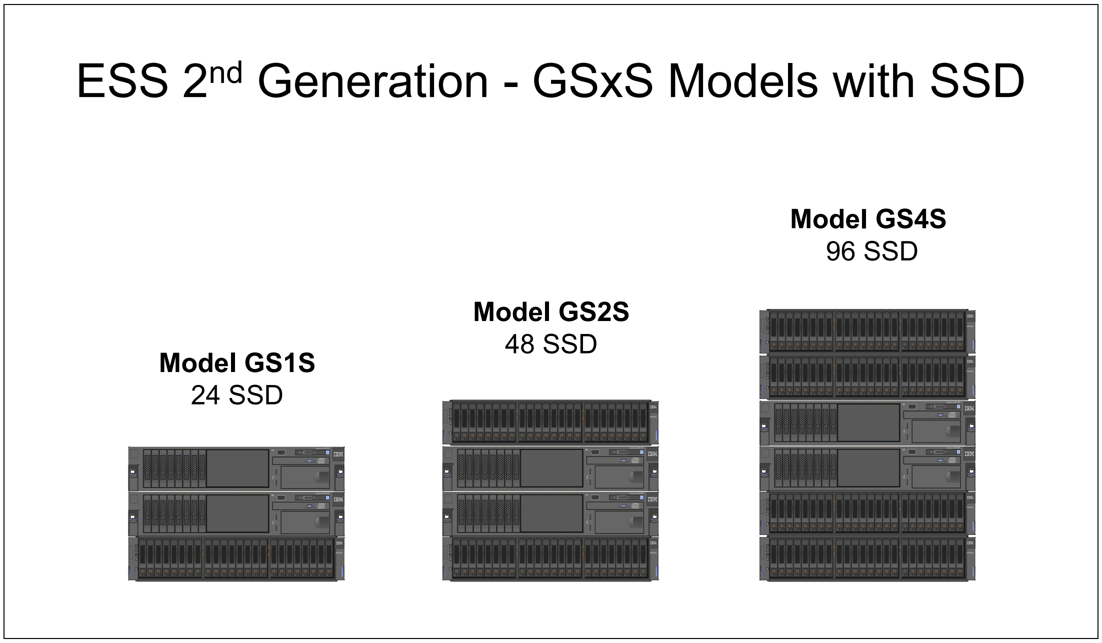

Figure A-5 Second-generation IBM ESS GSxS models

Note: IBM ESS performance data is available on request to your IBM or IBM Business Partner representative. They use the IBM File and Object Solution Design Engine to estimate performance that is based on your workload and network environment.

Optimum IBM ESS performance is derived from an unconstrained IOR benchmark for 100% sequential read numbers by using unconstrained InfiniBand networks. Other networks (such as 100 Gb Ethernet (GbE), 40 GbE, and 10 GbE) have more overhead than InfiniBand and typically lower aggregate bandwidth capabilities results. For more information, contact your IBM or IBM Business Partner representative.

IBM Elastic Storage System GHxy Hybrid models with SSD and HDD

The IBM ESS GHxy Hybrid models combined 1 - 2 drawers of SSDs with 2 - 4 drawers of HDD storage. A primary usage of the GHxy hybrid models is for a less expensive option to combine an IBM Spectrum Scale SSD storage tier with an IBM Spectrum Scale HDD capacity storage tier.

Second-generation IBM Elastic Storage System Hybrid models

IBM ESS Hybrid models are IBM ESS storage building blocks that consist of SSD flash 5147-024 Storage Enclosures and HDD 5147-084 Storage Enclosures. The IBM ESS Hybrid models (model GHxy) are a blend of 2U24 SSD enclosures and 5U84 HDD enclosures; “x” is the number of 2U24 SSD enclosures, and “y” is the number of 5U84 HDD enclosures.

The following models are available:

- GH12
- GH22
- GH14
- GH24

Figure A-6 shows the IBM ESS GHxy models.

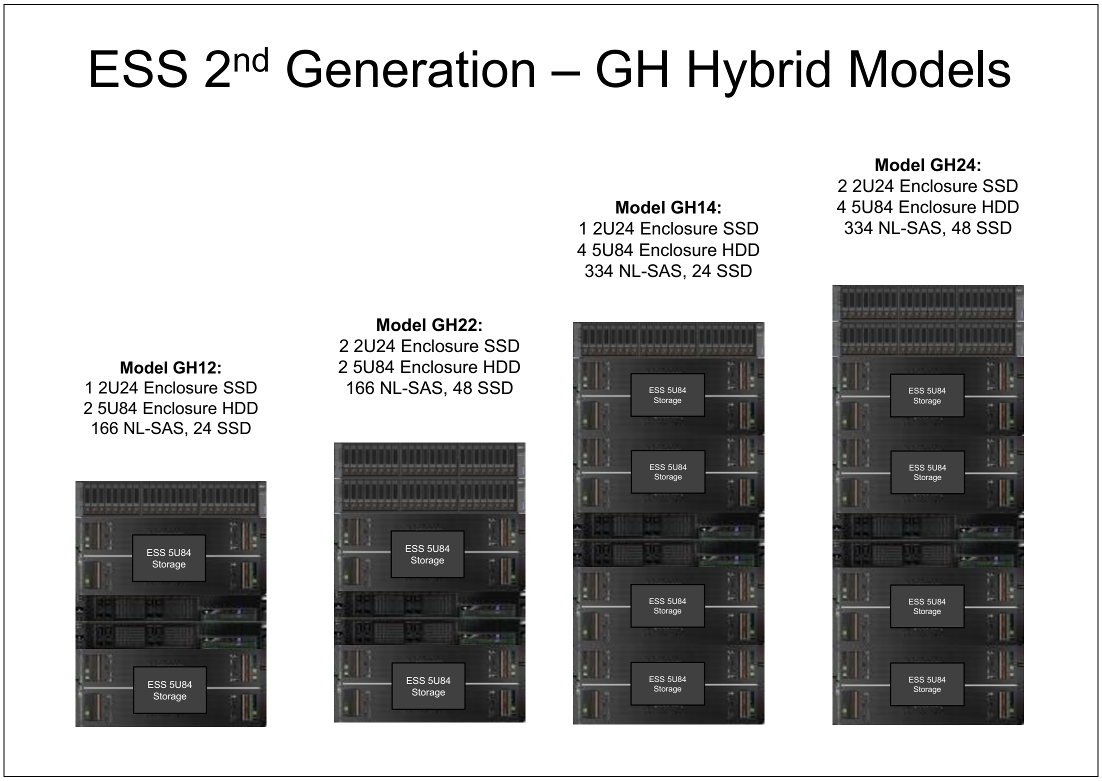

Figure A-6 Second-generation IBM ESS GHxy models

> **Note:** IBM ESS performance data is available on request to your IBM or IBM Business Partner representative. Use the IBM File and Object Solution Design Engine to estimate performance that is based on your workload and network environment.
>
> Optimum IBM ESS performance is derived from an unconstrained IOR benchmark for 100% sequential read numbers by using unconstrained InfiniBand networks. Other networks (such as 100 GbE, 40 GbE, and 10 GbE) have more overhead than InfiniBand and typically lower aggregate bandwidth capabilities results. For more information, contact your IBM or IBM Business Partner representative.

IBM Elastic Storage System GLxS models

The IBM ESS GLxS models provided high capacity HDD storage by using a 5U84 storage enclosure. A primary use of the GLxS models is to provide a high capacity IBM Spectrum Scale HDD storage tier.

Figure A-7 shows the various models that are available.

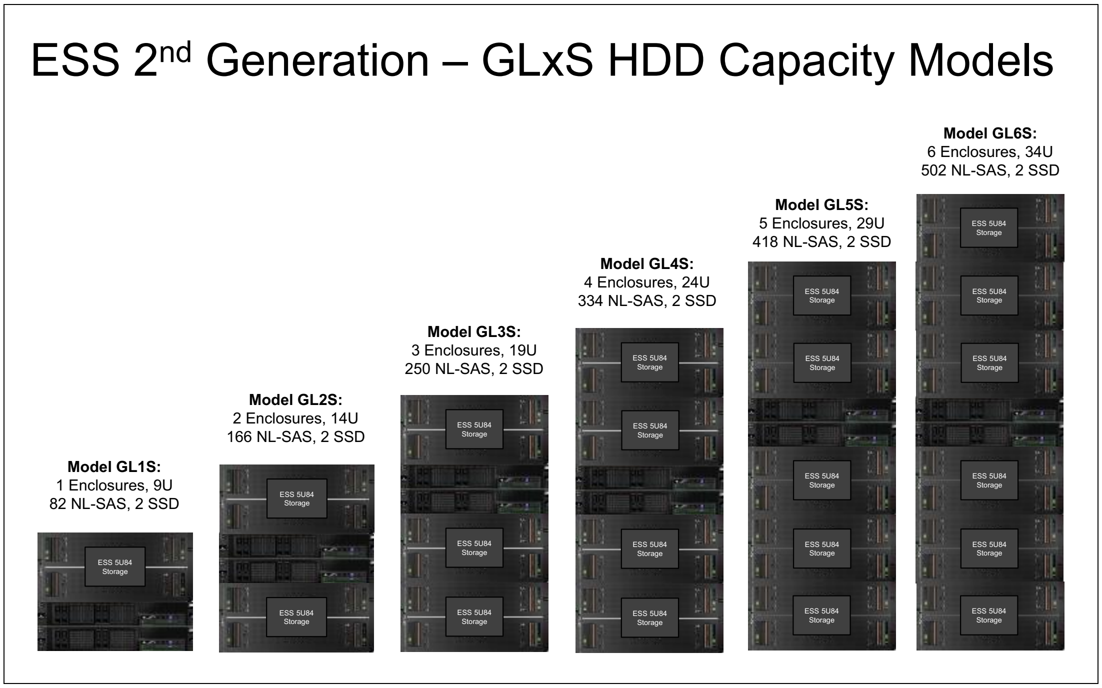

Figure A-7 Second-generation IBM ESS GLxS models

Note: IBM ESS performance data is available on request from your IBM or IBM Business Partner representative. They use the IBM File and Object Solution Design Engine to estimate performance that is based on your workload and network environment.

Optimum IBM ESS performance is derived from an unconstrained IOR benchmark for 100% sequential read numbers by using unconstrained InfiniBand networks. Other networks (such as 100 GbE, 40 GbE, and 10 GbE) have more overhead than InfiniBand and typically lower aggregate bandwidth capabilities results. For more information, contact your IBM or IBM Business Partner representative.

IBM Elastic Storage System GLxC models

IBM ESS GLxC models use the 4U106 storage enclosure that was pioneered and designed for use in the Sierra and Summit supercomputers at the US Department of Energy[1].

[1] https://www.ibm.com/thought-leadership/summit-supercomputer

Figure A-8 shows the IBM ESS GLxC models.

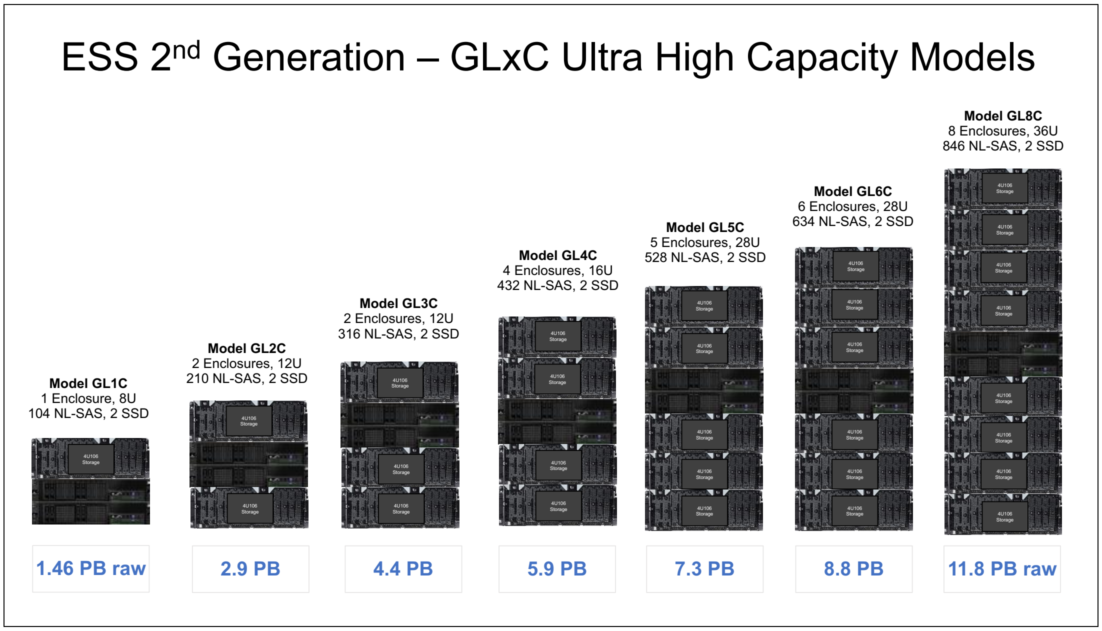

Figure A-8 Second-generation IBM ESS GLxC models

A primary use of the GLxC models is to provide an ultra high-capacity, high-density IBM Spectrum Scale HDD storage tier.

You can expect performance of the IBM ESS GLxC models to be similar to the corresponding IBM ESS GLxS model.

If you want to upgrade an existing IBM ESS GLxC model to an equivalent IBM ESS 5000 SC model, submit a Request for Price Quotation (RPQ) to IBM to replace the POWER8 processor-based data servers with POWER9 processor-based data servers. The upgrade is a destructive upgrade. Data must be backed up from the IBM ESS GLxC; the upgrade and replacement of the POWER8 processor-based data servers with POWER9 processor-based data servers is performed; the IBM ESS software stack must be reloaded; and the IBM Spectrum Scale file system is redefined on the upgraded IBM ESS 5000 SC model.

For more information about comparison tables with IBM ESS model specifications for SSD and HDD drive sizes and counts by model, see Appendix A, “IBM Elastic Storage System models” on page 65.

IBM Elastic Storage System 3200

The IBM ESS 3200 is a third-generation IBM ESS model. IBM ESS 3200 was announced in May 2021 and includes the following features:

- Based on a 2U24 storage enclosure with PCI Gen4-based x86 server canisters to provide improvements in throughput and bandwidth capability.
- Uses enterprise class NVMe drives:
  - Can be ordered half-populated with 12 NVMe drives or fully populated with 24 NVMe drives.
  - If ordered half-populated, you can add the remaining 12 NVMe drives non-disruptively (the added drives must be the same size as the first 12 drives).
- Designed to provide the following benefits:
  - High performance: NVMe flash storage with up to 80 GBps read throughput per 2U building block.
  - Designed to provide edge capability and global data access: This solution can be deployed in data centers or at the edge, incorporating and processing data that then uses IBM Spectrum Scale Active File Management (AFM) to share the data globally.
  - Simplicity: Containerized software installation and upgrade, plus a powerful management GUI, minimize the demands on IT staff time and expertise.
- Deployed by using containerized Red Hat Ansible playbooks that provide improved ease of use and orchestration of complex IBM ESS administration tasks, such as cluster configuration, file system creation, and code update.

The third-generation IBM ESS-3200 addresses the challenges of managing today's data. IBM ESS 3200 delivers high-performance, software-defined flash storage. The IBM ESS 3200 builds on years of experience and combines IBM Spectrum Scale software with fast NVMe storage technology to offer industry-leading file management capabilities.

The IBM ESS 3200 builds on and extends a track record of meeting the needs of the smartest and most demanding organizations. The IBM ESS 3200 is up to 100% faster than the previous generation of IBM ESS NVMe storage.

Figure A-9 shows the IBM ESS 3200 NVMe storage building block.

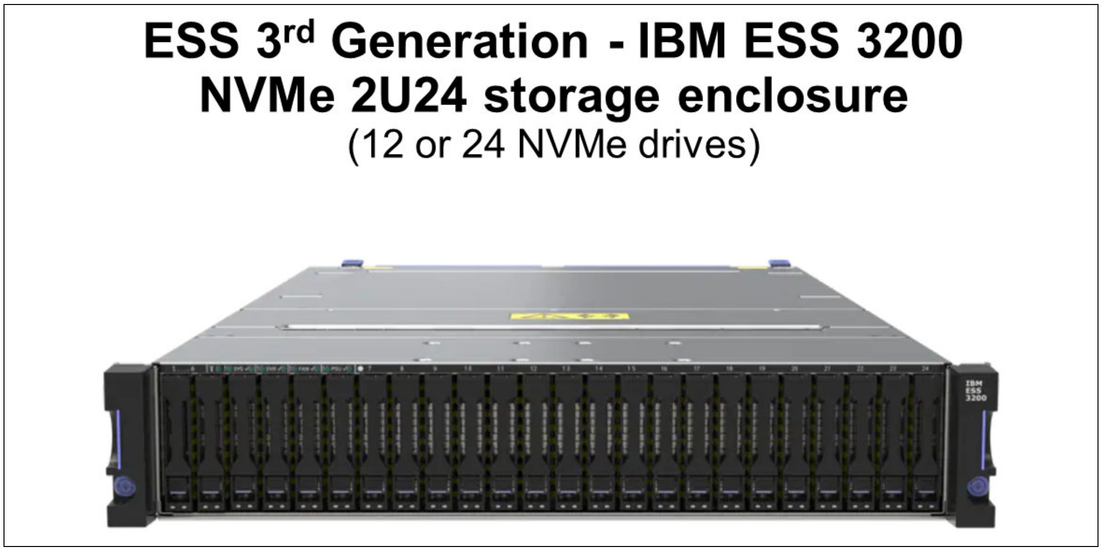

Figure A-9 Third-generation IBM ESS 3200 model

IBM ESS 3200 provides a high-performance tier of IBM Spectrum Scale file storage for various AI, analytics, and big data applications. IBM ESS 3200 keeps GPUs in AI workloads running at peak performance.

Like all IBM ESS models, IBM ESS runs IBM Spectrum Scale RAID erasure coding, which provides superior consistent high performance; mitigation of storage hardware failures; and intelligent monitoring, management, and dynamic tuning of all IBM ESS offerings for IBM Spectrum Scale data.

> **Note:** IBM ESS performance data is available on request from IBM or an IBM Business Partner representative. The IBM File and Object Solution Design Engine is used to estimate performance that is based on your workload and network environment.
>
> Optimum IBM ESS performance is derived from an unconstrained IOR benchmark for 100% sequential read numbers by using unconstrained InfiniBand networks. Other networks (such as 200 GbE, 100 GbE, 40 GbE, and 10 GbE) feature more overhead than InfiniBand and typically lower aggregate bandwidth capabilities results.
>
> For more information, contact your IBM or IBM Business Partner representative.

IBM ESS 3200 is the simplest way yet for users to deploy IBM Spectrum Scale. IBM Spectrum Scale is included in a pre-configured system. Installations and updates are delivered by using containerized software that speeds and simplifies the process.

A storage specialist from IBM System Lab Services implementation is not required if you have an IBM ESS and IBM Spectrum Scale system and you are comfortable with IBM ESS implementations. It is much easier to install, and maintenance can be performed by your IT staff.

If you are unfamiliar with IBM Spectrum Scale and IBM ESS, as a best practice, use IBM System Lab Services to ensure high satisfaction with your initial IBM ESS 3200 installation.

IBM Elastic Storage System 3000

The IBM ESS 3000 model was the initial model of the third-generation IBM ESS. IBM ESS 3000 was announced in October 2019 and includes the following features:

- Based on IBM FlashSystem® 9150 Storage Enclosure and x86 server canister technology.
- Uses enterprise class NVMe drives:
  - Can be ordered half-populated with 12 NVMe drives or fully populated with 24 NVMe drives.
  - If ordered half-populated, you can add the remaining 12 NVMe drives non-disruptively (the additional drives must be of the same size as the first 12 drives).
- Designed to provide the following benefits:
  - High-performance NVMe storage for High-Performance Tier (HPT) for IBM Spectrum Scale in a compact, energy-efficient 2U24 form factor with two highly available (HA), cross-connected, and x86-based IBM Spectrum Scale Data Servers.
  - The necessary bandwidth to fully use compute GPUs for AI and analytics workloads, which provide up to 40 GBps of throughput in a 2U24 form factor.
  - Improved ease of use, faster time to installation, and faster time to value.
- Deployed by using containerized Red Hat Ansible playbooks that provide improved ease of use and orchestration of complex IBM ESS administration tasks, such as cluster configuration, file system creation, and code update.

The third-generation IBM ESS 3000 model provides high-performance NVMe storage. Figure A-10 shows the IBM ESS 3000 NVMe storage building block.

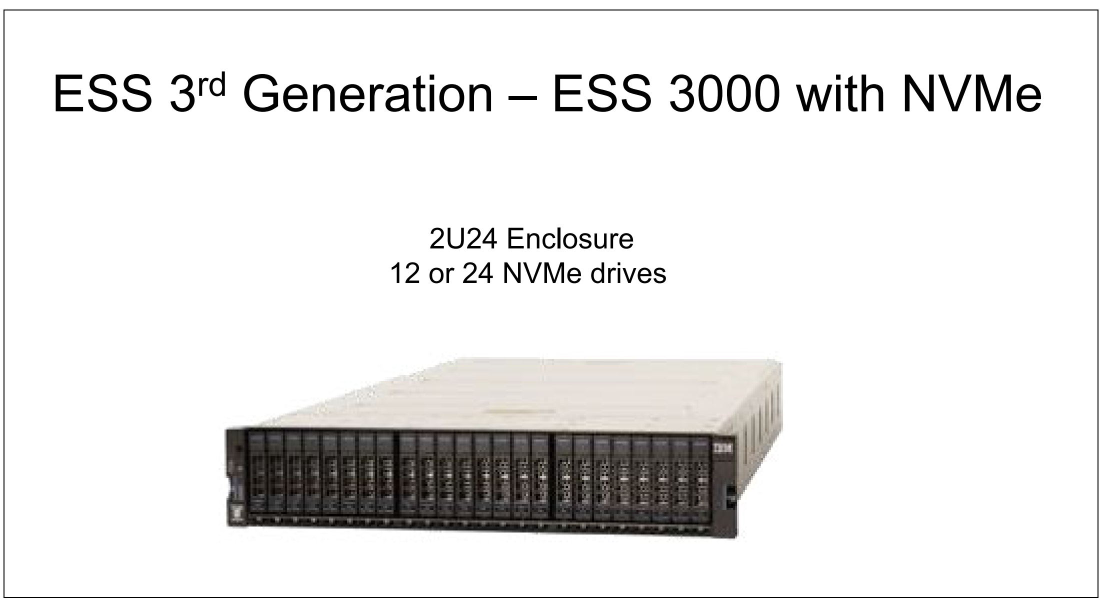

Figure A-10 Third-generation IBM ESS 3000 model

A primary use of the IBM ESS 3000 model with NVMe is for fast IBM Spectrum Scale metadata, or to provide a dense, high-speed, and potentially mobile NVMe storage tier.

Note: IBM ESS performance data is available on request from your IBM or IBM Business Partner representative. They use the IBM File and Object Solution Design Engine to estimate performance that is based on your workload and network environment.

Optimum IBM ESS performance is derived from an unconstrained IOR benchmark for 100% sequential read numbers by using unconstrained InfiniBand networks. Other networks (such as 200 GbE, 100 GbE, 40 GbE, and 10 GbE) have more overhead than InfiniBand and typically lower aggregate bandwidth capabilities results.

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

For ease of management and upgrading, IBM ESS 3000 provides a simplified, faster installation experience compared to previous-generation IBM ESS models. IBM ESS 3000 hardware is set up by an IBM System Services Representative (IBM SSR). If the client is unfamiliar with IBM Spectrum Scale and IBM ESS, as a best practice, use IBM System Lab Services to ensure high satisfaction with your initial IBM ESS 3000 installation.

Differences between the IBM Elastic Storage System generations

The following section describes the component differences between the first-, second-, and third-generation IBM ESS models. It is helpful if these different generations of IBM ESS are in the same IBM Spectrum Scale cluster.

Hardware

The IBM ESS generations feature the following major hardware differences:

- First generation:
  - Uses POWER8 processor-based 8247-21L for IBM EMS and POWER8 processor-based 8247-22L for the IBM ESS Data Servers, all running RHEL Big Endian. This generation IBM ESS has an overall solution machine type 5146.
  - The first-generation IBM ESS for 2U24 storage used the IBM EXP24 2U24 Storage Enclosure, which is specified as feature codes on the 8247-22L Data Server. The HDD storage used IBM 1818-80E Storage Enclosures for large capacity.
  - It required an IBM Power HMC to manage the IBM Power servers.
  - This generation requires an RHEL Big Endian-based IBM EMS (8247-21L). One RHEL Big Endian-based IBM EMS is required per IBM Spectrum Scale cluster that includes first-generation IBM ESS.
- Second generation:
  - Uses POWER8 processor-based 5148-21L for IBM EMS and POWER8 processor-based 5148-22L for the IBM ESS Data Servers, all running RHEL Little Endian. These machine types are unique to the IBM ESS and uniquely identify these POWER8 processor-based servers as part of an IBM ESS solution. The second-generation IBM ESS does not have an overall solution machine type.
  - This generation uses IBM 5147-024 Storage Enclosures for SSD storage. For high capacity HDD storage, it uses IBM 5147-084 Storage Enclosures or 5147-106 Storage Enclosures. These machine types are unique to IBM ESS and uniquely identify these storage enclosures as part of an IBM ESS solution.

- Third generation:
  - The IBM ESS 3000 has machine type 5141-AF8, and can be managed by a 5148-21L POWER8 processor-based IBM EMS or a 5105-22E POWER9 processor-based IBM EMS.
  - The IBM ESS 3200 has machine type 5141-FN1, and can be managed a 5105-22E POWER9 processor-based IBM EMS.
  - The IBM ESS 5000 has a common server IBM machine type of 5105-22E POWER9 processor-based servers for Data Servers, protocol nodes, and IBM EMS.
  - An IBM ESS 5000 storage building block can use 5147-092 Storage Enclosures (SL model) or 5147-106 Storage Enclosures (SC model). You cannot intermix 5147-092 and 5147-106 in the same IBM ESS 5000 storage building block.
  - IBM ESS 5000 must be managed by a 5105-22E POWER9 processor-based IBM EMS. The second-generation 5148-21L POWER8 processor-based management server cannot support the IBM ESS 5000 or the IBM ESS 3200.

Consider the following points:

- Only the first-generation IBM ESS models require an HMC for the IBM Power servers. In the second- and third-generation IBM ESS models, the HMC function was integrated into the IBM EMS ESA.
- IBM fully supports all three multiple generations of IBM ESS in the same IBM Spectrum Scale cluster. If doing so, make sure you satisfy the requirement that the appropriate IBM EMS generation also is attached to the IBM Spectrum Scale cluster.

> **Note:** All IBM ESS generations can all coexist in a IBM Spectrum Scale cluster.

A typical first-generation IBM ESS deployment is shown in Figure A-11.

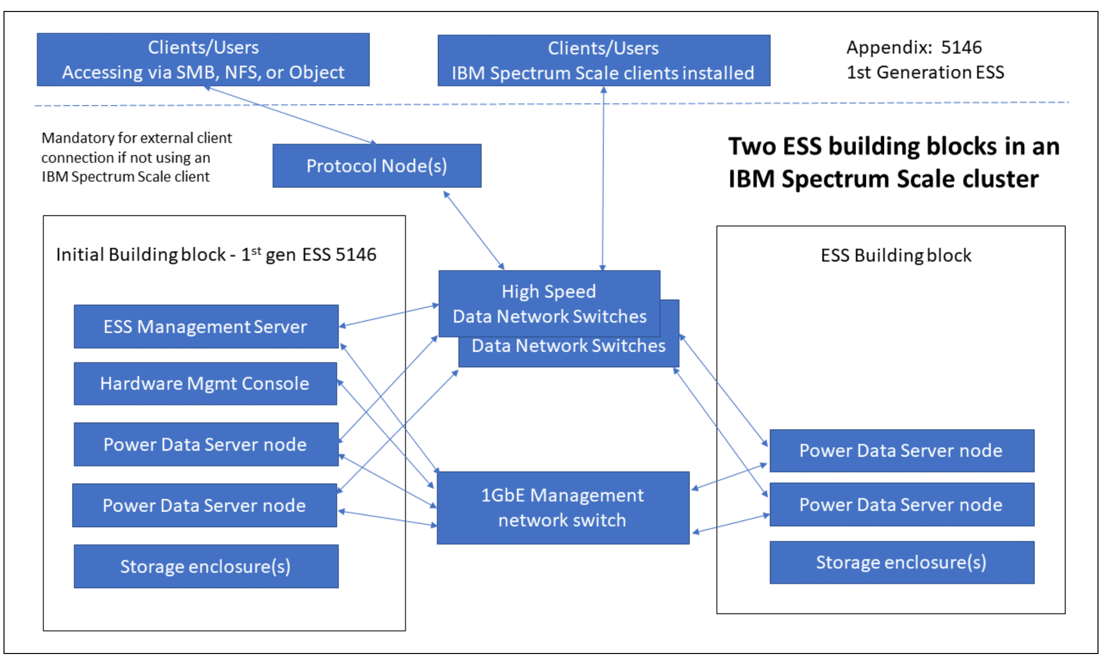

Figure A-11 Typical first-generation IBM ESS components and deployment

A typical IBM ESS second-generation deployment is shown in Figure A-12.

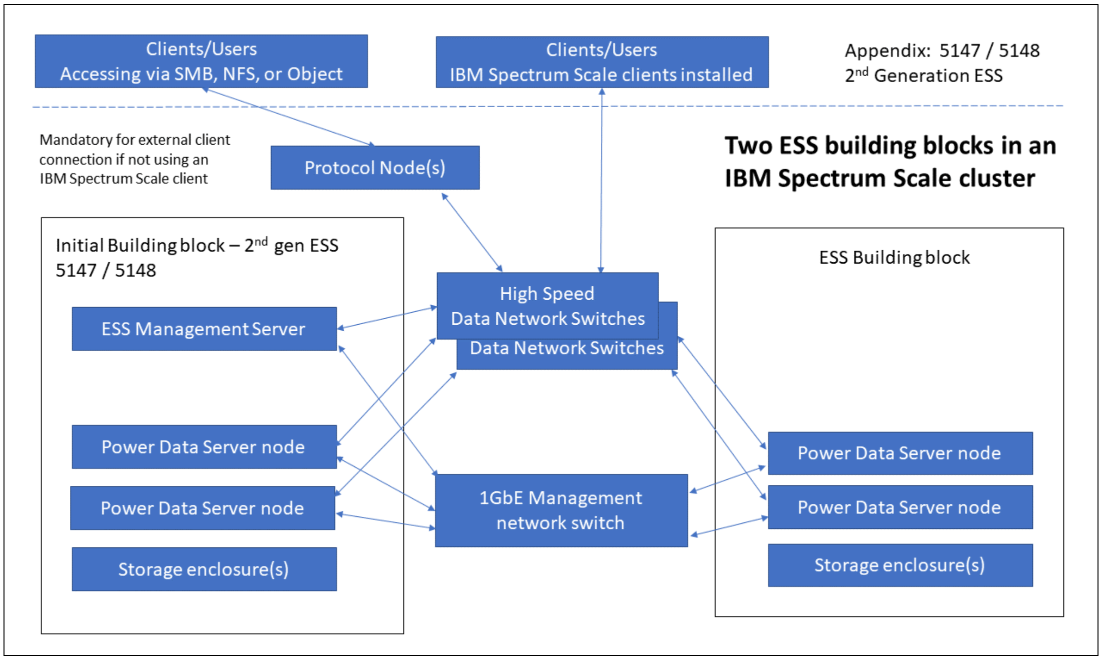

Figure A-12 IBM ESS second-generation components and deployment

The components and deployments of IBM ESS first generation is compared to IBM ESS second generation in Figure A-13. Consider the following major differences that are featured in the second generation:

- Does not have a solution machine type that is equivalent to “5146”. Instead, the second-generation IBM ESS components are uniquely identified as machine type 5148 servers or machine type 5147 Storage Enclosures.
- Runs RHEL Little Endian.
- Does not require an HMC for the IBM Power servers.

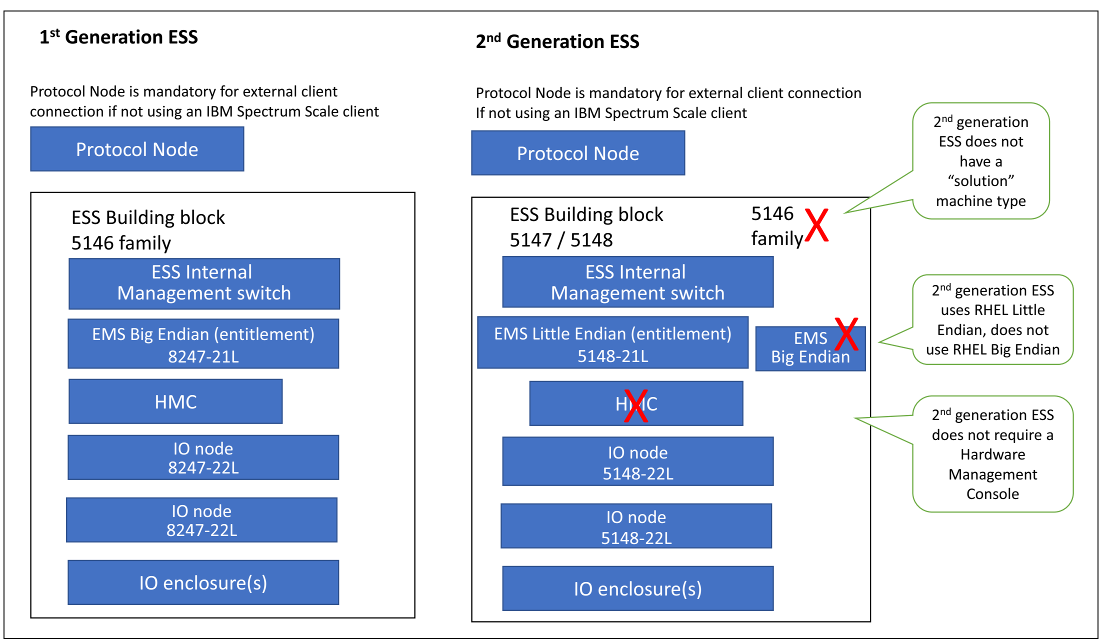

Figure A-13 Differences between IBM ESS first- and second-generation components

IBM third-generation IBM ESS 5000 models are similar in configuration to second-generation IBM ESS models, with the following differences or similarities:

- IBM 5148 POWER8 processor-based servers are replaced with IBM 5105-22E POWER9 processor-based servers.
- IBM ESS 5000 SL models use IBM 5147-092 Storage Enclosures, whereas the second-generation IBM ESS GLxS models used IBM 5147-084 Storage Enclosures.
- IBM ESS 5000 SC models use the same IBM 5147-106 Storage Enclosures, as the second-generation IBM ESS GLxC models.
- No third-generation IBM ESS Hybrid model exists that is equivalent to the second-generation IBM ESS GHxx models. Instead, use combinations of IBM ESS 3200, IBM ESS 5000, or IBM ESS 3000 as suitable.

IBM fully supports first-, second-, and third-generation IBM ESS models co-existing in an IBM Spectrum Scale cluster.

Software

The IBM ESS generations feature the following major software differences and nomenclature:

- First-generation IBM ESS runs RHEL Big Endian.
- Second- and third-generation IBM ESS runs RHEL Little Endian.
- IBM ESS software stacks are under the nomenclature *IBM ESS 5.x.x*.
- IBM ESS software stacks with the nomenclature *IBM ESS 5.2.x* have (now end of service) IBM Spectrum Scale 4.2.x.
- IBM ESS software stacks with the nomenclature *IBM ESS 5.3.x* have IBM Spectrum Scale 5.0.x.
- IBM ESS software stacks with the nomenclature *IBM ESS 6.x.x* have IBM Spectrum Scale 5.1.x.
- The various IBM ESS software stack component levels for IBM ESS software stacks are documented in the IBM Spectrum Scale RAID FAQ in the Online IBM Documentation.

Note: An IBM Spectrum Scale cluster can coexist nodes that run Big Endian and Little Endian operating systems.

IBM Elastic Storage System Data Server IO node

The IBM Spectrum Scale Data Server machine type and model for first-generation IBM ESS is IBM machine type 8247-22L. This machine type and model is the same as the equivalent POWER8 processor-based compute servers.

The second-generation IBM ESS uses IBM machine type 5148-22L for the IBM ESS Data Servers. This machine type is unique to IBM ESS and is a different machine type from the stand-alone POWER8 processor-based compute servers.

Two identical data servers always must be used in each IBM ESS building block.

The third-generation IBM ESS 3000 and IBM ESS 3200 Data Servers are based on x86 servers. They are integrated in a single 2U24 enclosure system with the NVMe storage that also is within the 2U24 IO enclosures.

The third-generation IBM ESS 5000 uses IBM machine type 5105-22E for the IBM ESS Data Servers. This POWER9 processor-based server is unique to IBM ESS.

IBM Elastic Storage System Management Server

The EMS machine type and model for first generation is IBM machine type 8247-21L. This machine type and model is the same as the equivalent POWER8 processor-based compute servers.

The second-generation IBM ESS uses IBM machine type 5148-21L for the IBM EMS. This IBM machine type is unique to IBM ESS to differentiate it from standard 8247-21L POWER8 processor-based compute servers.

An IBM Spectrum Scale cluster with IBM ESS must have at least one IBM EMS for each “Endianness”.

If you have first-generation IBM ESS in an IBM Spectrum Scale cluster, you must have a Big Endian IBM EMS. This Big Endian EMS 8247-21L supports the first-generation IBM ESS systems only.

If you have second- and third-generation IBM ESS systems in an IBM Spectrum Scale cluster, you must have at least one Little Endian IBM EMS. Use the suitable Little Endian EMS to support your second- and third-generation IBM ESS systems.

Consider the following points”

- The IBM ESS 3000 has machine type 5141-AF8 and can be managed by a 5148-21L POWER8 processor-based IBM EMS or a 5105-22E POWER9 processor-based IBM EMS.
- The IBM ESS 5000 has a common server IBM machine type of 5105-22E POWER9 processor-based servers for data servers, protocol nodes, and IBM EMS.
- An IBM ESS 5000 storage building block can use 5147-092 Storage Enclosures (SL model) or 5147-106 Storage Enclosures (SC model). You cannot intermix 5147-092 and 5147-106 in the same IBM ESS 5000 storage building block.

IBM ESS 5000 must be managed by a 5105-22E POWER9 processor-based IBM EMS. The second-generation 5148-21L POWER8 processor-based management server cannot support the IBM ESS 5000 or the IBM ESS 3200.

Hardware Management Console

An HMC is used by the first-generation IBM ESS systems only. For the second- and third-generation IBM ESS, the HMC features are integrated into the Little Endian IBM EMS.

IBM Spectrum Scale for IBM Elastic Storage System licensing

IBM Spectrum Scale and IBM ESS licensing is described in detail in the IBM Spectrum Scale FAQ Online at IBM Documentation, in Section 13 “Licensing and Pricing Questions” of IBM Spectrum Scale Overview.

For more information about IBM Spectrum Scale and IBM ESS Licensing, see *IBM Spectrum Scale: IBM Elastic Storage System Licensing Information*.

Capacity licensing on second- and third-generation IBM Elastic Storage System

All second- and third-generation IBM ESS models are licensed by capacity. Instead of the standard IBM Spectrum Scale software license with a “Per TiB” metric, a specific IBM Spectrum Scale for IBM ESS license with a specific IBM Program ID and with a “Per Disk” metric is used.

The “Per Disk” metric licenses the IBM Spectrum Scale for IBM ESS by the number of SSDs or HDDs in the IBM ESS. No difference exists in the license charge for the size of the SSD or HDD. Therefore, the licensing is the same for an IBM ESS model, regardless if it has small 4 TB HDDs or large 14 TB HDDs.

Note: IBM Spectrum Scale for IBM ESS (current IBM Program IDs are 5765-DAE for Data Access Edition and 5765-DME for Data Management Edition) is a specific IBM Program ID for use on IBM ESS. It uses a different metric than the IBM Spectrum Scale software-only capacity license (5641-DAx or 5641-DMx), which is licensed by TiB. (The “x” indicates 1, 3, or 5 years of IBM software service and support.)

Also, unlike IBM Spectrum Scale software (which IBM clients can manage through IBM Passport Advantage®), IBM Spectrum Scale for IBM ESS does not have a Passport Advantage part number. Renewals with 5765-DAE and 5765-DME can be managed with by working with your IBM Technical Support Services (TSS) representative.

It is possible to license your IBM ESS by using IBM Spectrum Scale “Per TiB” capacity pricing. For more information about how to use this licensing model on your IBM ESS, see IBM Spectrum Scale Licensing question 2.2 and *IBM Spectrum Scale: IBM Elastic Storage System Licensing Information*.

IBM Spectrum Scale RAID technical overview

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

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

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

IBM Spectrum 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 Spectrum Scale RAID implements erasure coding, which is embedded within the IBM Spectrum 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 Spectrum Scale RAID also implements an end-to-end checksum from the storage media all the way out to the IBM Spectrum Scale client.

IBM Spectrum Scale RAID erasure codes

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

The options for RAID configuration are eight stripes of data plus two or three parity stripes that use Reed-Solomon codes or one stripe of data plus two or three replica stripes. The data plus parity or replica stripes, which are known as *tracks*, are shown in Figure A-14.

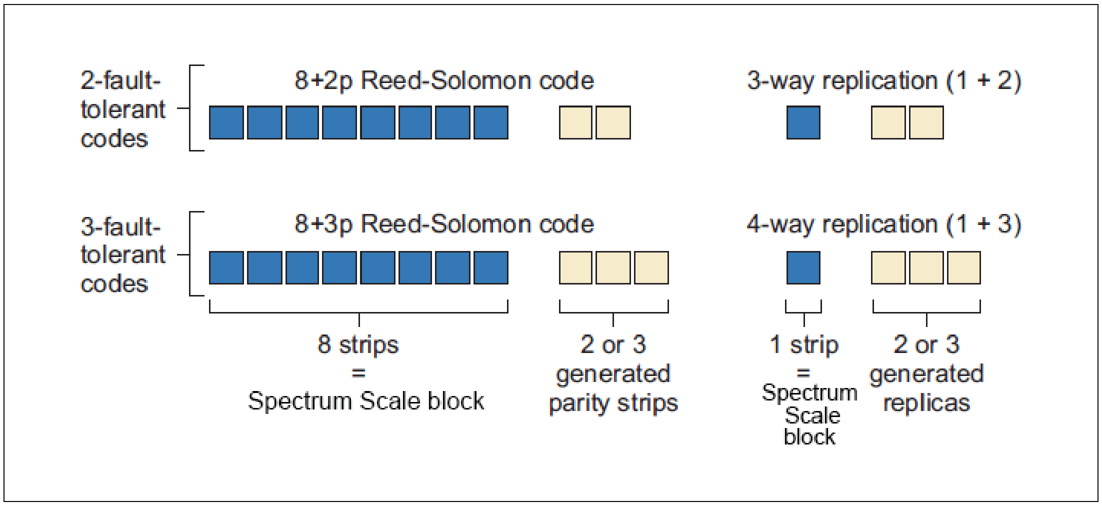

Figure A-14 IBM Spectrum Scale RAID tracks

End-to-end checksum

If you use the IBM Spectrum Scale software client to access data on the IBM ESS, the IBM Spectrum Scale client is aware if the IBM Spectrum Scale file system is based on IBM Spectrum 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 Spectrum Scale RAID server. The checksum is verified and then IBM Spectrum 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 Spectrum Scale RAID verifies checksum and version on its metadata. If it is OK, it sends the data to the client. If it is not OK, the data is rebuilt based on parity or replication and then sent to the client along with 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 Spectrum 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 Spectrum 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 Figure A-15.

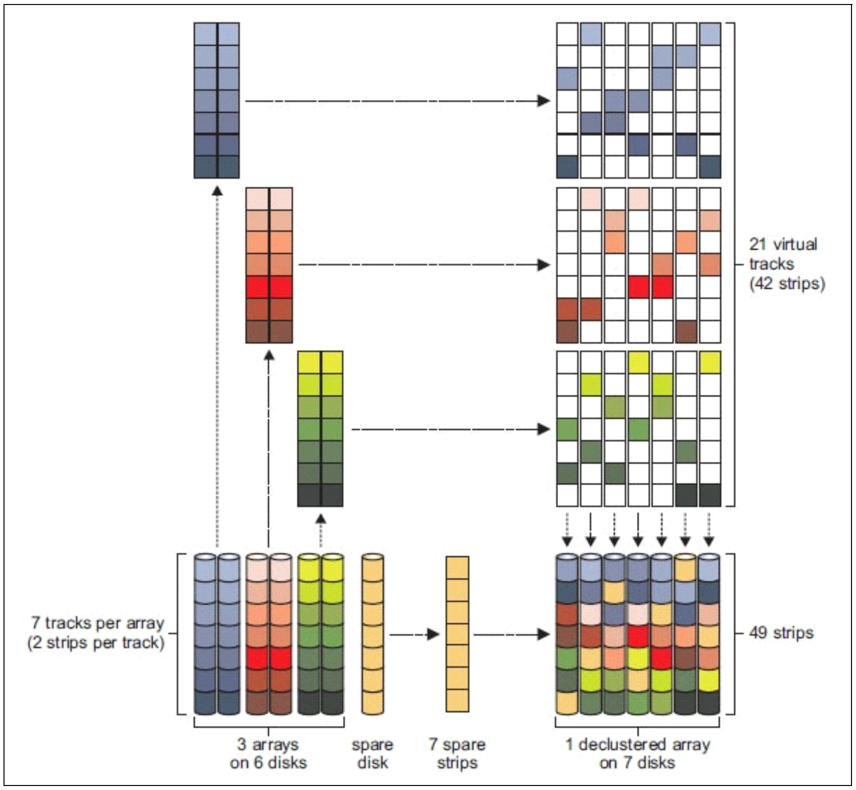

Figure A-15 Declustered array versus a 1+1 array

Consider an example where seven stripes of data are on each disk. Figure A-15 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 Spectrum Scale RAID overhead reduction is shown in Figure A-16.

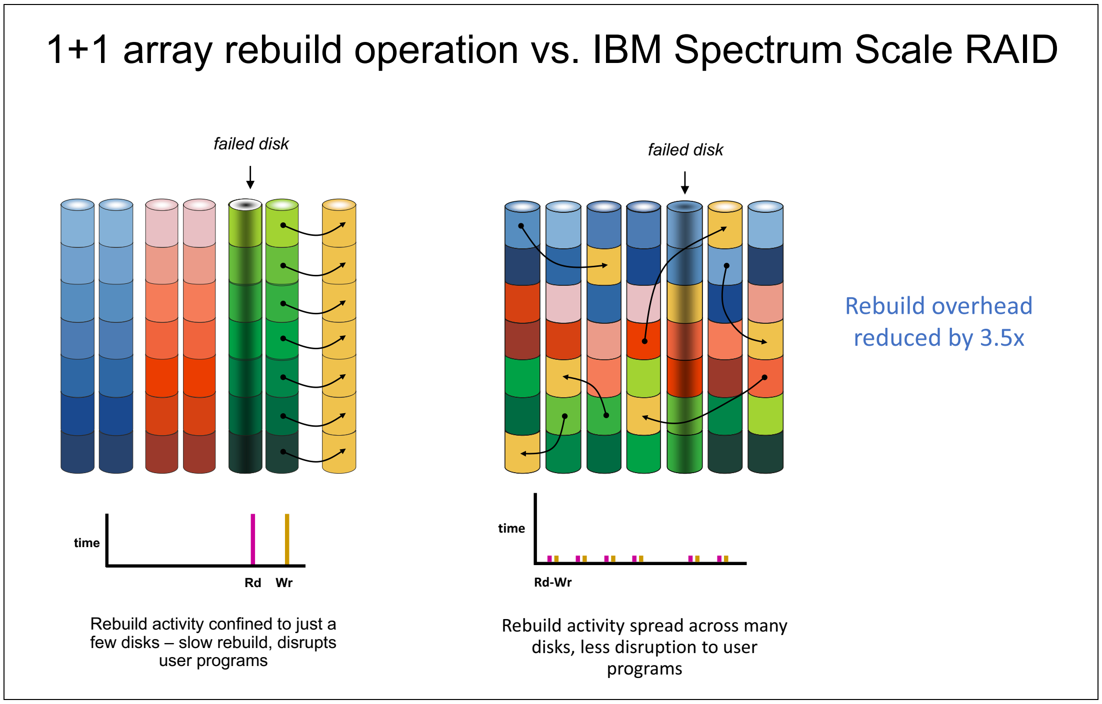

Figure A-16 Array rebuild operation

On the IBM ESS, 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 RAIDs or 3-way replication (1+3), the rebuild operation becomes critical only if three concurrent disk losses occur in the same declustered array.

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this paper.

IBM Redbook

Many IBM Redbooks publications and IBM Redpapers are available about IBM Spectrum Scale and IBM ESS. For more information, see the following search results on the IBM Redbooks publications website:

- IBM Elastic Storage Server
- IBM Elastic Storage System
- IBM Spectrum Scale

You can search for, view, download, or order documents and other Redbooks, Redpapers, web docs, drafts, and additional materials at the following website:

ibm.com/redbooks

Online resources

The following websites are also relevant as further information sources:

- IBM Documentation Frequently Asked Questions for IBM Elastic Storage Server:

  https://www.ibm.com/docs/en/ess-p8?topic=SSYSP8/gnrfaq.html

- IBM Documentation Frequently Asked Questions for IBM Spectrum Scale:

  https://www.ibm.com/docs/en/STXKQY/gpfsclustersfaq.html

- IBM Documentation for IBM Elastic Storage Server:

  https://www.ibm.com/docs/en/ess-p8

- IBM Documentation for IBM Spectrum Scale:

  https://www.ibm.com/docs/en/spectrum-scale

- IBM Education on IBM Elastic Storage Server or on IBM Spectrum Scale:

  https://www.ibm.com/training

  At this web page, search for:

  - Elastic Storage System
  - IBM Spectrum Scale

- IBM ESS classes from IBM Learning Services: *Introduction to IBM Elastic Storage Server* (no-charge course):

  https://www.ibm.com/services/learning/search?query=DL08017G

- IBM ESS IBM Spectrum Scale RAID Quick Deployment Guide:

  https://www.ibm.com/docs/en/ess-p8/6.1.4?topic=ess-deployment-guide

- IBM ESS manuals and publications in IBM Documentation:

  https://www.ibm.com/docs/en/ess-p8/6.1.4?topic=ess-pdfs

- IBM ESS Problem Determination Guide:

  https://www.ibm.com/docs/en/ess-p8/6.1.4?topic=ess-pdfs

- IBM ESS System Specifications:

  https://www.ibm.com/docs/en/power-sys-solutions/0008-ESS?topic=server-system-specifications-elastic-storage-5147-5148

- IBM Network Switches sourced from NVIDIA:

  - https://www.ibm.com/docs/en/power9?topic=networking-switches
  - https://www.ibm.com/docs/en/power8?topic=POWER8/p8hdx/p8hdx_network_switches.htm

- IBM Rack:
  https://www.ibm.com/docs/en/power9/9080-M9S?topic=hss-rack-installation-specifications-racks-that-are-not-purchased-from

- IBM Sales Manuals for IBM Elastic Storage Servers:

  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/3/897/ENUS5141-_h03/index.html
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/2/897/ENUS5141-_h02/index.html&lang=en&request_locale=en
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/897/ENUS5141-_h01/index.html&lang=en&request_locale=en
  - https://www.ibm.com/common/ssi/printableversion.wss?docURL=/common/ssi/rep_sm/1/897/ENUS5146-_h01/index.html&request_locale=en
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/897/ENUS5148-_h01/index.html&lang=en&request_locale=en
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/3/897/ENUS5147-_h03/index.html&lang=en&request_locale=en
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/2/897/ENUS5147-_h02/index.html&lang=en&request_locale=en
  - http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/4/897/ENUS5147-_h04/index.html&request_locale=en
  - https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/872/ENUS5141-_h01/index.html&lang=en&request_locale=en

- IBM Spectrum Scale Protocol Node Quick Deployment Guide as part of the IBM ESS solution:

  https://www.ibm.com/docs/en/ess-p8/6.1.4?topic=ess-pdfs

- Introducing the new IBM Elastic Storage Server 3500:

  https://www.ibm.com/support/pages/introducing-new-ibm-elastic-storage-server-3500

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Abbreviations and acronyms

AFM	Active File Management
APB	Ansible playbook bundles
ATG	Advanced Technology Group
CES	Cluster Export Services
CIFS	Common Internet File System
CTDB	Clustered Trivial Data Base
DHCP	Dynamic Host Configuration Protocol
EDR	Extended Data Rate
IBM EMS	IBM ESS Management Server
EoS	End of Service
ESA	Electronic Service Agent
IBM ESS	IBM Elastic Storage Server or IBM Elastic Storage System
EUS	Extended Update Support
FOS DE	IBM File Object Solution Design Engine
FSP	flexible service processor
GbE	Gb Ethernet
GPFS	General Parallel File System
HA	high availability or highly available
HBA	host bus adapter
HDD	hard disk drive
HDR	High Data Rate
HMC	Hardware Management Console
HPT	High Performance Tier
IBM	International Business Machines Corporation
ITSO	International Technical Support Organization
JBOD	just a bunch of disks
MES	miscellaneous equipment specification
MTM	machine type and model
NFS	Network File System
NIC	network interface card
NSD	Network Shared Disk
OFED	OpenFabrics Enterprise Distribution
PDU	power distribution unit
QDG	Quick Deployment Guide
RAS	reliability, availability, serviceability
RHEL	Red Hat Enterprise Linux
RoCE	RDMA over Converged Ethernet
RPQ	Request for Price Quotation
SAS	serial-attached SCSI
SATA	Serial Advanced Technology Attachment
SDS	software-defined storage
SMB	Server Message Block
SME	subject matter expert
SSD	solid-state drive
IBM SSR	System Services Representative
StorM	Storage Modeler
TDA	Technical and Delivery Assessment
TSS	IBM Technical Support Services
VPI	Virtual Protocol Interconnect

REDP-5253-04

ISBN 0738460931

Printed in U.S.A.