

IBM Storage FlashSystem 9500 Product Guide: Updated for IBM Storage Virtualize 8.7

Hartmut Lonzer

Jon Herd

Vasfi Gucer



Storage

Infrastructure Solutions



IBM Storage FlashSystem 9500 Product Guide

This IBM® Redpaper Product Guide describes the IBM Storage FlashSystem 9500 (IBM FlashSystem 9500) solution, which is a next-generation IBM FlashSystem® control enclosure. Leveraging IBM FlashCore® technology and Non-Volatile Memory Express (NVMe) architecture, it delivers high performance, reliability, and data protection. Ideal for AI, big data analytics, and cloud environments, the 9500 helps enterprises overcome resource constraints and maximize existing investments.

IBM FlashSystem 9500 provides a rich set of software-defined storage (SDS) features that are delivered by IBM Storage Virtualize, including the following examples:

- ▶ Data reduction and deduplication
- ▶ Dynamic tiering
- ▶ Thin-provisioning
- ▶ Snapshots
- ▶ Cloning
- ▶ Cyber resilience
- ▶ Transparent Cloud Tiering (TCT)
- ▶ Policy-based replication
- ▶ Policy-based high availability
- ▶ Scale-out and scale-up configurations that further enhance capacity and throughput for better availability
- ▶ Inline Data Corruption Detection

Overview

The IBM FlashSystem 9500 leverages IBM FlashCore technology with enhanced compression to boost capacity up to 116 TB per drive.

The IBM FlashSystem 9500 system with IBM FCMs includes built-in hardware data compression, and this data reduction is always on. This compression is implemented in hardware by using field-programmable gate arrays (FPGAs) within each module and uses a modified dynamic GZIP algorithm. With this approach, the solution can deliver the level of performance that you expect without compression, with the added benefit of better use of the physical storage.

IBM FCMs implement hardware-based compression without any performance penalty, and performance scales linearly with the number of instances.

In addition, the IBM FlashSystem 9500 control enclosure supports data reduction pool (DRP) compression and deduplication that can increase the effective capacity of your flash memory up to 5x, which decreases the cost of storing data up to 80%. DRPs support active data, unlike other data reduction solutions.

The IBM FlashSystem 9500 control enclosure offers several features for DRP compression workloads. These features include two 10 Intel core processors with up to 768 TB of memory per node and a built-in compression accelerator for hardware-assisted compression. In addition, the IBM FlashSystem 9500 system with IBM FCMs NVMe-type drives applies compression to any data that is not already compressed.

Agile integration

The IBM FlashSystem 9500 system features the following agile characteristics:

- ▶ Fully integrated system management.
- ▶ Application-aware data services.
- ▶ Advanced Encryption Standard (AES), data-at-rest encryption with all NVMe type drives, and IBM FCMs with FIPS 140-3.
- ▶ In-line hardware compression with IBM FCM.
- ▶ Tiering or mirroring to existing and public cloud storage.
- ▶ Mixed workload consolidation.
- ▶ Nondisruptive data migrations.
- ▶ Concurrent code load.

By accelerating physical and virtual applications, the IBM FlashSystem 9500 system can help organizations reduce costs, increase revenue, and improve customer satisfaction for all types of applications, which includes the following categories:

- ▶ Transactional
- ▶ Enterprise resource planning (ERP) and supply chain management
- ▶ Big data and analytics
- ▶ Server and desktop virtualization
- ▶ Cloud native
- ▶ Hybrid and multicloud

AI-empowered

The IBM FlashSystem 9500 system includes the following AI characteristics:

- ▶ AI-based data placement for optimal data center performance and zero-downtime data migration.
- ▶ IBM Storage Insights with AI-derived predictive analytics, storage resource management, and a support platform that is delivered over the cloud.

Multicloud enabled

IBM Storage Virtualize on-premises and IBM Storage Virtualize for Public Cloud together support mirroring between on-premises and cloud data centers or between cloud data centers. These functions can be used to perform the following tasks:

- ▶ Migrate data between on-premises and public cloud data centers or between public cloud data centers to provide consistent data management between on-premises storage and the public cloud
- ▶ Implement Disaster Recovery (DR) strategies between on-premises and public cloud data centers
- ▶ Enable cloud-based DevOps to help make replication of data from on-premises sources easier
- ▶ Improve cyber resilience with copies on Amazon Web Services (AWS) by using *air gap* snapshots to S3 and IBM Safeguarded Snapshot on Microsoft Azure

Customers can create hybrid, multicloud solutions for their traditional block data and workloads by using built-in IP replication capabilities.

Customers can deploy IBM Storage Virtualize for Public Cloud on AWS directly from AWS Marketplace. By using a predefined Cloud Formation Template, customers can automatically and securely install the software and deploy an HA two-node cluster on selected EC2 instances. Any Amazon EBS block storage can be attached.

On Azure, customers can deploy directly from the Azure Apps Marketplace by using predefined Azure Resource Manager templates that automatically and securely install the software. Customers can deploy an HA two-node cluster on selected Azure VMs and can attach shared Azure Managed Disk solid-state drive (SSDs) to the cluster for HA.

On IBM Cloud®, automated installation scripts assist deployment of the software on bare-metal servers. IBM Performance or Endurance block storage is supported behind the cluster. This setup enables clients to create clustered configurations, such as on-premises, and brings the optimization and virtualization capabilities of IBM Storage Virtualize to public cloud infrastructures.

IBM Storage Virtualize on-premises and IBM Storage Virtualize for Public Cloud together enable a hybrid multicloud deployment with a single data management layer between on-premises systems and the cloud across heterogeneous storage pools that might exist in the data center. IBM Storage Virtualize provides the following functions:

- ▶ Storage pooling and automated allocation with thin provisioning
- ▶ Easy Tier® automated tiering
- ▶ Deduplication and compression to reduce cloud storage costs
- ▶ IBM FlashCopy® and remote mirror for local snapshots and remote replication
- ▶ IBM Safeguarded Snapshot for cyber resiliency
- ▶ Support for the following virtualized and containerized server environments:
 - VMware
 - Microsoft Hyper-V
 - IBM PowerVM®
 - Red Hat OpenShift
 - CRI-O
 - Kubernetes

For more information about IBM FlashSystem systems and hybrid multicloud, see [Embracing hybrid cloud: Storage edition](#).

Data resilience

The IBM FlashSystem 9500 provides advanced capabilities that can help maximize data protection, security, and HA to significantly reduce the risk of disruption and financial losses that can occur because of user errors, malicious destruction, or ransomware attacks. FlashSystem 9500 delivers the capability to enable this level of protection while also delivering high performance for applications.

With Safeguarded Snapshot, the IBM FlashSystem 9500 adds a line of defense against cyberthreats by protecting your data from cyberattacks with immutable and isolated copies that are hidden, non-addressable, and cannot be altered. If an attack occurs, these copies can be quickly restored to support recovery. Customized to your specific application mix, Cyber Vault builds on IBM Safeguarded Snapshot to help reduce cyberattack recovery times.

A Cyber Vault solution runs continuously and monitors snapshots as they are created automatically by Safeguarded Snapshot. By using standard database tools and automation software, Cyber Vault checks Safeguarded snapshots for corruption. If Cyber Vault finds such changes, it is an immediate sign that an attack might be occurring.

When you prepare a response, if you know the last snapshots with no evidence of an attack, you can reduce the time to determine which snapshot to use. Because Safeguarded snapshots are on the same FlashSystem storage as operational data, recovery is fast by using the same snapshot technology. With these advantages, Cyber Vault is designed to help reduce cyberattack recovery time from days to just hours.

Refer to “IBM Cyber Vault solution” on page 51 for more information.

Advanced data services

The IBM FlashSystem 9500 system provides the following advanced data services:

- ▶ Business continuity with replication services across up to three sites
- ▶ Data protection with IBM FlashCopy services
- ▶ High availability with the IBM policy-based HA function
- ▶ Higher storage efficiency with thin provisioning
- ▶ DRPs that provide compression with deduplication
- ▶ Data-at-rest AES-XTS 256 encryption
- ▶ IBM Easy Tier
- ▶ External virtualization
- ▶ Data migration
- ▶ IP quorum support
- ▶ N_Port ID virtualization (NPIV) support
- ▶ VMware vSphere Virtual Volume (VVOL) support and space-efficient copies
- ▶ Transparent Cloud Tiering (TCT)
- ▶ Data and cyber resilience with IBM Safeguarded Snapshot
- ▶ Container support with the IBM block storage CSI driver
- ▶ Capacity and performance scaling through clustering

IBM FlashSystem 9500 enclosures overview

The IBM FlashSystem 9500 system features the following types of enclosures:

- ▶ A *control enclosure* manages your storage systems, communicates with the host, and manages interfaces. In addition, it can house up to 48 NVMe-capable flash drives. These drives can be industry-standard NVMe types or the exclusive IBM FlashCore Module (FCM) NVMe type or up to 12 optional Storage Class Memory (SCM) type drives. A control enclosure contains 2 node canisters which form an *I/O group*.
- ▶ An *expansion enclosure* increases the available capacity of an IBM FlashSystem 9500 cluster. It communicates with the control enclosure through a dual pair of 12 Gbps serial-attached SCSI (SAS) connections. This enclosure can house many of the flash (SSDs) SAS type drives, depending on which model of enclosure is ordered. FCM's are not permitted in any of the SAS Expansion Enclosures.

Figure 1 shows the IBM FlashSystem 9500 control enclosure front view where you can see the 48 NVMe drives in two rows of 24 drives.

Note: There are new rules for the plugging of the NVMe drives in the control enclosure. See the “IBM FlashSystem 9500 NVMe drive options” on page 30.



Figure 1 IBM FlashSystem 9500 control enclosure front view

Control enclosures

Each control enclosure can have multiple attached expansion enclosures, which expand the available capacity of the entire system. The IBM FlashSystem 9500 solution supports up to two control enclosures and up to two chains of SAS expansion enclosures per control enclosure.

Note: At the time of writing, there was a change in the machine type (MT) and model numbers because of the implementation of Expert Care. MT 4666 is still available through restricted ordering.

The IBM FlashSystem 9500 control enclosure supports up to 48 NVMe-capable flash drives in a 4U high form factor and consists of the following machine types and models:

- ▶ 4666-AH8: IBM FlashSystem 9500 control enclosure
- ▶ 4983-AH8: IBM FlashSystem 9500 control enclosure with Expert Care
- ▶ 4666-UH8: IBM FlashSystem 9500 control enclosure utility model

Notes:

- ▶ IBM Storage FlashSystem 9500, MT 4666, requires IBM Storage Virtualize Software for IBM Storage FlashSystem 9500 8.5.0, or later, for operation. Use of the software is entitled through the acquisition of IBM Storage Virtualize software licenses. IBM Storage FlashSystem 9500 system function capabilities are provided through IBM Storage Virtualize software. The software is licensed through IBM Storage Virtualize Software for FlashSystem 9500 (SW PID 5639- 011) for control enclosures and IBM Storage Virtualize Software for FlashSystem 9000 Expansions (SW PID 5639-EB2) for expansion enclosures. A three-month registration for software maintenance (SWMA) for 5639-012 is required if IBM Storage Expert Care is not ordered with the system.
- ▶ The 4666 UH8 model is an IBM FlashSystem 9500 solution with a 1-year warranty, optional Storage Expert Care service offerings, and is a 4666 offered in the Storage Utility Offering space. This model is physically and functionally identical to the IBM FlashSystem 4666 AH8 model, except for target configurations and variable capacity billing.
- ▶ The variable capacity billing uses IBM Storage Control or IBM Storage Insights to monitor the system usage, which allows allocated storage usage above a base subscription rate to be billed per terabyte per month. Allocated storage is storage that is allocated to a specific host and unusable to other hosts, whether or not data is written. For thin-provisioning, the data that is written is considered used. For thick-provisioning, the total allocated volume space is considered used.
- ▶ IBM Storage FlashSystem 9500, MT 4983, requires IBM Storage Virtualize Software for IBM Storage FlashSystem 9500 8.7.0, or later, for operation. Because the 4983 uses Licensed Internal Code (LIC), all licenses are included except encryption and external virtualization.

Each IBM FlashSystem 9500 control enclosure includes the following features:

- ▶ Two node canisters, each with two 24-core processors and options for up to 1.5 GB memory to deliver a system total of 3 TB memory.

- ▶ Up to 12 (six per canister) I/O adapter cages to add PCIe adapter features:
 - Four-port 32 Gbps FC-NVMe card
 - Four-port 64 Gbps FC-NVMe card (only 3 adapters per canister supported)
 - Two-port 25 Gbps Ethernet iSCSI or iSCSI Extensions for RDMA
 - Two-port 25 Gbps Ethernet iSCSI and PBHA over iWARP (RPQ only) card
 - Two-port 100 Gbps Ethernet iSCSI and NVMe RDMA card
 - 12 Gbps SAS ports for expansion enclosure attachment
- ▶ Support for up to 48 NVMe flash or FCM drive options.
- ▶ Hot-swappable batteries, boot drives, and AC power supplies.
- ▶ 4U, 19-inch rack mounted enclosure.

For more information about configuration and limit restrictions, see [V8.7.0.x Configuration Limits and Restrictions for IBM FlashSystem 9500](#).

Expansion enclosures

The IBM FlashSystem 9000 expansion enclosures consist of the following machine types and models:

- ▶ 4666 model AFF small form factor (SFF)
- ▶ 4983 model AFF small form factor (SFF)
- ▶ 4666 model A9F large form factor (LFF)
- ▶ 4983 model A9F large form factor (LFF)

The new SAS-based SFF and LFF expansion enclosures support flash SSDs as MDisks in a storage pool, which can be used for IBM Easy Tier. The following expansion enclosure models are available:

- ▶ IBM FlashSystem 9000 SFF Expansion Enclosure Model AFF offers drive options with up to 24 2.5-inch SSD flash drives. Up to 72 drives in three AFF SAS expansion enclosures are supported per IBM FlashSystem 9500 per SAS chain. The AFF expansion enclosure is 2U high.
- ▶ IBM FlashSystem 9000 LFF Expansion Enclosure Model A9F offers drive options with up to 92 3.5-inch (and 2.5-inch drives in carriers) SSD flash drives. The model A9F SAS expansion enclosure supports up to 92 drives per IBM FlashSystem 9500 Expansion enclosure per SAS chain.

SAS chain limitations

When you attach expansion enclosures to the control enclosure, you are not limited by the type of the enclosure. The only limitation for each of the two SAS chains is its chain weight. The IBM FlashSystem 9500 supports a maximum SAS chain weight of 3. Each type of enclosure has its own chain weight:

- ▶ IBM FlashSystem 9000 Expansion Enclosure Model AFF has a chain weight of 1.
- ▶ IBM FlashSystem 9000 Expansion Enclosure Model A9F has a chain weight of 2.5.

For example, with the IBM FlashSystem 9500, you can have three IBM FlashSystem 9000 Expansion Enclosure Model AFF, or you can have one IBM FlashSystem 9000 Expansion Enclosure Model A9F expansion per SAS chain.

Note: A 5U expansion enclosure cannot be mixed with 2U expansion enclosures on the same chain.

For more information about Enclosures, see the web link [Enclosures](#).

For more information about configuration and limit restrictions, see [V8.7.0.x Configuration Limits and Restrictions for IBM FlashSystem 9500](#).

IBM FlashSystem 9500R rack

The FlashSystem 9500R rack offering is a preassembled, installed, and configured by IBM product that contains the following components:

- ▶ A pair of clustered IBM FlashSystem 9500 Control Enclosures Models AH8, which can be specified by ordering a IBM FlashSystem 9502R.
- ▶ Two 32G Gbps FC switches for dedicated FC clustered network.

Other components, such as IBM FlashSystem 9000 Expansion Enclosure Models AFF and A9F, can be added to the rack after delivery to meet the growing needs of the business.

Figure 2 shows the legend that is used to denote the component placement and mandatory gaps for the figures that show the configurations.



Figure 2 9500R component placement legend

Figure 3 shows the layout of the components in the IBM FlashSystem 9500R rack enclosure and the positions that are assigned for the option expansion enclosures.




Figure 3 IBM FlashSystem 9500R rack enclosure

For more information about the IBM FlashSystem 9500R offering, see [IBM FlashSystem 9500R Data Sheet](#).

Placing your IBM FlashSystem 9500 in your infrastructure

Various possible use cases are available for the IBM FlashSystem 9500. All of which can address many technical or business requirements. In addition, the FlashSystem 9500 can be used to simplify the overall storage architecture.

The FlashSystem 9500 can be used in the following scenarios:

- ▶ First tier storage repository for production data.
- ▶ Primary or target system for data replication or disaster recovery (DR).
- ▶ HA storage within a PBHA configuration (where two FlashSystem 9500s are in a synchronous replication relationship).
- ▶ Externally virtualizing IBM or non-IBM storage by using IBM Storage Virtualize, which extends advanced functions (such as data reduction, encryption, and replication) to those storage subsystems.

- ▶ Data migration from IBM or non-IBM storage with IBM Storage Virtualize.
- ▶ Hybrid cloud storage by:
 - Enabling communication between an on-premises deployment of IBM Storage Virtualize and IBM Storage Virtualize for Public Cloud.
 - Deploying the Container Storage Interface (CSI) driver for Red Hat OpenShift Container Platform to provide persistent storage for on-premises or cloud-based containerized applications.
 - Using transparent cloud tiering to convert data into an object store for back up to certain available cloud instances.
- ▶ Support and interoperability within the IBM Storage Software Suite, including:
 - IBM Storage Scale where the FlashSystem 9500 can be used as the back-end storage repository for metadata, primary data, or in initial microprogram load (IML) processing.
 - IBM Storage Protect and IBM Storage Protect Plus where the FlashSystem 9500 can be used as cache or data repository.
- ▶ IBM Copy Data Management where the FlashSystem 9500 can be used as cache or a copy services target.

Figure 4 shows an IBM FlashSystem 9500 that is fully deployed in a hybrid multicloud environment in which it is also externally virtualizing various on-premises storage subsystems.

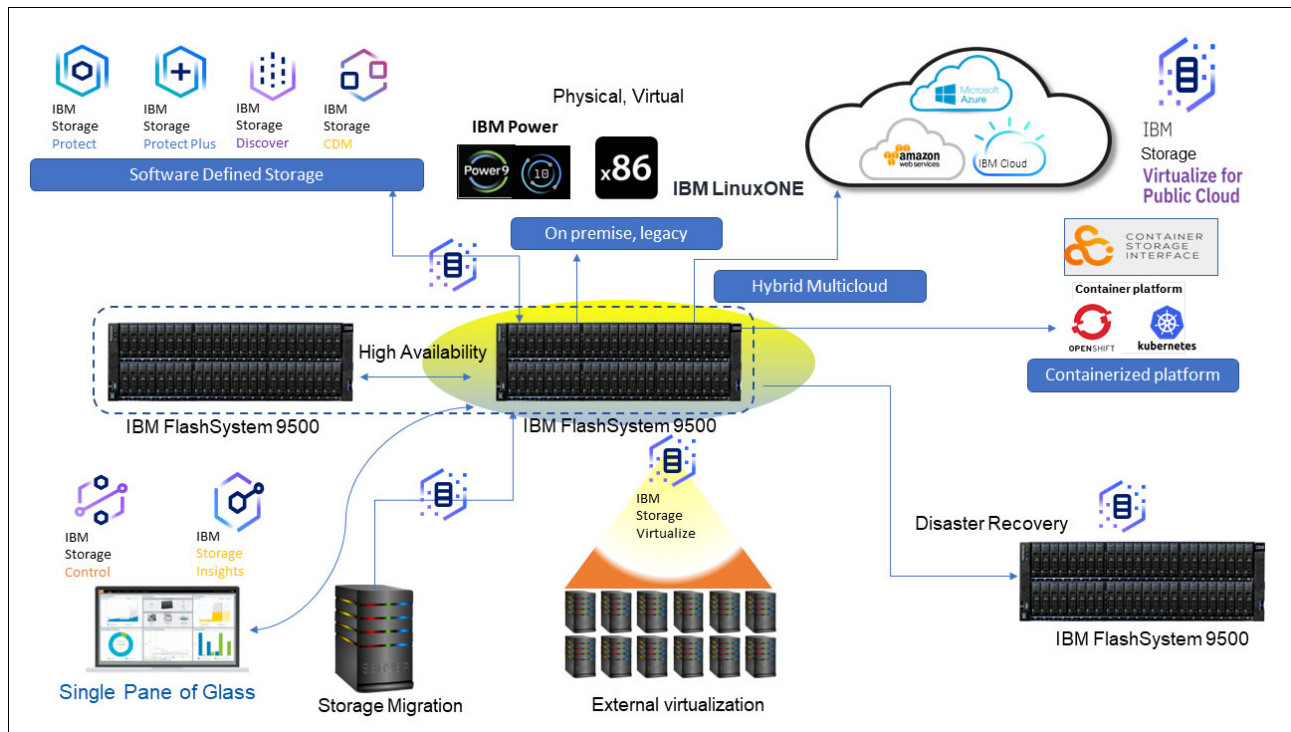


Figure 4 Fully deployed IBM FlashSystem 9500 sample scenario

The IBM FlashSystem 9500 can support multiple Open System hosts and interoperability with all of the software from the IBM Storage Suite of products.

In this scenario, the IBM FlashSystem 9500 provides the following benefits:

- ▶ Investment protection of established storage by using IBM Storage Virtualize to migrate from older storage subsystems and externally virtualize existing heterogeneous storage.
- ▶ A hybrid multicloud architecture by connecting to IBM Storage Virtualize for Public Cloud, which is available in Amazon S3, OpenStack Swift, Microsoft Azure, and IBM Cloud.
- ▶ Data protection by extending HA and DR to heterogeneous storage that is externally virtualized.
- ▶ Support for containerized workloads by interfacing with Red Hat OpenShift and Kubernetes for on-premises and off-premises infrastructures.
- ▶ Ease of management and seamless integration with the IBM Storage Insights and all the products in the IBM Storage Software Suit.

Features of IBM FlashSystem 9500

IBM Storage Virtualize 8.7 delivers support for IBM FlashSystem 9500 machine type 4666-AH8, the new 4983-AH8 and utility model 4666-UH8.

This release provides the advanced software functions of IBM Storage Virtualize software to FlashSystem 9500. and includes the following features:

- ▶ Control enclosure support for SSDs, including up to 48 industry 2.5-inch SFF standard NVMe drives and IBM FlashCore Modules and up to 12 SCM drives.

The following NVMe drives are supported:

- 2.5-inch SFF 4.8 TB, 9.6 TB, 19.2 TB, or 38.4 TB IBM FlashCore Module NVMe drives
 - 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB industry-standard drives
 - A mixture of both
 - Storage Class Memory drives of 1.6 TB
- ▶ Autonomic tiering through Easy Tier, which enables workload data to be placed on the most suitable storage media (various classes of NVMe for the control enclosure and HDDs and commodity SSDs for the expansion enclosures and externally virtualized storage).
 - ▶ Hardware-based data reduction and encryption from the IBM FlashCore Modules with negligible impact on system performance.
 - ▶ Scale-out capacity options with the 12 Gbps SAS-attached model A9F and AFF expansion enclosures.
 - ▶ Software-based AES-256 data-at-rest encryption and key management through IBM Security® Guardium® Key Lifecycle Manager or USB key encryption. Up to four IBM Security Guardium Key Lifecycle Manager key servers are supported.
 - ▶ Remote support assistance (RSA), which enables IBM support personnel to access the system to complete troubleshooting tasks.
 - ▶ Multiple data-protection solutions, such as FlashCopy (point-in-time-copy), Metro Mirror (synchronous replication), and Global Mirror (asynchronous replication).
 - ▶ Cyber resiliency with Safeguarded Snapshot, which provides rapid ransomware recovery by using immutable, untouchable snapshots that are based on FlashCopy technology.

- ▶ Multiple management modes by using a graphical user interface (GUI) or a Linux-based command-line interface (CLI).
- ▶ Performance throttling that allows control of used resources when the system is processing host I/O, advanced functions, or copy services.
- ▶ Transparent cloud tiering (TCT), which enables a FlashCopy volume that can be backed up to an S3 object store, such as AWS or IBM Cloud.
- ▶ Support for 32 / 64 Gbps FC and 10 or 25 Gbps Ethernet or 100 Gbps Ethernet iSCSI and NVMe RDMA connectivity. The ability to intermix these port options is also supported.
- ▶ Simplified connectivity by using portsets, which enable host connectivity ports to be grouped for more effective zoning and management.
- ▶ Easier access and compatibility with IBM Storage Insights, which offers advanced performance monitoring for one or more IBM FlashSystem 9500 storage systems.
- ▶ Additional access security by using multi-factor authentication.

In addition to the base warranty, the IBM FlashSystem 9500 offers IBM Storage Expert Care with two levels of support options. Either IBM Storage Expert Care Advanced 1-5 years or IBM Storage Expert Care Premium 1-5 years.

Current release functions

This section describes several new functions and features that are available, at the time of writing, in the IBM FlashSystem 9500 release.

Products that support IBM Storage Virtualize software follow a continuous development release model, beginning with IBM Storage Virtualize 8.4.0.

You can choose from the long-term support (LTS) model that remains stable so that you receive only security and defect fixes with no incremental functional enhancements. The other option is the Continuous Delivery (CD) support offering that receives regular updates with fixes and new features. Each update incrementally delivers new capabilities and defect fixes.

LTS users demand the highest levels of stability during the support period with updates containing fixes to identified problems and vulnerabilities. LTS users prefer stability and a longer horizon. LTS updates provide PTFs for security and defect fixes, but do not provide new functions. LTS releases are offered approximately every 1–2 years.

CD users prefer to receive the new functions as they are made available throughout the year. CD users prefer to use the most recently released features and capabilities of IBM Storage Virtualize without waiting for the next LTS release cycle. CD releases are offered approximately every few months.

Important: After you move to a new CD release, you are unable to move back to a previous release.

Figure 5 describes a continuous development model with both options.

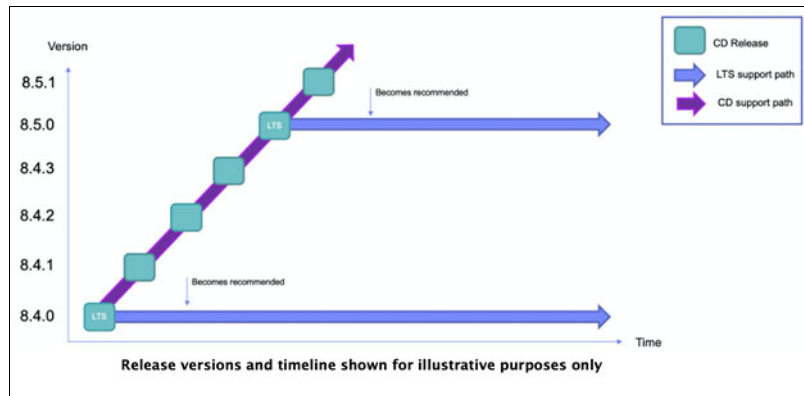


Figure 5 Release Types

New in IBM FlashSystem 9500 Software V8.7.0

- ▶ Policy-based high availability:
 - Define host location
 - Local-site I/O access to avoid cross-ISL traffic

The High Availability (HA) solution uses storage partitions, a configuration object that is the single point of management for HA objects. These partitions contain volumes, volume groups, hosts, and host-to-volume mappings.

Within a partition:

- All volumes are in volume groups.
- Mappings can be created only between volumes and hosts in the same partition.
- HA replication policies are associated with storage partitions to define the paired systems providing high availability. Partition membership automatically configures objects across both systems, enabling host access from either system. Hosts can prioritize the local system for data access. Before creating partitions, partnerships, linked pools, replication policies, and IP quorum must be established. Multiple HA partitions are supported, each with a preferred and active management system that can fail over. All partition elements are highly available with management and I/O failover. Partitions can be expanded with additional volumes, volume groups, hosts, and mappings. IP Quorum prevents split-brain scenarios by designating an active management system.
- ▶ TLS connection to send syslog notifications and messages.
 - The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The system can send syslog messages that notify personnel about an event. You can set up syslog event notifications with either the management GUI or the command-line interface (CLI).
- ▶ New Flash Grid
 - You can scale out a storage system by using a Flash Grid, which combines multiple independent systems into a federated cluster. Unlike traditional scale-out systems based on I/O groups, Flash Grid supports up to eight systems, delivering linear increases in performance, capacity, volumes, hosts, and snapshots. Each system within a Flash Grid can be updated independently, allowing for hardware diversity and

upgrades. Storage Partitions can be seamlessly migrated across systems within a Flash Grid using non-disruptive Storage Partition Migration.

Systems in a Flash Grid have the following requirements

- The system supports a maximum of a single I/O group and the compatibility mode for multiple I/O groups must be disabled.
 - The system topology must be standard.
 - The system layer must be replication.
- Systems in a Flash Grid exchange certificates to enable secure communication between members. System-signed certificates are the default, however a CA-signed certificate can be used instead providing it meets the requirements. For more information, see [Requesting and installing an externally signed certificate](#). A Flash Grid is supported on the following systems:
- IBM Storage FlashSystem 5200
 - IBM Storage FlashSystem 5300
 - IBM Storage FlashSystem 7200
 - IBM Storage FlashSystem 7300
 - IBM Storage FlashSystem 9100
 - IBM Storage FlashSystem 9200
 - IBM Storage FlashSystem 9500
 - SAN Volume Controller
- ▶ Automatic download of compatible security patches to the system.
- Under some scenarios, IBM releases a patch that fixes or updates a part of the product. Patch installation does not cause any interruption to the IO processing. Patches are made available on FixCentral with the help of direct download link (for more information, see [Obtaining packages](#)) or as part of a general collection of available patches.
- ▶ An increase in the number of short-distance partnerships over RDMA.
- ▶ To convert thin-clone volume to clone volume.
- ▶ Fast path Ethernet and VLAN for Management IP addresses.
- ▶ Non-disruptive storage partition migration from one FlashSystem to another using IBM Storage Insights.
- ▶ Ability to view volume group snapshot restore timestamp details.
- ▶ Ability for policy-based replication to be used on a different partnership for different volume groups while policy-based high availability is configured.
- ▶ Snapshot population for non-high availability partitions.
- ▶ RESTful APIs to provide system performance statistics.
- The IBM Storage Virtualize REST model API offers Common Information Model (CIM) and ways to extract general performance statistics from a Storage Virtualize system for new users Application programmers, who are migrating away from CIM to REST API, and new REST API users can complete the following tasks to extract system performance statistics:
- Understand the now decommissioned CIM-based performance statistics for the IBM Storage Virtualize systems.
 - Discover and connect to the REST API to list and download performance data files.

- Retrieve and extract metrics and apply calculations that were previously applied by the CIM agent.
- CIM, a DMTF standard for system management, is increasingly being replaced by RESTful APIs due to their performance, scalability, and widespread adoption. IBM Storage Virtualize offers a high-performance REST API. To optimize system management and reduce script complexity, it is recommended to migrate from CIM to REST. This document provides guidance on the migration process, CIM agent calculations, and REST API code examples.
- ▶ Ransomware threat detection:
 - File-system awareness for ransomware threat detection.
 - Storage Insights Pro entitlement information.
- ▶ Automatic update of candidate FlashCore Modules (FCM).
- ▶ Policy-based high availability and policy-based replication for the IBM FlashSystem 5045.
- ▶ User-configurable default grain size for new snapshots and FlashCopy mappings.

For a complete list of all changes that affect the IBM FlashSystem 9500, see [What's new in Version 8.7](#).

IBM FlashCore technology

When IBM FlashCore Module (FCM) flash drives are installed and configured in the IBM FlashSystem 9500, the combination of FCMs and the IBM FlashSystem 9500 provide the following benefits:

- ▶ Hardware-accelerated architecture that is engineered for flash, with a hardware-only data path.
- ▶ A modified dynamic GZIP algorithm for IBM FlashSystem data compression and decompression. Because the algorithm is implemented completely in hardware, no processor intervention is required.
- ▶ Design for low latency, density, and reliability.
- ▶ IBM Advanced Flash Management, which improves flash endurance over standard implementations without increasing latency.
- ▶ Capability to run a maximum of 48 FCM drives in a single control enclosure.

Figure 6 shows IBM FlashCore technology. For more information about IBM FlashCore technology, see [What is flash storage?](#) or the IBM Redpaper *IBM FlashCore Module (FCM) Product Guide: Features the newly available FCM4 with AI-powered ransomware detection*, REDP-5725.

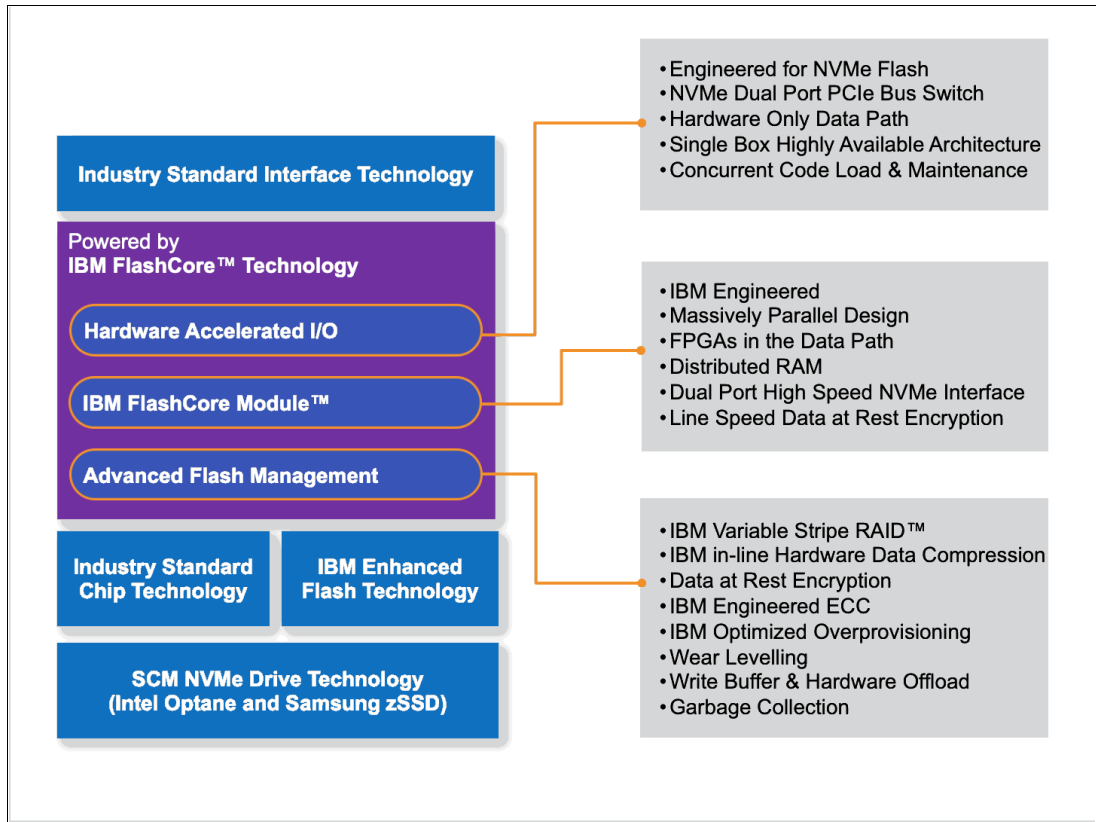


Figure 6 IBM FlashCore technology

Storage Class Memory

Storage Class Memory (SCM) is another way to describe nonvolatile memory devices that perform faster at approximately 10 μ s than traditional NAND SSDs, which perform at 100 μ s. Both are slower than dynamic random access memory (DRAM), which has response times of approximately 100 ns.

The typical cost for SCMs is between the costs for DRAM and traditional NAND. The price for SCMs is significantly more expensive than traditional NAND drives.

The IBM FlashSystem 9500 system supports the new low-latency, high-speed SCM drives in any of the slots of the control enclosure. The control enclosure can contain up to 12 x 1.6 TB 2.5-Inch NVMe Storage Class Memory Drives.

Note: SCM and other NVM Express (NVMe) drive types can be installed in any drive slot. However, the highest capacity drives must be installed in the lowest available drive slots.

SCM and Easy Tier

Because of their faster speed, SCM drives are placed in a new top tier of Easy Tier. This new tier is ranked higher than the existing tier0_flash that is used for NVMe NAND drives that are supported today. The following Easy Tier levels are available:

Storage Class Memory tier

Exists when the pool contains drives that use persistent memory technologies that improve the endurance and speed of current flash storage device technologies.

Tier 0 flash

Exists when the pool contains high-performance flash drives.

Tier 1 flash

Exists when the pool contains tier 1 flash drives. Tier 1 flash drives typically offer larger capacities, but slightly slower performance and lower write endurance characteristics.

Enterprise tier

Exists when the pool contains enterprise-class MDisks, which are disk drives that are optimized for performance.

Nearline tier

Exists when the pool contains nearline-class MDisks, which are disk drives that are optimized for capacity.

Hot data is placed in the SCM tier when Easy Tier is enabled. Extent allocation for data reduction pools (DRPs) and Volume Groups (VGs) is tuned to use SCM drives, especially for metadata (directory volume lookups). This use case is one of the main use cases for SCM drives with an IBM FlashSystem 9500 system.

SCM and RAID configurations

SCM drives feature the following rules concerning RAID supported configurations:

- ▶ Distributed RAID 1 (DRAID 1) with two drives (including distributed spare capacity) or more and is the best practice recommendation and configuration.
- ▶ Distributed RAID 6 (DRAID 6) with six drives (including distributed spare capacity) or more is supported.
- ▶ SCM drives have their own SCM technology type and drive class.
- ▶ SCM and Standard NVMe (or SAS) cannot intermix in the same array.
- ▶ Easy Tier tier_scm arrays can take lower tech types as spare drives when no tier_scm drives are available.
- ▶ Lower tier arrays can accept tier_scm drives as superior drives.

Note: DRAID 5 and traditional RAID are not supported by FS9500.

SCM drive formatting and UNMAPS

Because of its media technology, SCM drive formats can take 15 minutes, which is typically more time than is needed for formatting an NVMe drive.

SCM drive formats occur when one of the following conditions exists:

- ▶ A drive is replaced.
- ▶ An enclosure is managed.
- ▶ An array or MDisk is removed.

Note: Intel Optane drives do not support UNMAPs because it is not beneficial for these drives. Instead, the system writes zeros to deallocate drive sectors.

IBM Storage Expert Care

IBM FlashSystem 9500 is a client-installed product and is being released with a one-year warranty with the option to include IBM Storage Expert Care Advanced or IBM Storage Expert Care Premium.

Note: IBM FlashSystem 9500 does not include Basic Expert Care. A three-month software maintenance (SWMA) offering is defaulted when no IBM Storage Expert Care Service and Support option is selected.

Figure 7 shows IBM Storage Expert Care levels.

	Warranty	Basic 5200, 7200, 7300	Advanced 5200, 7200, 7300, 9200, 9500	Premium 7200, 7300, 9200, 9500
IBM Spectrum Virtualize fixes, updates and new releases	1 year	Yes	Yes	Yes
Guidance on installation, usage and configuration		Yes	Yes	Yes
Automated ticket management and alerting		Yes	Yes	Yes
Use of Storage Insights for collaborative problem resolution		Yes	Yes	Yes
Predictive issue alerting			Yes	Yes
Storage Insights Pro entitlement				Yes
IBM Installation	Standard with 9500	Additional paid service	Additional paid service	Additional paid service for 7300
Remote code upgrades (2x year) ***				Yes
Dedicated Technical Account Manager (TAM)				Yes
30 minutes Severity 1/2 response				Yes
Hardware service / parts replacement	9x5 NBD* or 24x7 Same day**	9x5 NBD, IBM on-site	24x7 Same day, IBM on-site	24x7 Same day, IBM on-site

* Next business day, parts only for FS5200, FS 7200 and FS7300 ** Same day, IBM on-site for FS 9200 and FS9500 *** On-site available as additional paid service

Figure 7 IBM Storage Expert Care levels

IBM Storage Expert Care options

IBM Storage Expert Care Advanced

The following support features are included with the IBM Storage Expert Care Advanced option:

- ▶ Available in 3–5 years duration.
- ▶ IBM onsite repair coverage includes 24x7, same-day support that is based on the duration that is chosen.
- ▶ IBM Software Maintenance (SWMA) software support and services agreement.
- ▶ IBM Storage Expert Care services:
 - Installation, usage, and configuration support line
 - Proactive issue resolution
 - Predictive alerting
 - Automated ticket management

IBM Storage Expert Care Premium

The following support features come with the IBM Storage Expert Care Premium option:

- ▶ Available in 3–5 years duration.
- ▶ IBM onsite repair coverage includes 24x7, same-day support that is based on the duration that is chosen.
- ▶ Enhanced response time for defect support:
 - *Response time* is defined as the elapsed time between technical support receipt of the customer problem submission and the acknowledgment of the submission.
 - 30-minute targeted response time objectives for Severity 1 and Severity 2 problem reports.
- ▶ IBM Storage Expert Care services:
 - Installation, usage, and configuration support line
 - Proactive issue resolution
 - Predictive alerting
- ▶ Automated ticket management
- ▶ IBM Storage Insights Pro with more detailed reports and planning
- ▶ Remote Code Load:
 - Code upgrades are provided by IBM remote support representatives up to twice per year to keep your systems updated.
 - On-site code load services that are available through feature code #AHY2 - Remote code load exception when required.
- ▶ Technical Account Manager to serve as the key client interface.

Note: A Technical Account Manager serves as the key client interface for in-scope hardware and software, delivering partnership and consultancy, and direct engagement on high-priority support cases.

Storage Assurance

IBM Storage Assurance Perpetual offers a subscription model for IBM FlashSystem that ensures consistent modernization, budget predictability, and premium support. This option includes full-system refreshes, eliminating lifecycle challenges like terminal code levels, end-of-support concerns, disruptive migrations, and procurement hurdles. Available as a 4 or 8-year contract, it simplifies storage management and reduces costs.

For more information see: [IBM Storage Assurance](#).

Table 1 Overview of the available IBM Storage Assurance features. Valid for all FlashSystem models

Years	Upfront	Annual	Quarterly	Monthly
1	ALF0	ALF8	ALFG	ALFO
2	ALF1	ALF9	ALFH	ALFP
3	ALF2	ALFA	ALFI	ALFQ
4	ALF3	ALFB	ALFJ	ALFR
5	ALF4	ALFC	ALFK	ALFS

Years	Upfront	Annual	Quarterly	Monthly
6	ALF5	ALFD	ALFL	ALFT
7	ALF6	ALFE	ALFM	ALFU
8	ALF7	ALFF	ALFN	ALFV

Note: Availability may vary depending on country offering.

The Perpetual Model offers transparent, flat-rate pricing for budget predictability. It ensures access to cutting-edge technology, boosting operational agility. With a future-proof all-flash platform, continuous innovation is unlocked. Enjoy flexible payment options (upfront, annual, quarterly, or monthly) on 4 or 8-year contracts, backed by premium ExpertCare support, comprehensive software licensing, and automatic hardware refreshes aligned with SLAs:

- ▶ Refresh if the system is not meeting your performance¹ or energy efficiency needs.
- ▶ Refresh if your system is no longer compatible with the latest SW release or if your system has reached its end-of-support (EOS) date.
- ▶ Refresh if your system cannot add more capacity.
- ▶ Refresh at least once during an 8-year contract.

IBM Storage Assurance Perpetual delivers continuous innovation, operational agility, and budget certainty with IBM FlashSystem. Enjoy seamless access to advanced features, including AI-powered data management, ransomware protection, and non-disruptive upgrades. Optimize storage efficiency, performance, and resilience while simplifying management across hybrid cloud environments.

IBM Storage Assurance – Perpetual model feature overview

Software upgrades get full access to all-inclusive array and monitoring features at the time of purchase and in the future as they become available. Hardware upgrades experience all-inclusive full-system hardware refreshes if your current device cannot upgrade to the latest software, cannot add more capacity, is end of support, or does not meet your performance needs.

Premium Support Premium-level ExpertCare

Predictive analytics support, capacity planning, monitoring and reporting with IBM Storage Insights Pro entitlement. Guarantees:

- ▶ Non-disruptive system-to-system migrations guarantee.
- ▶ Price protection guarantee.
- ▶ Workload performance guarantee.
- ▶ End of service guarantee.
- ▶ Energy efficiency guarantee.
- ▶ Eight-year contract full-system guarantee.
- ▶ Terminal code level guarantee.
- ▶ Capacity upgrade eligibility guarantee.

Considerations:

- ▶ Performance guarantee only after year 3.
- ▶ Optional feature, priced separately.
- ▶ For high-level contract term details and country availability, see this [link](#).

IBM FlashWatch overview

IBM FlashWatch is an offering from IBM to complement the purchase of the IBM FlashSystem 9500. It features these main focus areas:

Cyber recovery

Once configured by IBM Expert Labs, the FlashSystem Cyber Recovery Guarantee states that immutable snapshots from Safeguarded Snapshot can be recovered in just 60 seconds. This offering includes FlashSystem implementation, Safeguarded Snapshot configuration and recovery validation. The scope of work also enhances security options such as Two Person Integrity and Multi-Factor authentication. This also brings capabilities such as our new ransomware threat detection capabilities as part of the tools to detect an attack.

3:1 data reduction

IBM offers a 3:1 self-certified data reduction guarantee to give organizations peace of mind when they acquire a specific amount of raw capacity for their workloads. Organizations will benefit from the hardware assisted data reduction technology on IBM FlashCore Modules (FCMs). FCMs are a unique computational storage platform that due to advanced hardware assistance delivers a great compression ratio on the business's workloads without any performance impact

100% data availability

For critical workloads, IBM guarantees 100% high availability commitment through the IBM Policy based HA configuration. Zero down time during a 3-year period on a client's on-premises Storage FlashSystems. PBHA is an IBM high availability feature that provides dual-site, active-active access to a volume. PBHA functions are available on systems that can support more than one I/O group or machine, as applicable.

Storage Expert Care

All IBM FlashSystems come with 1-year limited warranty for all the components. See "IBM Storage Expert Care" on page 19.

Sustainability

IBM FlashSystems are the first primary block storage vendor to offer an energy efficiency / sustainability guarantee on client-owned devices. The IBM FlashSystem Sustainability Self-Certified Guarantee with FlashCore Modules can deliver energy efficiency of 1.7 W/TB on selected configurations. IBM FlashSystem can deliver a lower footprint, more TB capacity per watt and 10x more IOPs per watt than other competitors. IBM is all about sustainability; IBM's supply chain, products and efforts are directed to it. Learn more in [IBM Sustainability in Business](#).

IBM Storage Insights

IBM Storage Insights is another part of the monitoring capability of the IBM FlashSystem 9500 system and supplements the views that are available in the GUI.

IBM strongly suggests that all customers install and use this no-charge, cloud-based IBM application because it provides a single dashboard that provides a clear view of all your IBM block storage. You can make better decisions by seeing trends in performance and capacity.

With storage health information, you can focus on areas that need attention. When IBM support is needed, IBM Storage Insights simplifies uploading logs, speeds resolution with online configuration data, and provides an overview of open tickets all in one place.

IBM Storage Insights includes the following features:

- ▶ A unified view of IBM systems:
 - Provides a single view to see all your system's characteristics.
 - Displays all of your IBM storage inventory.
 - Provides a live event feed so that you know in near real-time the status of your storage.
- ▶ IBM Storage Insights collects telemetry data and Call Home data and provides real-time system reporting of capacity and performance. Overall storage monitoring, which provides the following information:
 - The overall health of the system.
 - Monitoring of the configuration to see whether it meets preferred practices.
 - System resource management to determine which system is overtaxed and provides proactive recommendations to fix it.
- ▶ IBM Storage Insights provides advanced customer service with an event filter that you can use to accomplish the following tasks:
 - You and IBM Support can view, open, and close support tickets, and track trends.
 - You can use the autolog collection capability to collect the logs and send them to IBM before IBM Support looks into the problem. This capability can save as much as 50% of the time to resolve the case.

In addition to the no-charge version of IBM Storage Insights, IBM offers IBM Storage Insights Pro, which is a subscription service that provides longer historical views of data, more reporting and optimization options, and supports IBM file and block storage with EMC VNX and VMAX.

Customer dashboard

Figure 8 shows a view of the IBM Storage Insights main dashboard and the systems that it monitors.

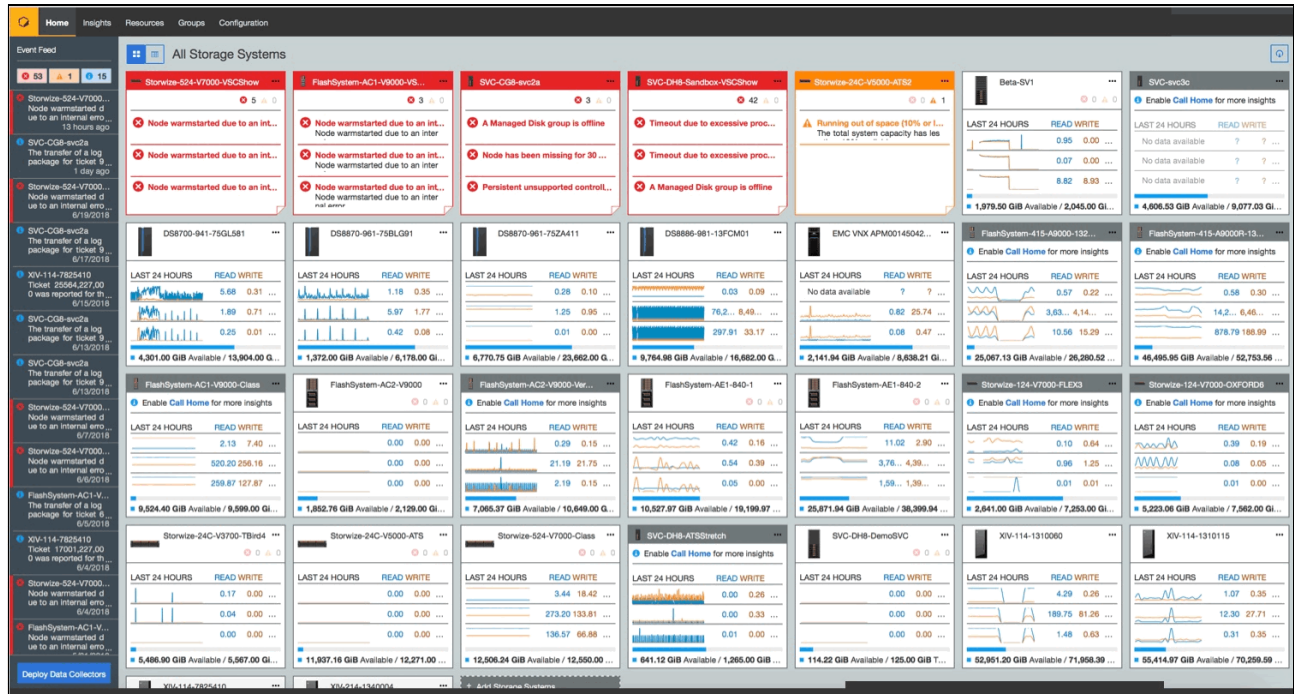


Figure 8 IBM Storage Insights dashboard

For more information about the architecture and design overview of IBM Storage Insights, see IBM Redbooks® *Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller*, [SG24-8561](#).

IBM Storage Insights: Information and registration

For more information about IBM Storage Insights, see the following resources:

- ▶ [IBM Storage Insights Fact Sheet](#)
- ▶ [Product demonstration](#) (log in required)
- ▶ [IBM Storage Insights Security Guide, SC27-8774](#)
- ▶ [IBM Storage Insights](#)

To register for free service, see [Sign Up for Storage Insights](#).

IBM Storage Connect

IBM FlashSystem 9500 Software supports integrating VASA and VAAI by using IBM Storage Connect V3.11.0 or later. For more information about this feature, see [IBM Storage Connect](#).

IBM Storage Control

IBM Storage Control consolidates a range of IBM storage provisioning, virtualization, cloud, automation, and monitoring solutions through a unified server platform. IBM Storage Control provides insight and awareness about the configuration capabilities, storage health, and events of a storage system regarding VMware and vSphere. With this capability, VMware administrators can independently and centrally manage their storage resources on IBM storage systems. For more information, see [IBM Storage Control](#).

IBM FlashSystem 9500 GUI

Because IBM FlashSystem 9500 control enclosures cluster to form a system, a single management interface is used for IBM FlashSystem 9500 control enclosures. Each IBM FlashSystem 9500 node canister is an individual server in an IBM FlashSystem 9500 clustered system on which the IBM Storage Virtualize software runs.

You can access the GUI by opening any supported web browser and entering the management IP addresses. You can connect from any workstation that can communicate with the IBM FlashSystem 9500 system. The IBM FlashSystem 9500 control enclosure is delivered in a 4U 19-inch rack-mount enclosure. The IBM FlashSystem 9500 system features an IBM Service Support Representative (SSR) installation as part of the product offering.

Figure 9 shows the IBM FlashSystem 9500 GUI dashboard.

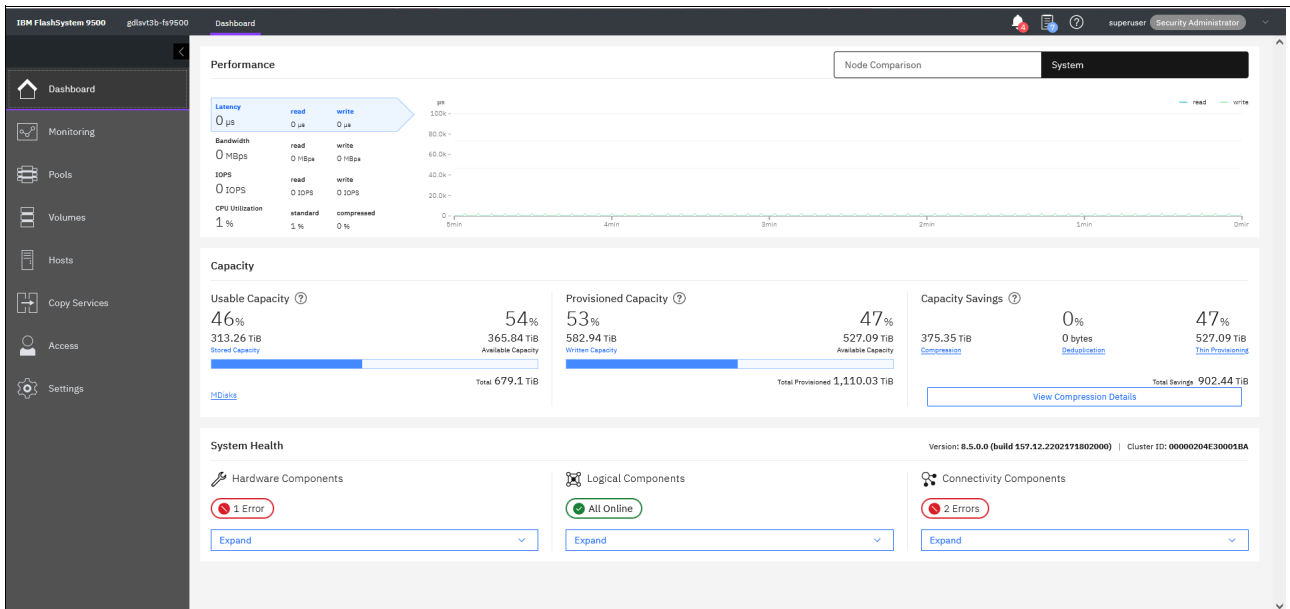


Figure 9 IBM FlashSystem 9500 GUI showing the dashboard

Figure 10 shows the Control Enclosure 1 window. You can open this window by selecting **Monitoring** → **System Hardware** from the left panel.

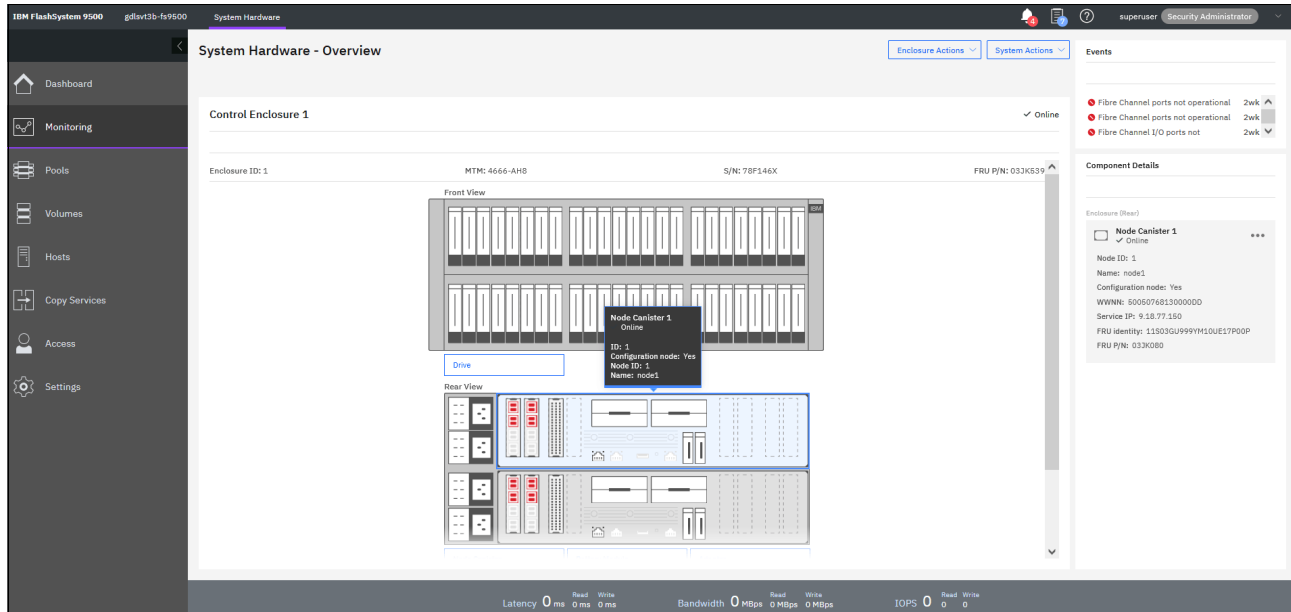


Figure 10 IBM FlashSystem 9500 system overview of the control enclosure

The IBM FlashSystem 9500 system includes a CLI, which is useful for scripting, and an intuitive GUI for simple and familiar management of the product. The IBM FlashSystem 9500 system supports SNMP email forwarding that uses Simple Mail Transfer Protocol (SMTP), and syslog redirection for complete enterprise management access.

By using the GUI, you can quickly deploy and manage storage. The GUI runs on the IBM FlashSystem 9500 control enclosure; therefore, a separate console is not needed. When you enter the system IP address in a web browser, you can manage all of the expansion enclosures from one place.

The IBM FlashSystem 9500 control enclosure node canisters are configured for active-active redundancy. The node canisters run a highly customized Linux based operating system that coordinates and monitors all significant functions in the system.

The node canisters provide a web interface, Secure Shell (SSH) access, and SNMP connectivity through external Ethernet interfaces. By using the web and SSH interfaces, administrators can monitor system performance and health metrics, configure storage, and collect support data, among other features.

The storage configuration includes defining logical units with capacities, access policies, and other parameters. No software must be installed on host computers to administer the IBM FlashSystem 9500 system beyond a web browser or a standard SSH client.

Supported platforms

The IBM FlashSystem 9500 system features extensive interoperability with support for a wide range of operating systems that includes Microsoft Windows Server, Linux, and IBM AIX® and IBM i. Supported hardware includes IBM Power Systems and x86 & x86_64 servers, host bus adapters (HBAs), and SAN fabrics. For more information, see [V8.7.0.x Configuration Limits and Restrictions for IBM FlashSystem 9500](#) and [IBM System Storage Interoperation Center](#).

IBM FlashSystem 9500 hardware component overview

The IBM FlashSystem 9500 control enclosure is a 4U rack-mounted NVMe flash memory enclosure that is based on IBM flash technology. It provides the primary management interface (GUI) and the host interface configuration. The IBM FlashSystem 9500 control enclosure supports FC Protocol (FCP and FC-NVMe) and iSCSI interfaces. For iSCSI, the RoCE and iWARP protocols are supported. Figure 11 shows the front view of the IBM FlashSystem 9500 control enclosure.



Figure 11 IBM FlashSystem 9500 control enclosure front view

Note: There are new rules for the plugging of the NVMe drives in the control enclosure. See the “**IBM FlashSystem 9500 NVMe drive options**” on page 30.

The IBM FlashSystem provides the following NVMe technologies:

- ▶ Supports:
 - Unique world class IBM FCM drives with inline compression
 - Industry-standard NVMe drives
 - Industry-standard NVMe drives with up to 1.47 PB of maximum raw capacity
- ▶ Option for SCM type drives
- ▶ Macroefficiency with up to 5.52 PB of raw maximum protected capacity with inline hardware data compression ratio of up to 3:1 if you use IBM FCM drives
- ▶ Latency measured in microseconds with IBM FCMs
- ▶ Optional expansion enclosures that provide tiering options with SSD flash drives

The IBM FlashSystem 9500 system is offered as three models:

1. 4666 model AH8: IBM FlashSystem 9500 NVMe control enclosure
2. 4666 model UH8: IBM FlashSystem 9500 NVMe control enclosure utility model
3. 4983 model AH8: IBM FlashSystem 9500 NVMe control enclosure

Note: The 4666-UH8 utility-based model features fixed configurations because they are Capacity on Demand (CoD)-based offerings.

A newer machine type 4983 model AH8 being introduced is physically identical to the 4666 except it is sold with LIC, which is in line with the other products in the IBM FlashSystem product line. This ensures that all features are included in the product price except for encryption.

The IBM FlashSystem 9500 configuration can consist of the following components:

- ▶ A total of 1–2 IBM FlashSystem 9500 control enclosures
- ▶ A total of 1–6 IBM FlashSystem 9000 SFF expansion enclosures
- ▶ A total of 1–2 IBM FlashSystem 9000 LFF High Density (HD) expansion enclosures

The following machine warranties are offered for the IBM FlashSystem 9500 system:

- ▶ Machine type 4666-AH8 with a 12-month, parts-only
- ▶ Machine type 4666-UH8 with a 12-month, parts-only
- ▶ Machine type 4983-AH8 with a 12 month, parts only

All models are offered with optional Storage Expert care to enhance the service from IBM with various levels of cover and duration. For more information, see “IBM Storage Expert Care” on page 19.

The IBM FlashSystem 9500 control enclosure includes integrated AC power supplies (PSU) and battery units inside each of the canisters. These batteries supply power to the control enclosure during a sudden power loss or failure so that the system can correctly commit all transactions to the storage medium.

The IBM FlashSystem 9500 control enclosure includes the following features:

- ▶ Control enclosure:
 - Two canisters that are placed one on top of the other
 - 48 NVMe drive slots
- ▶ Full internal redundancy:
 - Redundant and hot-swappable:
 - Canisters
 - Batteries within each canister
 - Power supplies and boot drives
 - Hot-swappable PCI Express (PCIe) adapters
 - Concurrently replaceable DIMMs and fans
- ▶ Each canister contains the following components and quantities:
 - CPU (2)
 - DIMM Slots (24)
 - Battery (2)
 - Canister Fans (5)
 - Power Supply PSU (2)
 - PCIe adapters cages (4)
 - PCIe adapter cards (0-6)
 - PCIe Compression QAT card (1)
 - PCIe adapter blanking plates (0-6)

Note: The number of PCIe adapters is configurable at product ordering time and can be added or removed by a sales MES.

Figure 12 shows a top-down view of the IBM FlashSystem 9500 enclosure. Highlighted are the various components of the control enclosure and the two canisters.




Figure 12 Top view of the IBM FlashSystem 9500 enclosure

Figure 13 shows the rear view of the IBM FlashSystem 9500 control enclosure. All components are concurrently maintainable, except for the passive midplanes. All external connections are from the rear of the system.



Figure 13 Rear view of IBM FlashSystem 9500 control enclosure

Figure 14 shows a more detailed view of the rear of a canister.

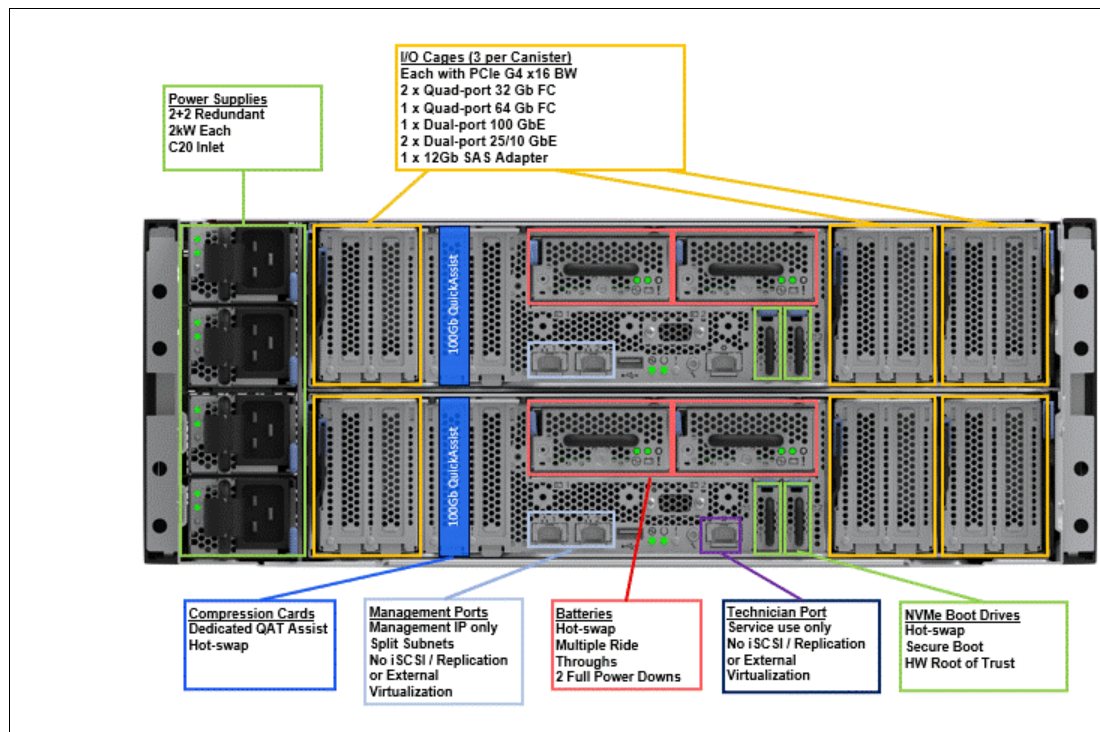


Figure 14 IBM FlashSystem 9500 enclosure detailed rear view

Figure 14 shows the two canisters, which are located one above the other three PCIe G4 interface cages, four power supply units (PSU), the two batteries and boot drives per canister and the USB and management Ethernet ports.

The PCIe adapters plug into the adapter cages and any unused PCIe adapter slots include a blanking plate assembly in place to ensure and maintain the correct air flow for cooling through the canister. A dedicated compression QAT card is also included in each canister.

IBM FlashSystem 9500 NVMe drive options

The IBM FlashSystem 9500 control enclosure supports up to 48 NVMe 2.5-inch drives, which can be the IBM FlashCore Module NVMe type drives or the industry-standard NVMe drives.

With partially populated control enclosures, we have some drive slot plugging rules that must be adhered to, ensuring the best possible operating conditions for the drives.

Figure 15 on page 31 shows the logical NVMe drive placement, starting from the center of the enclosure (slot 12) on the upper 24 slots. Any slots that do not have an NVMe drive present must have a blank filler installed.

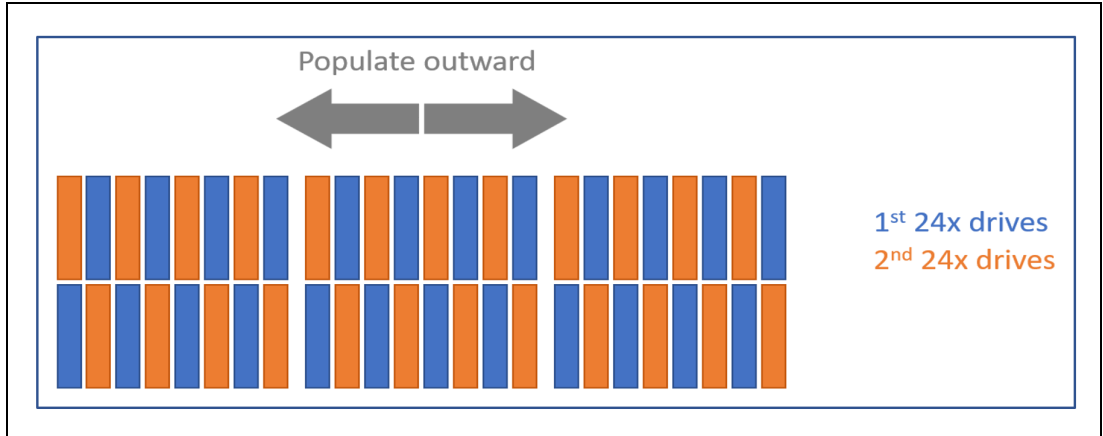


Figure 15 Logical NVMe drive placement

Figure 16 shows the actual drive population with numbering. This shows how the drives are populated from center out, and then distributing them from top and bottom, as the number of drives increase over time.

Note: The layout in Figure 16 has been split at slots 12 and 13 for better clarity on this page, but in reality slots 1 to 24 and slots 25 to 48 are contiguous.

Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9	Slot 10	Slot 11	Slot 12
36	11	34	9	32	7	30	5	28	3	26	1
Slot 25	Slot 26	Slot 27	Slot 28	Slot 29	Slot 30	Slot 31	Slot 32	Slot 33	Slot 34	Slot 35	Slot 36
24	47	22	45	20	43	18	41	16	39	14	37
Slot 13	Slot 14	Slot 15	Slot 16	Slot 17	Slot 18	Slot 19	Slot 20	Slot 21	Slot 22	Slot 23	Slot 24
25	2	27	4	29	6	31	8	33	10	35	12
Slot 37	Slot 38	Slot 39	Slot 40	Slot 41	Slot 42	Slot 43	Slot 44	Slot 45	Slot 46	Slot 47	Slot 48
13	38	15	40	17	42	19	44	21	46	23	48

Figure 16 NVMe drive population with numbering

Supported drive types

The following 2.5-inch (SFF) IBM FlashCore Module NVMe types are supported in IBM FlashSystem 9500 control enclosures:

- ▶ 4.8 TB 2.5-inch NVMe FlashCore Module
- ▶ 9.6 TB 2.5-inch NVMe FlashCore Module
- ▶ 19.2 TB 2.5-inch NVMe FlashCore Module
- ▶ 38.4 TB 2.5-inch NVMe FlashCore Module

The following 2.5-inch (SFF) NVMe industry-standard flash drives are supported in IBM FlashSystem 9500 control enclosures:

- ▶ 1.92 TB 2.5-inch
- ▶ 3.84 TB 2.5-inch
- ▶ 7.68 TB 2.5-inch
- ▶ 15.36 TB 2.5-inch
- ▶ 30.72 TB 2.5-inch

It also supports up to 12 1.6 TB NVMe SCM drives.

The following DRAID rules must be observed:

- ▶ DRAID 1 is recommended for array configurations of 3–5 drives.
- ▶ DRAID 1 and DRAID 6 are recommended for array configurations of six drives, depending on capacity and performance requirements.
- ▶ DRAID 6 is recommended for array configurations with seven or more drives.
- ▶ All flash modules must be the same type and capacity within the same DRAID array.

The following DRAID guidelines are recommended:

- ▶ FlashCore Modules:
 - DRAID 6 (6 drives minimum; 48 drives maximum)
 - DRAID 1 (2 drives minimum; 16 drives maximum)
 - It is a best practice that all drives in a single array be of the same type, technology, and size.
- ▶ Industry standard NVMe drives:
 - DRAID 6 (6 drives minimum; 128 drives maximum)
 - DRAID 1 (2 drives minimum; 16 drives maximum)
 - It is a best practice that all drives in a single array be of the same type, technology, and size.
- ▶ Storage-class memory drives:
 - DRAID 6 (6 drives minimum; 12 drives maximum)
 - DRAID 1 (2 drives minimum; 12 drives maximum)
 - Storage Class Memory drives in the same RAID array must be of the same capacity

All drive slots must be filled regardless of the number of drive modules. If a slot does not have a drive module that is installed, a filler must be used to ensure correct air flow through the enclosure. All drives are dual-port and hot-swappable.

Drives of the same form factor and connector type can be intermixed within an enclosure. All flash modules must be the same type and capacity within the same DRAID array.

Consider the following points:

- ▶ Each IBM FCM contains IBM enhanced flash chips, FPGA chips, an IBM PowerPC® processor, and DRAM devices that are connected to the flash controllers and processor.
- ▶ Cache options from 1.0 TB (512 GB per canister) to 3.0 TB (1.5 TB per canister) per control enclosure.
- ▶ Four 1 Gbps Ethernet ports standard for management and technician connectivity only.
- ▶ 32 Gbps / 64 Gbps FC, 25 Gbps Ethernet, and 100 Gbps Ethernet ports are available for FC and iSCSI connectivity through the use of a PCIe adapter.
- ▶ 12 Gbps SAS ports for expansion enclosure attachments are available by using a PCIe adapter.

The IBM FlashSystem 9500 control enclosure supports FCP with point to point (FC-P2P), arbitrated loop (FC-AL), and switched fabric (FC-SW) topologies. FC interfaces can be configured as N_port or NL_port types.

The FC port is configured in NPIV mode; therefore, the user must check that this configuration is defined for their installation. If not, this configuration must be changed from the default when the SAN infrastructure is set up. Full active-active multipathing across all interfaces also is supported, although host software support for this function can vary.

The IBM FlashSystem 9500 control enclosure also includes the following expansion features:

- ▶ Capability for adding into clustered systems with one more IBM FlashSystem 9500 control enclosure.
- ▶ Up to six SAS attached expansion enclosures are supported per IBM FlashSystem 9500 controller pair, which provides up to 144 SSDs with IBM FlashSystem 9000 Expansion Model AFF.
- ▶ Up to two IBM FlashSystem 9000 LFF HD model A9F expansion enclosures are supported per IBM FlashSystem 9500 control enclosure pair, which provides up to 184 SSD drives.

For more information about mixing expansion enclosures and the maximum configurations that are allowed, see [V8.7.0.x Configuration Limits and Restrictions for IBM FlashSystem 9500](#).

IBM FlashSystem 9500 utility model UH8

IBM FlashSystem 9500 utility models UH8 provide a variable capacity storage offering. These models also offer a fixed capacity, with a base subscription of approximately 35% of the total capacity.

IBM Storage Insights monitors the system and capacity that is used. The amount that is billed is based on the capacity-used basis. With this billing structure, you can increase or decrease usage, and pay for the configured capacity only.

IBM FlashSystem utility models are provided for customers who can benefit from a variable capacity system, where billing is based on provisioned space only. The hardware is leased through IBM Global Finance on a three-year lease. Leased hardware entitles the customer to use approximately 30—40% of the total system capacity at no extra cost depending on the specific terms of the contract. If storage must increase beyond that initial capacity, the extra usage is billed based on the average daily provisioned capacity per terabyte per month on a quarterly basis.

Example: A total system capacity of 115 TB

A customer has an IBM FlashSystem 9500 utility model with 12 9.6 TB NVMe drives for a total system capacity of 115 TB. The base subscription for such a system is 40.25 TB. No extra billing occurs during the months where the average daily usage is less than 40.25 TB.

The system monitors daily provisioned capacity and averages those daily usage rates over the month. The result is the average daily usage for the month.

If a customer uses 45 TB, 42.5 TB, and 50 TB in three consecutive months, IBM Storage Insights calculates the overage (see Table 2), rounding to the nearest terabyte.

Table 2 Billing calculations based on customer usage

Average daily	Base	Overage	To be billed
45 TB	40.25 TB	4.75 TB	5 TB
42.5 TB	40.25 TB	2.25 TB	2 TB
50 TB	40.25 TB	9.75 TB	10 TB

The total capacity that is billed at the end of the quarter is 17 TB per quarter in this example.

Flash drive expansions can be ordered with the system in all supported configurations. Table 3 lists the feature codes that are associated with the UH6 utility model billing.

Table 3 9500 UH8 utility model billing feature codes

Feature code	Description
#AE00	Variable Usage 1 TB/month
#AE01	Variable Usage 10 TB/month
#AE02	Variable Usage 100 TB/month

These features are used to purchase the variable capacity that is used in the utility models. The features #AE00, #AE01, and #AE02 provide terabytes of capacity beyond the base subscription on the system. Usage is based on the average capacity that is used per month. The total of the prior three months' usage is totaled, and the corresponding number of #AE00, #AE01, and #AE02 features ordered quarterly.

Billing

The local project office compiles the usage information from IBM Storage Insights on a quarterly basis. This data is compared to the base system capacity subscription. Any provisioned capacity beyond that base subscription is billed per terabyte, per month on a quarterly basis.

The calculated usage is based on the average use over a month. In a highly variable environment, such as managed or cloud service providers, this subscription enables the system to be used only as much as is necessary during any month. Usage can increase or decrease, and is billed.

Provisioned capacity is considered capacity that is reserved by the system. In thick-provisioned environments, this capacity is the capacity that is allocated to a host, whether it has data that is written.

For thin-provisioned environments, the data that is written is used because of the different ways in which thick-provisioning and thin-provisioning use flash drive space.

IBM FlashSystem 9500 supported expansion enclosures

This section describes the supported expansion enclosures of the IBM FlashSystem 9500.

IBM FlashSystem 9000 SFF Expansion Enclosure Model AFF

IBM FlashSystem 9000 Expansion Enclosure Model AFF holds up to 24 2.5-inch SAS flash drives in a 2U 19-inch rack mount enclosure. An intermix of capacity drives is allowed in any drive slot, and up to three AFF enclosures can be attached to the control enclosure to a total of 72 drives maximum.

Figure 17 shows the front view of the IBM FlashSystem 9000 Expansion Enclosure Model AFF.



Figure 17 IBM FlashSystem 9000 Expansion Enclosure Model AFF

Figure 18 shows the rear view of IBM FlashSystem 9000 SFF Expansion Enclosure Model AFF.



Figure 18 Rear view of IBM FlashSystem 9000 Expansion Enclosure Model AFF

Model AFF: Supports up to 24 SFF SSD 2.5-inch drives that can be formatted as DRAID 1 or DRAID 6. It supports the following drive types:

- ▶ 1.92 TB 2.5-inch flash drive
- ▶ 3.84 TB 2.5-inch flash drive
- ▶ 7.68 TB 2.5-inch flash drive
- ▶ 15.36 TB 2.5-inch flash drive
- ▶ 30.72 TB 2.5-inch flash drive

The AFF model of IBM FlashSystem 9000 SFF Expansion Enclosure includes the following features:

- ▶ Up to 24 2.5-inch flash SSDs or HDDs
- ▶ Two expansion canisters
- ▶ 12 Gbps SAS ports for attachment to the IBM FlashSystem 9500 control enclosures
- ▶ 2U 19-inch rack-mount enclosure with AC power supplies

IBM FlashSystem 9000 LFF Expansion Enclosure Model A9F

IBM FlashSystem 9500 LFF Expansion Enclosure Model A9F includes the following features:

- ▶ Up to 92 3.5-inch drives that are top-loaded into drive slots of the expansion enclosure
- ▶ 5U 19-inch rack-mount enclosure with slide rail and cable management assembly

- ▶ High-performance SSD support, which is available in 1.92 TB, 3.84 TB, 7.66 TB, 15.36 TB, and 30.72 TB capacity versions
- ▶ Redundant 200–240 V AC power supplies (new C19/C20 PDU power cord required)

Figure 19 shows the front view of the IBM FlashSystem 9000 Expansion Enclosure Model A9F.



Figure 19 IBM FlashSystem 9000 Expansion Enclosure Model A9F

Model A9F: Supports up to 92 LFF 2.5-inch SSDs in a 3.5-inch carrier. It can be formatted as DRAID 1 or DRAID 6 and supports the following drive types:

- ▶ 1.92 TB 12 Gbps SAS flash drive (2.5-inch and 3.5-inch form factor features)
- ▶ 3.84 TB 12 Gbps SAS flash drive (2.5-inch and 3.5-inch form factor features)
- ▶ 7.68 TB 12 Gbps SAS flash drive (2.5-inch and 3.5-inch form factor features)
- ▶ 15.36 TB 12 Gbps SAS flash drive (2.5-inch and 3.5-inch form factor features)
- ▶ 30.72 TB 12 Gbps SAS flash drive (2.5-inch and 3.5-inch form factor features)

Figure 20 shows the rear view of IBM FlashSystem 9000 LFF Model A9F Expansion Enclosure.




Figure 20 Rear view of IBM FlashSystem 9000 LFF Model A9F Expansion Enclosure

Scalability and performance

Multiple expansion enclosures are supported for each IBM FlashSystem 9500 control enclosure, which provides up to 72 drives with Expansion Enclosure Model AFF and up to 92 drives with Expansion Enclosure Model A9F.

On each SAS chain, the system can support up to a SAS chain weight of three:

- ▶ Each 4666-A9F expansion enclosure adds a value of 2.5 to the SAS chain weight.
- ▶ Each 4666-AFF expansion enclosure adds a value of 1 to the SAS chain weight.

For example, each of the following expansion enclosure configurations has a total SAS weight of three:

- ▶ Three 4666-AFF expansion enclosures per SAS chain.
- ▶ One 4666-A9F expansion enclosure per SAS chain.

The IBM FlashSystem 9500 system includes the following scalability and performance features:

- ▶ Up to 1.84 PB usable and 5.52 PB effective flash capacity in only 4U with 3:1 data reduction.
- ▶ Up to 3.68 PB usable and 18.4 PB maximum flash capacity in only 8U with 5:1 data reduction.
- ▶ Extra scalability through Expansion Enclosures Models AFF and A9F, and increased raw capacity of up to a maximum of 32 PB.

- ▶ The IBM FlashSystem 9500 system can deliver up to 6.3 million IOPS, with latency down to 50 microseconds (μ s) and bandwidth up to 65 GBps from a single system.
- ▶ A clustered IBM FlashSystem 9500 system (including the IBM FlashSystem 9500R system) can scale linearly and deliver 12.6 million IOPS and 130 GBps on a fully configured IBM FlashSystem 9500 cluster with two enclosures.

Advanced functions

The IBM FlashSystem 9500 system provides the following advanced functions:

- ▶ Policy-based HA for IBM FlashSystem 9500
- ▶ IBM Flash Grid
- ▶ Safeguarded Snapshot
- ▶ IBM FlashSystem Cyber Vault
- ▶ Increased security features
- ▶ NVMe over Fabrics
- ▶ Portsets
- ▶ IP quorum
- ▶ Data reduction tools
- ▶ N-Port ID virtualization support
- ▶ VMware integration
- ▶ External virtualization

Policy-based HA for IBM FlashSystem 9500

Policy-based HA for IBM FlashSystem 9500 is available with IBM Storage Virtualize V8.7.0 and later. This ensures continuous application data access even in case of primary storage system failures. Policy-based HA maintains synchronized data copies on a peer system, allowing applications to access data from either system. This active/active solution enables simultaneous I/O to both copies while maintaining data synchronization. Synchronous replication guarantees data consistency between production volumes. Volume groups manage consistency across application-dependent volumes. Storage partitions facilitate easy management of highly available resources by grouping related volume groups, hosts, and mappings. To prevent conflicts, IP Quorum determines the active management system. You can flexibly add or modify volumes, volume groups, hosts, and mappings within existing partitions or create new ones. See Figure 21 on page 39.

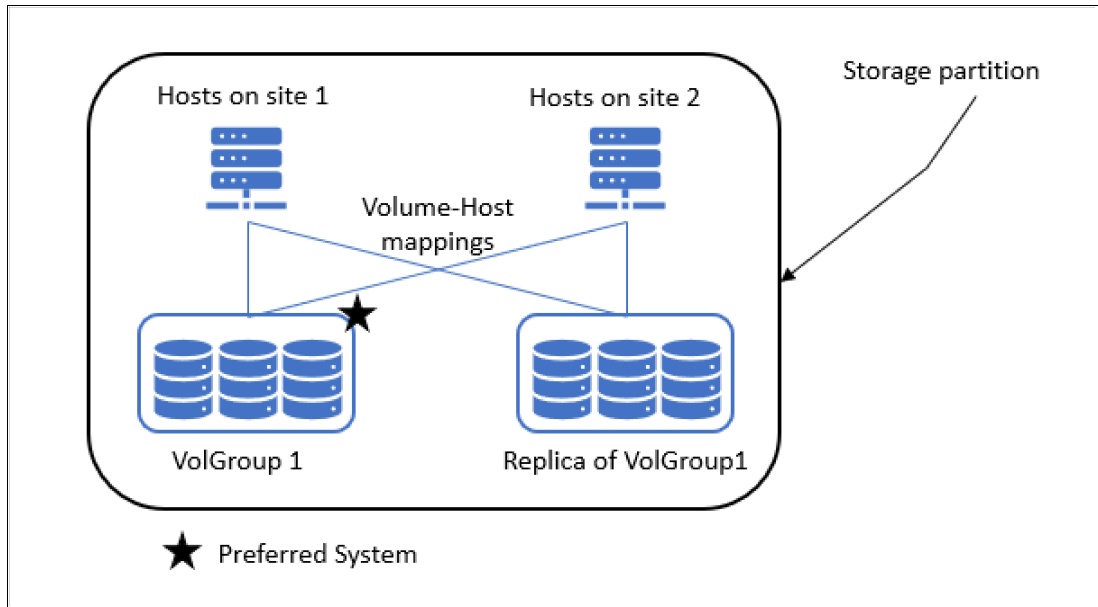


Figure 21 Storage partitions

Storage partitions simplify HA management by automating host mapping to volume group replicas across both systems. Hosts seamlessly access volumes with identical UIDs. To maintain consistency, HA replicates not only data but also storage partition configurations, including host definitions and mappings. Configuration changes are typically managed on the preferred system and synchronized to both. The preferred system is usually the system where the partition is managed and remains accessible during disconnections. Hosts can access volumes through the preferred system or, if host locations are set, directly from the local storage system. Volume groups, storage partitions, and policies are centrally managed on the active management system, which can fail over in case of an outage. Local hosts automatically switch to the remote system using ALUA multipathing during local site failures.

For more information on policy-based HA refer to *Ensuring Business Continuity: A Practical Guide to Policy-Based Replication and Policy-Based High Availability for IBM Storage Virtualize Systems*, [SG24-8569](#).

IBM Flash Grid

IBM Storage Virtualize, with storage partitions and volume groups, separates business continuity requirements like HA and replication from hardware systems, enabling multiple software-defined virtual storage systems within a single FlashSystem deployment. Flash Grid empowers users to create federated, scalable clusters of independent storage devices, surpassing traditional scale-out limitations. It offers linear performance, capacity, and resource growth with up to eight systems. Unlike previous "per I/O group" clustering, Flash Grid focuses on system-level scale-out, simplifying management and reducing hardware compatibility constraints. This approach enables easier hardware replacement, capacity balancing, and non-disruptive application data migration through storage partition migration. Clients can aggregate IBM FlashSystem or SVC systems into a single, highly available, manageable storage grid.

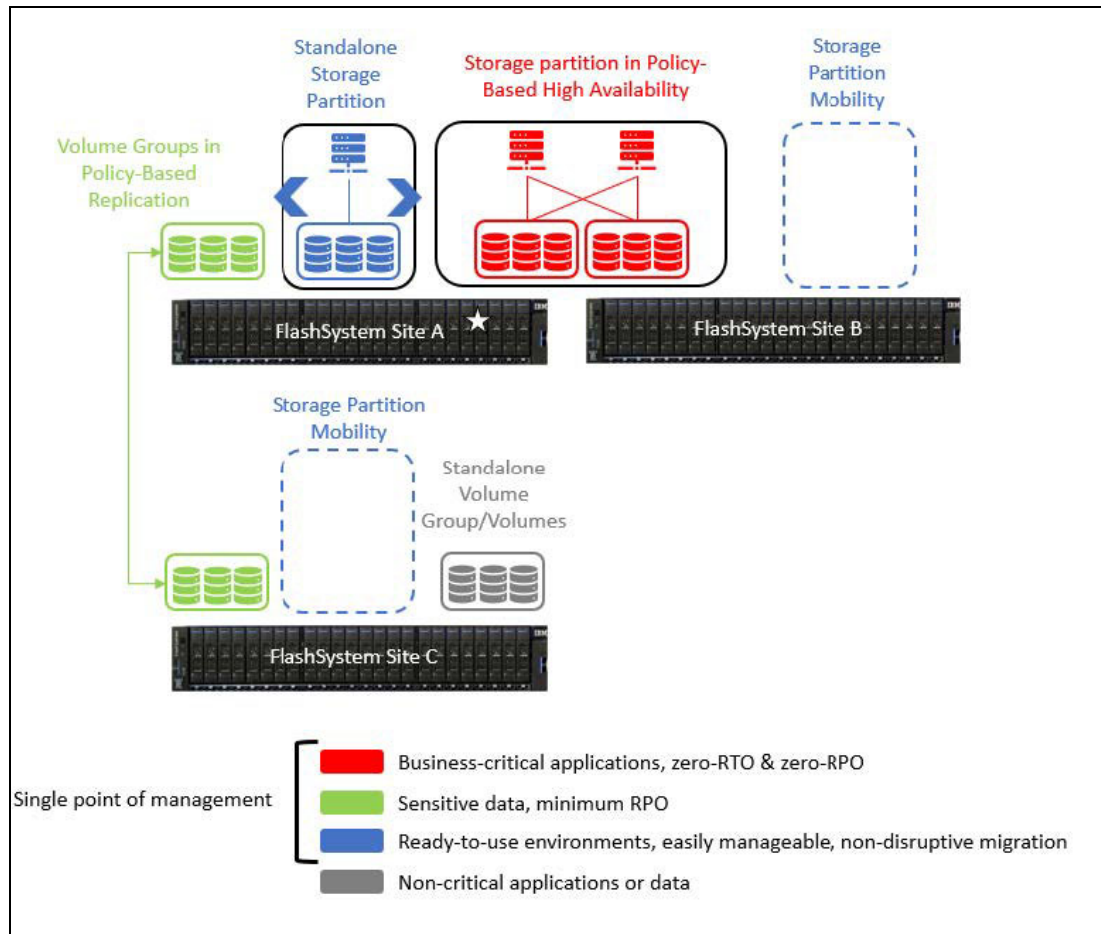


Figure 22 IBM Flash Grid concept

Once storage partitions are configured, they can easily be moved from a system to another, manually balanced by users over several systems and sites. They can also be stretched over two sites for high availability. See Figure 22. At the time of this writing, Flash Grid features (partitions mobility) are manageable with IBM Storage Insights Pro only.

It is possible to use the CLI to create a Flash Grid and add or remove systems in a Flash Grid.

Safeguarded Snapshot

IBM Safeguarded Snapshot (previously called Safeguarded Copy) regularly creates isolated, separated from server, immutable snapshots of data to help protect against cyberattacks, malware, acts of disgruntled employees, and other data corruption. Because Safeguarded snapshots are on the same FlashSystem storage as operational data, recovery is faster than restoring from copies stored separately.

The Safeguarded child-pool capability on the IBM Storage Virtualize family of products was introduced in an earlier version. The IBM Storage Virtualize family of products includes IBM FlashSystem products, IBM SAN Volume Controller, and IBM Storage Virtualize for Public Cloud.

HA configurations mitigate against physical component failure and provide small Recovery Point Objective (RPO) and Recover Time Objective (RTO). Continuous protection and operation occur if a component fails. IBM Storage Virtualize HA configurations, such as PBHA, provide such protection.

DR is designed to have a slightly higher RPO and RTO. Therefore, DR is in a position to provide an airgap to protect against corruption to the data that is replicated immediately in an IBM Storage Virtualize HA configuration, such as policy-based high availability.

However, because DR solutions at the storage layer are replicating only data, any logical corruption eventually is replicated to the DR site as well.

One way to provide adequate protection against logical data corruption is to take periodic snapshots of the data and to have that data stored in a non-modifiable state that is inaccessible to administrators, servers, and applications. These Safeguarded snapshots can then serve as recovery points from which the data can be restored to a pre-corruption state. The corruption can be the result of an errant batch job, a disgruntled employee, or a ransomware attack.

For more information about Safeguarded Snapshot, see *Data Resiliency Designs: A Deep Dive into IBM Storage Safeguarded Snapshots*, [REDP-5737](#).

IBM FlashSystem Cyber Vault

The IBM FlashSystem Cyber Vault solution complements IBM Safeguarded Snapshot. IBM FlashSystem Cyber Vault automatically scans the copies that are created regularly by Safeguarded Snapshot, looking for signs of data corruption that are introduced by malware or ransomware. Refer to “IBM Cyber Vault solution” on page 51 for more information.

Increased security features

IBM Storage Virtualize enhances security with multi-factor authentication (MFA), requiring multiple forms of verification for login. This strengthens protection against unauthorized access, especially in cloud environments. Integrated with IBM Security Verify, MFA provides granular control over user access and simplifies monitoring. Additionally, the software supports single sign-on (SSO) via Microsoft Active Directory Federation Services (AD FS) for efficient user logins.

For more information, see *IBM Storage Virtualize, IBM FlashSystem, and IBM SAN Volume Controller Security Feature Checklist*, [REDP-5678](#).

NVMe over Fabrics

NVMe is a high-speed interface optimized for flash storage, designed to deliver low latency and high throughput. This standard enables direct communication between the host system and flash-based storage devices, bypassing traditional bottlenecks.

The NVMe protocol is an interface specification for communicating with storage devices and it is functionally the same as other protocols, such as SATA and SAS. However, the NVMe interface was designed for fast storage media, such as flash-based SSDs and low-latency nonvolatile storage technologies.

Typically, NVMe storage devices are directly attached to a host system over a PCIe bus, and the NVMe controller is contained in the storage device. This architecture eliminates the need

for another I/O controller between the CPU and the storage device and results in reduced latency, throughput scalability, and simpler system designs.

This design was implemented in the IBM FlashSystem 9500, with its dual-ported PCIe-based NVMe drives.

However, the PCIe bus has a limit to the number of NVMe drives that can be attached to a host over the bus to a few tens of devices. Also, because the maximum length of PCIe cabling is a few meters, the flexibility of deploying PCIe NVMe drives outside the host server is limited and data center level scalability is not feasible.

NVMe over Fabrics (NVMe-oF) overcomes the limitations of the SCSI protocol and the limited number of concurrent queues by extending the benefits of low latency and high efficiency of the NVMe technology across network fabrics. This feature supports sharing of NVMe storage at a large scale of 100s or 1000s of devices and over distance.

Figure 23 on page 42 shows that the NVMe architecture supports many different network fabric technologies.

The NVMe transport layer can be mapped to different network fabric technologies.

As of this writing, the NVMe transport layer supports the NVMe over Fabrics by using Fibre Channel (referred to as FC-NVMe or NVMeFC) main fabric transport.

FC-NVMe uses FCP as its transport mechanism, which is similar to RDMA, because it places the data transfer in control of the target and transfers data direct from host memory. In addition, FC-NVMe allows for a host to send commands and data together (first burst), which eliminates the first data read by the target and provides better performance at distances.

Figure 23 shows that the NVMe architecture supports many different network fabric technologies.

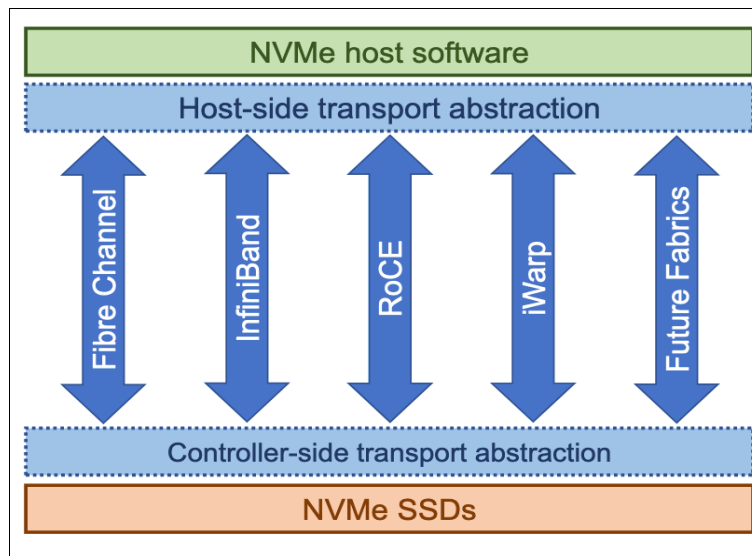


Figure 23 NVMe technology support options

The NVMe transport layer can be mapped to different network fabric technologies.

At the time of writing, the NVMe transport layer supports the NVMe over Fabrics by using FC (referred to as *FC-NVMe* or *NVMeFC*) main fabric transport.

FC-NVMe uses FC Protocol (FCP) as its transport mechanism, which places the data transfer in control of the target and transfers data direct from host memory, which is similar to RDMA. In addition, FC-NVMe allows for a host to send commands and data together (first burst), which eliminates the first data “read” by the target and provides better performance at distances.

The NVMe over Fabrics interface uses the same model of submission and completion queues as PCIe NVMe. As such, it maintains the same asynchronous submission and completion model and achieves similar benefits in terms of latency, efficiency, and scalability as the NVMe technology because of the shortened code paths and lockless concurrency in multi-core environments.

IBM FlashSystem 9500 supports NVMe protocols by way of the various PCIe adapters that can be ordered and installed in the control enclosure. For more information about these options, see “Host I/O connectivity and expansion enclosure adapters” on page 57.

For more information about NVMeoF protocols and architecture, see [IBM Storage and the NVM Express Revolution, REDP-5437](#).

Portsets

IBM FlashSystem 9500 is designed for multi-tenancy where multiple clients can share storage resources. In multi-tenancy environments, it might be a requirement that clients use storage ports that are connected to different networks to isolate traffic from other clients.

Portsets are groupings of logical addresses that are associated with the specific traffic types. The system supports IP portsets for host attachment, back-end storage connectivity, and IP replication traffic. In addition, FC portsets can be configured for host attachment.

After you create portsets and assign IP addresses, you can assign hosts, host clusters, and IP partnerships to the portset for those traffic types.

IP quorum

For lower implementation and operation costs for a HA solution, you can use an IP quorum to use lower-cost IP network-attached hosts as a quorum disk. Typically the quorum devices at site 1 and site 2 are quorum disks and the quorum device at site 3 is an IP quorum application. However, the system can be configured to use either quorum disks or IP quorum applications at any site. This configuration ensures that a quorum device is always available, even after a single-site failure. When you are using an IP quorum application at a third site, you can configure a preference for which site continues operation if there is a loss of connectivity between the two sites. If only one site runs critical applications, you can configure this site as preferred. If a preferred site is configured and a failure causes an outage at the preferred site, the other site wins the tie-break and continues operating and processing I/O requests. A quorum device also is used to store a backup copy of important system configuration data. Just over 256 MB is reserved for this purpose on each quorum device.

A system can have only one active quorum device that is used for a tie-break situation. However, the system uses up to three quorum devices to record a backup of system configuration data to be used if a disaster occurs. The system automatically selects one quorum device to be the active quorum device.

For more information about configuring quorum disks within two or three sites, see [Configuring quorum](#).

Note: FC over IP (FCIP) connectivity is not supported between nodes when a PBHA system is configured without the use of inter-switch links (ISLs).

Data reduction tools

Compression and de-duplication are a key part of the IBM FlashSystem 9500 system. The IBM Comprestimator and the IBM Data Reduction Estimator Tool are the key sizing tools that are used to estimate how much capacity savings that a client can expect. The IBM Comprestimator and the IBM Data Reduction Estimator Tool can recognize the patterns of the client data, and estimate the compressibility of data per volume.

IBM FlashSystem models are supported by the IBM Comprestimator, which is available as a stand-alone tool and is found in the FlashSystem 9500 GUI. The Data Reduction Estimator Tool (DRET) is a host-based application that estimates the amount of compression and de-duplication on the IBM FlashSystem 9500 system for specific workloads.

For more information about DRET, see [IBM Data Reduction Estimator Tool \(DRET\) for SVC, Storwize and FlashSystem products](#).

For more information about the IBM Comprestimator, see [IBM FlashSystem Comprestimator](#).

Choose your data reduction approach and use the tools to estimate the amount of usable storage that is required by reviewing the following tools.

1. From the main menu, click **Volumes** → **Volumes** and then select one volume.
2. Right-click the volume to open the menu.
3. From the menu, select **Capacity Savings** → **Estimate Compression Saving**.

Figure 24 shows how to start the Estimate Compression Saving option from the GUI:

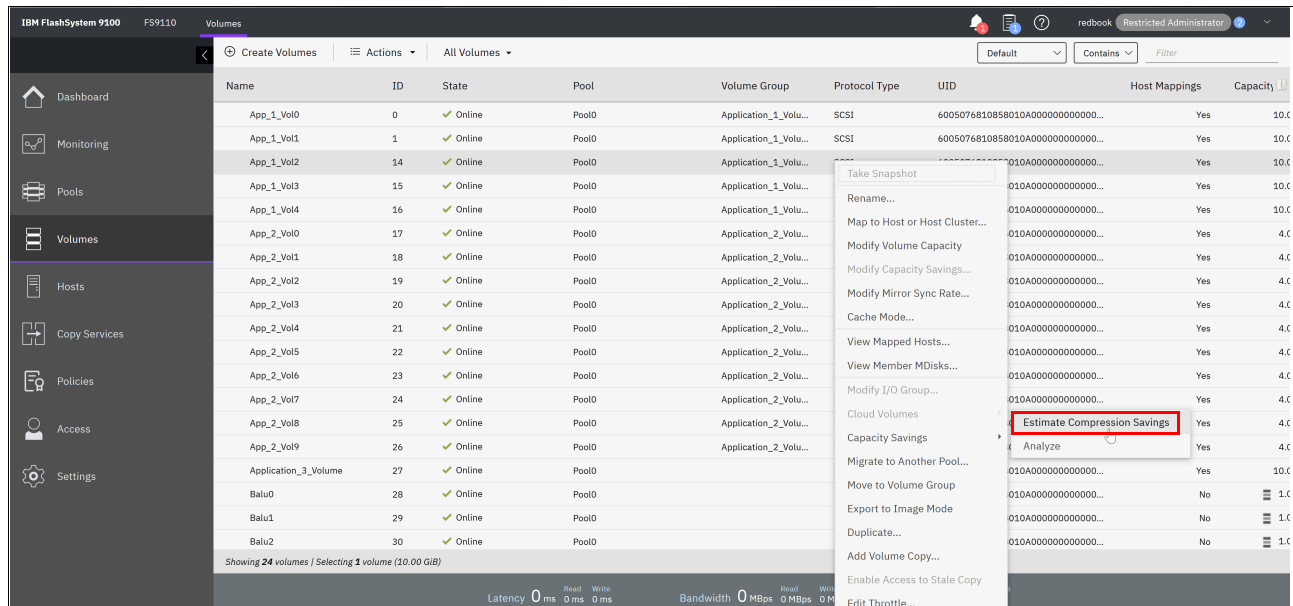


Figure 24 Estimate Compression Savings from the GUI

To estimate compression and savings on FCM and DRP use the following methods:

- ▶ IBM FCM compression:
 - Use the FCM option.
 - Do not use the Estimate Compression Saving option in the GUI to calculate the IBM FlashCore Module savings.
- ▶ DRP compression:
 - Use the DRP option.
 - Workloads that are running on any IBM Storage Virtualize devices can use the Estimate Compression Saving option in the GUI.
- ▶ Data Reduction Pool compression and deduplication:
 - IBM Comprestimator and Data Reduction Estimator Tool shows the savings for thin-provisioning, compression, and deduplication.
 - IBM Comprestimator and Data Reduction Estimator Tool reads entire volumes to identify de-duplicated data, so it takes longer to run.

For more information about data reduction pool compression and setup, see [Introduction and Implementation of Data Reduction Pools and Deduplication, SG24-8430](#).

N-Port ID virtualization support

IBM Storage Virtualize software offers N-port ID virtualization (NPIV) support. Use NPIV to virtualize WWPNs, which increases redundancy during firmware updates and scheduled maintenance where WWPNs transparently move to the controller that is not being maintained. As a consequence, FC-attached hosts experience zero path reduction during controller outages.

Important: On the IBM FlashSystem 9500 system, NPIV is enabled by default; therefore, if the customer does not want to use NPIV, the customer must disable NPIV before configuring FC ports for host communications.

VMware integration

IBM Storage Virtualize software includes the following features, which enable tight integration with VMware:

- ▶ vCenter plug-in enables monitoring and self-service provisioning of the system from within VMware vCenter.
- ▶ vStorage APIs for Array Integration (VAAI) support supports hardware-accelerated virtual machine (VM) copy and migration and hardware-accelerated VM initiation, and accelerates VMware Virtual Machine File System (VMFS).
- ▶ Microsoft Windows System Resource Manager (SRM) for VMware Site Recovery Manager supports automated storage and host failover, failover testing, and failback.
- ▶ VVOL integration for better usability:
 - The migration of space-efficient volumes between storage containers maintains the space efficiency of volumes.
 - Cloning a VM achieves a full independent set of VVOLs.
 - Resiliency is also improved for VMs if volumes start running out of space.
- ▶ Support for version 2 metadata volume for VMware Virtual Volumes (vVols)

Before the availability of VVOLs, a VM in a VMware environment was presented as a VMware disk that was called a VMware Virtual Machine Disk (VMDK). This file represented a physical disk to the VM, which can be accessed by the operating system that is installed on the VM in the same way that a physical volume on a regular server was installed.

The VMDK file was placed in a file system that is called VMFS, which is hosted by a standard volume (LUN). For example, it can be implemented on an external storage system, such as the IBM FlashSystem 9500 system. With the availability of the VVOL technology, each VM disk can now be mapped to an external storage volume (for example, an IBM FlashSystem 9500 volume).

With VVOL, the IBM FlashSystem 9500 solution recognizes individual VMDK files. Therefore, data operations, such as snapshot and replication, can be performed directly by the IBM FlashSystem 9500 system at the VMDK level rather than the entire VMFS data store.

Note: The integration of VVOL with the IBM FlashSystem 9500 system is based on the VMware APIs for Storage Awareness (VASA). IBM support for VASA is delivered as part of the embedded VASA Provider starting with IBM Storage Virtualize 8.5.1.0.

For more information about VMware integration, see *IBM Storage Virtualize and VMware: Integrations, Implementation and Best Practices*, [SG24-8549](#).

External virtualization

The IBM FlashSystem 9500 can manage external storage arrays through virtualization, offering features like thin provisioning, data replication, and easy data migration. By incorporating external storage into a unified pool, it simplifies management and enables efficient capacity utilization. This approach also protects applications from underlying storage changes, minimizing downtime. Additional licensing is required for external storage virtualization.

Scaling up and scaling out

The IBM FlashSystem 9500 system's scalable architecture enables flash capacity to be added (scaled up) to support multiple applications by adding expansion enclosures. Scale out expands NVMe capacity, with the processing power to use that capacity to linearly scale IOPS and bandwidth. As a result, your organization can gain a competitive advantage through a more flexible, responsive, and efficient storage environment.

The IBM FlashSystem 9500 system has the following flexible scalability configuration options:

- ▶ Base configuration
- ▶ Scale out: Add Control Enclosures and capacity
- ▶ Scale up: Add capacity

As previously mentioned, each IBM FlashSystem 9500 system has two canisters (sometimes also known as *nodes* or *controllers*). Each canister contains the CPUs, cache memory, PCIe cards, and other hardware to communicate to the NVMe drives and connected host systems. These two canisters are housed in a chassis that is known as a *control enclosure*. The pair of canisters in a control enclosure is also called an I/O group.

It is possible to connect up to a maximum of two control enclosures of the same machine type to form a cluster.

Note: The FS9500 now has two machine types: the 4666 and the new 4983. If you are going to cluster these types of systems, then the machine types *must* match. You cannot cluster a 4666 with a 4983 and vice versa.

Table 4 lists the values of the CPUs, cache memory, raw NVMe drive size, and PCIe adapters cards for each number of control enclosures.

Table 4 Maximum values for each canister and control enclosure

Element Description	Number of CPUs	Maximum cache memory (GB)	Max NVMe 48 drive capacity raw (TB)	Inclusive Onboard Ethernet ports	Max PCIe ports (FC, Ethernet)
One Canister	2	1536	N/A	2	24
One Enclosure	4	3072	922	4	48
Two Enclosures	8	6144	1844	8	96

A basic configuration of an IBM FlashSystem 9500 storage platform consists of one IBM FlashSystem 9500 control enclosure. For a balanced increase of performance and scale, up to two IBM FlashSystem 9500 control enclosures can be clustered into a single storage system, which multiplies performance and capacity with each addition.

IBM is also offering a pre-configured clustered system, which is known as the IBM FlashSystem 9500R system. This configuration features two IBM FlashSystem 9500 control enclosures, FC switches, and optional expansion enclosures in a rack that are approved by IBM and ready for the IBM Engineer and Lab Services to install and start.

Note: For more information about the FS 9500R rack mount offering, see "IBM FlashSystem 9500R rack" on page 9.

The clustered IBM FlashSystem 9500 system can include dedicated internal FC switches for internal communications. However, other methods are available to configure the switches and ports to provide performance improvements.

For more information about how to restrict the FC ports for control enclosure inter-node connections and host connections, see [Fibre Channel port masking](#).

With the scalable IBM FlashSystem 9500 configurations, you can add IBM FlashSystem 9500 control enclosures to the storage system. If 1.84 TB usable capacity and up to 9.21 PB maximum effective capacity from two control enclosures is not enough, capacity can be increased by adding SAS expansion enclosures with SSDs.

The IBM FlashSystem 9500 system is expandable up to the architectural limit of IBM Storage Virtualize of 32 PB by using SSDs in expansion enclosures.

Maximum capacity configurations

Table 5 lists the maximum raw capacity, usable capacity, and maximum effective capacity for scalable configurations that are built on the IBM FlashSystem 9500 control enclosures.

Table 5 IBM FlashSystem 9500 maximum usable and effective capacity.

Number of control enclosures	Maximum usable capacity by using DRAID 6 with 48 drives (terabytes)	Maximum effective capacity (terabytes) with inline IBM FlashCore Modules Drive compression @ 3:1 ratio	Maximum effective capacity (terabytes) with software data reduction @ 5:1 ratio
One control enclosure	1834	5502	9216
Two control enclosures	3668	11004	18432

Note: IBM FlashCore Module inline compression is hardware-based and operates at line speed, which has no effect on performance.

Capacity expansion through expansion enclosures

The IBM FlashSystem 9500 system supports the addition of expansion enclosures to extend capacity.

For more information about the available drive module options that can be installed in expansion enclosures, see “IBM FlashSystem 9500 hardware component overview” on page 27.

Expansion Enclosure Model A9F

HD Enclosure Model A9F accepts only SAS SSDs. With this tiering option, 1 expansion enclosure can be used per control enclosure with a maximum capacity of 2.82 PB when using 30.72 TB flash drives.

Note: To support SAS-attached expansion enclosures, an AHBA - SAS Enclosure Attachment adapter must be installed in the control enclosure of the IBM FlashSystem 9500 system.

Expansion Enclosure Model AFF

IBM FlashSystem 9000 SFF Expansion Enclosure Model AFF offers new tiering options with flash drives. Each SFF Expansion Enclosure supports up to 24 2.5-inch flash drives.

Up to 3 expansion enclosures are supported per IBM FlashSystem 9500 control enclosure, which provides up to 72 drives with up to 2.21 PB of SAS SSD capacity.

Note: The IBM Storage Virtualize maximum storage limit is 32 PB.

RAID types

The IBM FlashSystem 9500 system supports the following DRAID levels:

- ▶ DRAID 6 for all drive types - SCM Class drives, FCMs, and Industry-standard flash drives.
- ▶ DRAID 1 arrays contain 2–6 drives initially, and can be expanded up to 16 drives of the same capacity.
- ▶ It is a best practice that all drives in a single array be of the same type, technology, and size.
- ▶ Only one DRAID array per pool.

The IBM FlashSystem 9500 does not support traditional RAID.

For more information about the types of supported RAID configurations, see [Planning for array configurations](#).

Manageability and security

The IBM FlashSystem 9500 system offers many improvements on the manageability and security features as compared to previous models. In general, the manageability and security of the new range of systems greatly improved.

The following list describes manageability and security changes that were implemented on the FlashSystem 9500 version 8.7:

- ▶ Safeguarded copy, a technology that is derived from the technology that is used in the IBM DS8000® line. It consists of a copy that cannot be changed or deleted after it is created. Storing sensitive copies in immutable storage, cloud environments, or offline write-once read many (WORM) tape devices is also an option to provide physical air-gap protection.
- ▶ Advanced security for data at rest with hardware-accelerated AES-XTS 256 encryption.
- ▶ IBM obtained FIPS 140-3 Level 1 certification for generation 4 IBM FCMs in the IBM FlashSystem 9500. This level of certification is fully supported by the FCM4 cards in the system.
- ▶ A GUI and a CLI to manage the IBM FlashSystem 9500 control enclosures and the IBM FlashSystem 9500 expansion enclosures are available. The GUI is supported in any supported browser.

The IBM FlashSystem 9500 CLI supports a collection of commands that you can use to manage the IBM FlashSystem 9500 system. In the 9500 system, the GUI and CLI support MFA for increased security.

- ▶ The IBM FlashSystem 9500 system supports the Representational State Transfer (REST) model API. The REST API consists of command targets that are used to retrieve system information and to create, modify, and delete system resources.
- ▶ SRA provides a secure connection for IBM Remote Support, which can perform remote troubleshooting and code load, and obtain diagnostic logs.
- ▶ Email and SNMP alerts are part of the standard alerting options.
- ▶ Syslog redirect to send system log messages to another host.
- ▶ Combined password and SSH key authentication are now supported as a first factor for local users.
- ▶ Restricts methods of access to the system at a user group level. GUI, CLI, and Rest access can now be restricted (Role Based Access Control).
- ▶ Log-in grace time and session timeout duration can now be configured according to the customer's policies.
- ▶ Multifactor authentication (MFA):
 - MFA with IBM Security Verify

With IBM Security Verify, security administrators can configure the system as an application that requires two factors for users and user groups to access the system with either the management GUI or CLI. MFA can be used to protect both local users, including superuser, and remote users. Remote users are users who are defined on a remote LDAP server. For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify. For more information, see IBM Security Verify Bridge for Directory Sync in the IBM Security Verify documentation.

- MFA with Duo Security

The FlashSystem 9500 integrates with Duo Security, which is a cloud-based identity and access management (IAM) service provider. You can manage every aspect of your Duo two-factor authentication system from the Duo Admin Panel including creating and managing applications, enrolling and activating users, issuing and managing SMS pass codes and bypass codes, managing mobile devices, fine-tuning the user experience of your Duo installation.

For more information, see [Multifactor authentication](#).

- ▶ For the IBM FlashSystem 9500, features were added to enable a hardware root of trust to protect against physical access attacks and to prevent running foreign code. Attempting to modify the code on the boot drive is *not* possible.

When an attacker attempts to replace the boot drive with their own, the UEFI secure boot detects that effort and rejects the boot drive so that the system does not boot. Therefore, a customer can have confidence that if code is running on an IBM FlashSystem 9500, the code is running from an encrypted disk and was signed by IBM.

- ▶ An interactive management GUI with excellent management and monitoring features.
- ▶ Native scheduler for snapshots. External software is no longer required for creating and managing crash-consistent copies of data.

For application-consistent copies, an external (third party) orchestrator with application awareness is still needed.

IBM Cyber Vault solution

IBM Cyber Vault is a combination of hardware and software that provides a comprehensive approach to cyber resiliency.

IBM Cyber Vault includes the following pillars:

Safeguarded snapshots

Immutable copies of data, which is created by IBM Safeguarded snapshots.

Pro-active monitoring

By using IBM Storage Insights (a recommended service from IBM), with which early warnings of a possible attack are sent. Integration with IBM's QRadar® is also recommended.

Rapid Recovery

The monitoring of the Cyber Vault snapshots that are created by Safeguarded copies is a continuous process. Monitoring is done by using standard database tools and other software. In this way, Cyber Vault checks Safeguarded Copy snapshots for corruption. If any corruption is found, it is considered a sign of a possible attack.

Data copy test and validation

Data copies are in an isolated environment to check whether they are corruption free and recovery procedures are in place. Also, Forensics and Diagnostics Services are available to check the validity of the data copies.

All of these features can be implemented by IBM Lab services or an IBM Business Partner by using the IBM Frameworks for IT cyber resiliency. However, more software might be required such as IBM QRadar, IBM Security Guardium Data Protection, and so on.

Figure 25 shows the time that is needed to come back to operational status without IBM Cyber Vault.

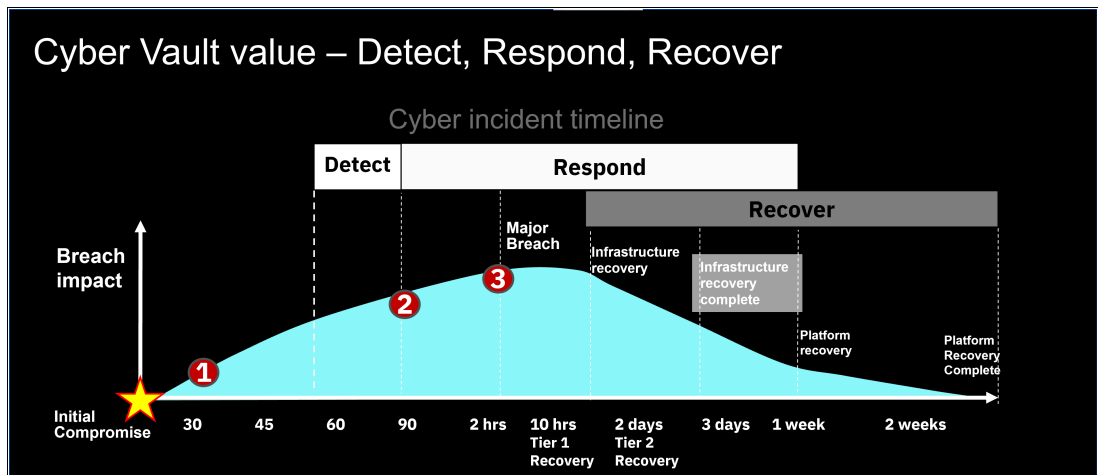


Figure 25 Cyberattack without IBM Cyber Vault

Figure 26 on page 52 shows the same attack *with* protection against Cyber Vault. It shows clearly how much time can be saved when you use the IBM Cyber Vault implementation.

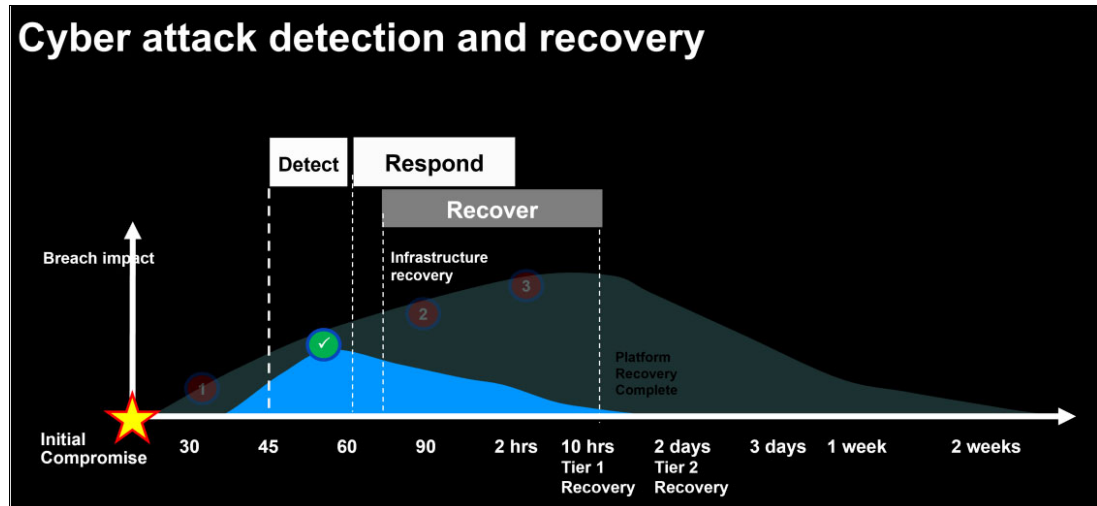


Figure 26 Cyberattack with IBM Cyber Vault

Encryption

Like its predecessors, IBM FlashSystem 9500 data encryption, based on the industry standard AES-XTS 256 encryption, as defined in the IEEE 1619-2007 standard and NIST Special Publication 800-38E as XTS-AES-256. The data encryption key is protected by a 256-bit AES key wrap of a key that is derived from the access key that is stored on the USB flash drive. The wrapped key is stored in the system in nonvolatile form.

Encryption on the IBM FlashSystem 9500 system requires the following feature codes:

- ▶ Encryption Enablement (#ACE9/ACEE)

This feature enables the encryption function. A single instance of this feature enables the function on the entire IBM FlashSystem 9500 system (IBM FlashSystem 9500 control enclosure, all attached IBM FlashSystem 9500 expansion enclosures), and on externally virtualized storage subsystems.

USB flash drives (#ACEA) IBM Security Guardium Key Lifecycle Manager, or Gemalto Safenet Keysecure are required for encryption key management.

- ▶ Encryption USB Flash Drives (Four Pack) Optional (#ACEA)

This feature provides four USB flash drives for storing the encryption master access key.

Unless IBM Security Guardium Key Lifecycle Manager or Gemalto Safenet Keysecure is used for encryption keys management, a total of three USB flash drives are required per IBM FlashSystem 9500 cluster when encryption is enabled in the cluster, regardless of the number of systems in the cluster.

If encryption is used in a cluster, this feature must be ordered on one IBM FlashSystem 9500 system, which results in a shipment of four USB flash drives.

Encryption can be applied to virtualized storage arrays, even if the virtualized array does not have encryption capabilities. In this scenario, the encryption is done by using IBM Storage Virtualize software. Encrypted volumes are transparent to applications, which ease implementation and operation. In addition, the IBM FlashSystem 9500 system has the following functions:

- ▶ Encryption Activation: Adding an encryption license to a system is not concurrent and must be done at array initialization time.

- ▶ Encryption Deactivation: Removing encryption is also non-concurrent and destroys any data on the array.
- ▶ Encryption Rekey: Changing the encryption key on a previously initialized system is concurrent and can be done while the array is in use.

These operations require that you purchase an Encryption Enablement Pack (#ACE9/ACEE).

Self-encrypting drives

The IBM FlashCore Modules and NVMe drives, including the SCMs, in the IBM FlashSystem 9500 control enclosure are self-encrypting drives (SEDs). With SEDs, you can encrypt the data on the drive within the hardware.

SEDs include the following features:

- ▶ Encryption of data is done in the electrical circuit of the drive; therefore, it is not affected by performance issues from software encryption.
- ▶ Data encryption keys never leave the confines of the SED, and are never loaded into CPU or memory.
- ▶ You can perform a fast cryptographic erasure of a SED by using a single CLI command to replace the data encryption key or revert the entire device to factory settings.
- ▶ Supports a security feature that is called *auto-lock*, which protects against someone inserting your drive into another system and accessing your data.
- ▶ Drives automatically lock themselves on power loss and require an access key at start time to unlock and allow I/O operations.
- ▶ If an SED drive is removed from a system with encryption and placed in another system, the drive data is not readable. The system posts an error message that indicates it is locked. The only way to use the drive is to format it. This formatting also performs a cryptographic erase by removing any encryption keys; therefore, all of the data on the drive is destroyed.

Combining system encryption with self-encrypting drives

For control enclosures that support NVMe architecture, NVMe-attached drives are self-encrypting and self-compressing. With SEDs that use NVMe architecture, data encryption is completed in the drive. Data encryption keys remain on the drive without being stored in system memory.

In addition, the system supports automatic locks of encrypted drives when the system or drive is powered down. When the drive or system restarts, a master key is required to unlock the drive and continue I/O operations.

Because the encryption of data is done in the electrical circuit of the drive, it is not affected by any potential performance issues from software encryption.

If SEDs are encrypting the data, why do you need to enable system encryption in IBM Storage Virtualize? The answer is that you can use SEDs without enabling encryption on the system, but SEDs are unlocked by default at start time, unless configured with extra protection.

With system-level encryption in IBM Storage Virtualize you can use USB flash drives or IBM Security Guardium Key Lifecycle Manager to manage access to encrypted objects on the system. This feature ensures that when a system is powered on, this extra encryption key is required to read the data on the drives.

Consider the following points:

- ▶ SEDs are always encrypting, and you cannot stop them from being encrypted.
- ▶ You can use SEDs without enabling encryption on the system, but SEDs are unlocked by default unless they are configured with extra protection.
- ▶ With system encryption in IBM Storage Virtualize, you can use USB flash drives or IBM Security Guardium Key Lifecycle Manager to manage access to encrypted objects on the system.
- ▶ Software in the operating system is required to manage an access key that can be used to lock and unlock the SEDs and bring them online for I/O.

Therefore, the best solution is to use the SEDs with the Encryption Enablement Pack and USB or IBM Security Guardium Key Lifecycle Manager type encryption, or a mixture of both. This configuration ensures the maximum level of encrypting for your data that is on the system.

Transparent cloud tiering and encryption

Transparent cloud tiering is a function that enables volume data to be copied and transferred to cloud storage. The system supports creating connections to cloud service providers to store copies of volume data in private or public cloud storage.

With Transparent cloud tiering, administrators can move older data to cloud storage to free up capacity on the system. Point-in-time snapshots of data can be created on the system and then copied and stored on cloud storage. After the system is authenticated, it can then access cloud storage to either copy data to the cloud storage or restore data that is copied to cloud storage back to the system. The system supports one cloud account to a single cloud service provider. Migration between providers is not supported.

An external cloud service provider manages the cloud storage, which reduces storage costs for the system. At the time of this writing, IBM supports the following cloud service providers:

- ▶ **IBM Cloud**

The system can connect to IBM Cloud, which is a cloud computing platform that combines platform as a service (PaaS) with infrastructure as a service (IaaS).

- ▶ **OpenStack Swift**

OpenStack Swift is a standard cloud computing architecture from which administrators can manage storage and networking resources in a single private cloud environment. Standard application programmable interfaces (APIs) can be used to build customizable solutions for a private cloud solution.

- ▶ **Amazon S3**

Amazon Simple Storage Service (Amazon S3) provides programmers and storage administrators with flexible and secure public cloud storage. Amazon S3 is also based on object storage standards and provides a web-based interface to manage, back up, and restore data over the web.

- ▶ **Microsoft Azure**

Microsoft Azure is an object storage solution that allows large volumes of unstructured data to be stored in the Azure public cloud. It provides an API and web-based interface to help you manage your cloud solution. Microsoft Azure is suitable for archiving, disaster recovery, and backup data storage.

The following considerations apply to Transparent cloud tiering and encryption:

- ▶ When a cloud account is created, it must continue to use the same encryption type throughout the life of the data in that cloud account. Even if the cloud account object is removed and remade on the system, the encryption type for that cloud account might not be changed while backup data for that system exists in the cloud provider.
- ▶ When performing rekeying operations on a system that has an encryption-enabled cloud account, perform the commit operation immediately after the prepare operation. Retain the previous system master key (on a USB or in the keyserver) because this key might be needed to retrieve your cloud backup data when performing a T4 recovery or an import.
- ▶ The `restore_uid` option should not be used when the backup is imported to a new cluster.
- ▶ Transparent cloud tiering uses Sig V2 when connecting to Amazon regions, and does not currently support regions that require Sig V4.

For more information, see [Transparent cloud tiering](#).

Secure drive erasure process

The IBM FlashSystem 9500 system that is running IBM Storage Virtualize V8.5.0 or higher provides methods to securely erase data from a drive or boot drive when a control enclosure is decommissioned or before a drive is removed from the system during a repair activity.

Secure data deletion effectively erases or overwrites all traces of data from a data storage device. The original data on that device becomes inaccessible and cannot be reconstructed. You can securely delete data that is on individual drives and on a boot drive of a control enclosure. The methods and commands that are used to securely delete data enable the system to be used in compliance with European Regulation EU2019/424.

The following types of drives can be used for this process:

- ▶ Expansion enclosure SAS SSDs
- ▶ NVMe drives including IBM FCMs and industry standard
- ▶ Control enclosure node canister SSD boot drives

The methods that the system uses to securely delete data from the drives varies according to the CLI commands that each type of drive can support. The completion time for the erase procedure also varies, depending on the amount of data and the method that is used to delete the data. In each case, when the operation completes, the result is that the data on the drive effectively becomes impossible to access.

Table 6 lists the types of erasure, the methods used, and the time taken.

Table 6 Comparison of methods to securely delete data from drives

Priority	Deletion type	Method	Completion time
1	Cryptographic erase	Changes the encryption key and makes the data inaccessible.	Instant
2	Block erase	Quickly raises and lowers the voltage level of the storage element. Physical blocks are altered with a vendor-specific value.	Fast
3	Data Overwrite	Replaces the existing data with random data.	Slow

The methods that are used to securely delete data vary according to manufacture, drive type, and drive firmware. For more information, see the documentation that is provided by the drive manufacturer.

If a drive supports more than one data deletion method, the system uses the highest-priority method.

For more information about the CLI commands that are used to run this secure erase function, see [Secure Data Deletion](#).

Reliability, availability, and serviceability

IBM FlashSystem 9500 delivers the following enterprise-class reliability features, much like its predecessors:

- ▶ IBM FlashSystem 9500 control enclosures support concurrent code load. Onsite and remote code upgrades are supported. However, an extra feature Remote Code Load (RCL) support #AHY2 must be purchased for remote code loads, which include code loads to the FCM drives. The RCL support #AHY2 feature is available with Expert Care support.
- ▶ The components of the IBM FlashSystem 9500 are one of the following types:
 - Hot-swappable: All system functions remain the same.
 - Concurrently replaceable: All functions remain, but elements might be offline, so performance or availability might be affected.

The IBM FlashSystem 9500 control enclosure features the following clustered node canisters that contain hot-swappable and concurrently replaceable components:

- Hot-swappable:
 - Secure boot drives
 - Batteries
 - PSUs (2 per canister)
 - NVMe or FCM drives
- Concurrently replaceable:
 - Fan modules
 - Memory DIMMs
 - Trusted Platform Module (TPM)
 - PCIe adapters
 - Node canister

If an IBM FlashCore Module failure occurs, critical customer applications can remain online while the defective module is replaced because a spare module is available.

- ▶ IBM enhanced technology in the SCM drives for higher storage density and improved performance can coexist with FCM4 drives in the IBM FlashSystem 9500 for even more demanding workloads.
- ▶ By using 3-site data copies, with synchronous and asynchronous data copies, greater flexibility exists in creating DR and HA solutions.

Options and feature codes

This section describes the options and feature codes of the IBM FlashSystem 9500 control enclosure.

Memory options

Table 7 lists the following memory options by feature code:

Table 7 IBM FlashSystem 9500 memory options per control enclosure / two node canisters.

Base Memory (GB)	#ACGN Plant or Field Upgrade (GB)	#ACGP Plant or Field Upgrade (GB)	Total Memory (GB)
1024	N/A	N/A	1024
1024	1024	N/A	2048
1024	1024	1024	3096

- ▶ (#ACGN): 1024 GB Cache upgrade

This feature indicates that the base machine is configured with 1024 GB (512 GB for each node canister). Selecting this feature determines the cache upgrade paths that the machine uses.

- ▶ (#ACGP): 1024 GB Cache upgrade

This feature provides another 1024 GB of cache (512 GB for each node canister) to increase the total system cache by 1024 GB to the total memory of 3096 GB.

Host I/O connectivity and expansion enclosure adapters

This section includes discussion of a single canister.

Remember: An IBM FlashSystem 9500 enclosure consists of *two* node canisters. The canisters have identical configurations.

Six PCIe slots are available for port expansions in the IBM FlashSystem 9500 node canister. Each canister has six PCIe adapter slots and both canisters must have the same configuration. The PCIe adapter feature codes offer a pair of adapters to ensure that they are supplied symmetrically in each canister.

The node canister can be configured with six I/O adapter features to provide up to twenty-four 32 Gbps FC ports or up to ten 10/25 Gbps Ethernet ports and six 100 Gbps Ethernet. The node canister includes two 1 Gbps Ethernet ports for system management. A feature code is also available to include the SAS Expansion card if the user wants to use optional expansion enclosures.

Figure 27 shows the connectivity options for the different adapters at time of this writing.

Adapter	Supported Protocols	Supported Optics	Ports
32/64 Gb Fibre Channel	Fibre Channel FC-NVMe	32Gb / 64Gb SW SFP 32Gb / 64Gb LW SFP	4
12G SAS	SAS Expansion	mini-HD	2
25Gb iWARP Ethernet	iSCSI Cluster/hyperswap over iWARP (RPQ only)	10Gb RJ45 SFP+ SFP28 and more*	2
25Gb RoCE Ethernet	iSCSI NVMe/RDMA	10Gb RJ45 SFP+ SFP28 and more*	2
100Gb RoCE Ethernet	NVMe/RDMA iSCSI**	QSFP28 MPO QSFP28 LC 100Gb DAC and more*	2

Figure 27 Connectivity options

For more information about the limits and rules for adapter placement to ensure correct best practices, see [IBM FlashSystems & SAN Volume Controller FAQ](#).

Table 8 lists the maximum host port count per building block configuration in a clustered configuration. At the time of writing, 2 control enclosures is the maximum number that can be clustered.

Table 8 Maximum host port count per control enclosure

Number of control enclosures	32 Gbps FC	64 Gbps FC	100 Gbps iSCSI (RoCE)	10 / 25 Gbps iSCSI (RoCE)	10 / 25 Gbps iSCSI (iWARP)
One	48	24	12	20	20
Two	96	48	24	40	40

Note: All connectivity ports are NVMe-oF compatible.

Table 9 lists the current features for host and connectivity for the IBM FlashSystem 9500 Control Enclosure 4666-AH8, 4666-UH8 and 4983-AH8 machine types.

Table 9 Supported expansion enclosure and interface components MTMs 4666-AH8, 4666-UH8 and 4983-AH8

Item	Feature code	Description	Ports
100 Gbps Ethernet (RoCE) adapters (Pair)	#AHB8	This feature provides two I/O adapters, each with two 100 Gbps Ethernet ports. It is used to add 100 Gbps Ethernet connectivity to the FlashSystem 9500 control enclosure and supports RDMA with RoCE 2.	Each adapter has two 100 Gbps Ethernet ports. Supplied without optics or cables.
SAS Expansion Enclosure Attach Card (Pair)	#AHBA	This feature provides two 4-port 12 Gbps SAS Expansion Enclosure attachment adapters. This feature is used to attach up to 3 AFF or 1 A9F Expansion Enclosure.	Each adapter has 2 active SAS ports.
32 Gbps FC 4-Port adapter pair	#AHBB	This feature provides two I/O adapters. It is used to add 32 Gbps FC connectivity.	Each adapter has four 32 Gbps FC ports and shortwave SFP transceivers.
10/25 Gbps Ethernet (RoCE) adapter pair	#AHBH	This feature provides two I/O adapters. It is used to add 10/25 Gbps Ethernet connectivity and supports RoCE V2.	Each adapter has two 25 Gbps Ethernet ports. Supplied without optics or cables.
10/25 Gbps Ethernet (iWARP) adapter pair	#AHBJ	This feature provides two I/O adapters. It is used to add 10/25 Gbps Ethernet connectivity. and supports RDMA with iWARP.	Each adapter has two 25 Gbps Ethernet ports. Supplied without optics or cables.
64 Gbps FC 4-Port adapter pair	#AHBK	This feature provides two I/O adapters. It is used to add 64 Gbps FC connectivity.	Each adapter has four 64 Gbps FC ports and shortwave SFP transceivers.
25 Gbps Ethernet InfoSphere® MDM Server Workbench SFP28 Transceivers (Four)	#ACHP	This feature provides four 25 Gbps shortwave SFP28 transceivers to populate a pair of 25/10 Gbps Ethernet host interface cards. These transceivers are used to set the card ports to 25 Gbps.	Requires feature AHBH, AHBJ, or AHB8.
10 Gbps Ethernet InfoSphere MDM Server Workbench SFP+ Transceivers (Four)	#ACHQ	This feature provides four 10 Gbps shortwave SFP+ transceivers to populate a pair of 25/10 Gbps Ethernet host interface cards. These transceivers are used to set the card ports to 10 Gbps.	Requires feature AHBH, AHBJ, or AHB8.
32 Gbps FC LW SFP Transceivers (Pair)	#ACHV	This feature provides two 32 Gbps longwave SFP transceivers for use with 32 Gbps FC I/O ports.	#AHBB is a prerequisite. The maximum allowed is four for each instance of #AHBB.
100 Gbps Ethernet QSFP28 SR4 Transceivers (Four)	#ACHX	This feature provides four 100 Gbps Ethernet QSFP28 transceivers to populate a pair of 100 Gbps Ethernet host interface cards.	Requires feature AHB8.

Cables

The following cable feature codes are available:

- ▶ (#ACSQ): 1 m OM3 Fiber Cable (LC)
- ▶ (#ACSU): 2 m OM3 Fiber Cable (LC)
- ▶ (#ACSR): 5 m OM3 Fiber Cable (LC)
- ▶ (#ACSS): 10 m OM3 Fiber Cable (LC)
- ▶ (#ACST): 25 m OM3 Fiber Cable (LC)

Cables that must be ordered with the expansion enclosure

The following cable feature codes are available with the expansion enclosures only:

- ▶ (#ACUA): 0.6 m 12 Gbps SAS Cable only for SFF Expansion (mSAS HD)
- ▶ (#ACUB): 1.5 m 12 Gbps SAS Cable only for SFF Expansion (mSAS HD)
- ▶ (#ACUC): 3 m 12 Gbps SAS Cable SFF/HD LFF Expansion (mSAS HD)
- ▶ (#ACUD): 6 m 12 Gbps SAS Cable SFF/HD LFF Expansion (mSAS HD)

IBM Storage Class Memory, FlashCore Modules, and industry-standard NVMe drive options

The IBM FlashSystem 9500 system supports Storage Class Memory (SCM) drives, IBM FlashCore Modules, industry-standard flash drives, and an intermix of all three.

IBM FlashCore Modules combine IBM MicroLatency technology, advanced flash management, and reliability into a 2.5-inch SFF NVMe with built-in hardware compression and encryption that do not affect performance.

For improved flexibility, IBM FlashSystem 9500 systems also support various industry-standard, self-encrypting NVMe flash drives.

The following NVMe flash drive feature codes are available:

- ▶ (#AHS9): 4.8 TB NVMe FlashCore Module 3
- ▶ (#AHTA): 9.6 TB NVMe FlashCore Module 3
- ▶ (#AHSB): 19.2 TB NVMe FlashCore Module 3
- ▶ (#AHSC): 38.4 TB NVMe FlashCore Module 3
- ▶ (#AHSE): 4.8 TB NVMe FlashCore Module 4
- ▶ (#AHSF): 9.6 TB NVMe FlashCore Module 4
- ▶ (#AHSG): 19.2 TB NVMe FlashCore Module 4
- ▶ (#AHSH): 38.4 TB NVMe FlashCore Module 4
- ▶ (#AHTH): 1.92 TB NVMe Flash Drive
- ▶ (#AHTI): 3.84 TB NVMe Flash Drive
- ▶ (#AHTJ): 7.68 TB NVMe Flash Drive
- ▶ (#AHTK): 15.36 TB NVMe Flash Drive
- ▶ (#AHTL): 30.72 TB NVMe Flash Drive
- ▶ (#AHTM): 1.6 TB NVMe SCM Drive

Consider the following points about limitations and drives:

- ▶ IBM FlashCore Modules:
 - DRAID 6 (recommended with a minimum of 6, and a maximum of 48)
 - DRAID 1 (minimum of 2, maximum of 16)
 - IBM FlashCore Modules in the same RAID array must be of the same capacity

Important: FCM2 and FCM3 + FCM4 can not work together in the same RAID Array. If you plan to upgrade / expand your FCM2 system you have to create a separate new RAID Array with either FCM3 or FCM4 and upgrade your System onto a supported software level. **Note:** FCM2 and FCM3 are already withdrawn.

FCM3 and FCM4 can work together in the same RAID Array, but the FCM4 drives will work like FCM3 drives without the enhanced Ransomware Detection features.

If you plan to have all of the new Ransomware Detection functionalities you need **ALL** drives on FCM4 and minimum IBM Storage Virtualize Code level 8.6.3.

- ▶ Industry-standard NVMe drives:
 - DRAID 6 (minimum of 6; maximum of 128)
 - DRAID 1 (minimum of 2; maximum of 16)
 - Industry-standard NVMe drives in the same RAID array must be of the same capacity
- ▶ SCM NVMe drives:
 - DRAID 6 (minimum of 6, maximum of 12)
 - DRAID 1 (minimum of 2, maximum of 12)
 - SCM NVMe drives in the same RAID array must be of the same capacity

IBM FlashSystem 9000 Expansion Enclosure options for models AFF and A9F

The following feature codes are available for SAS flash drives:

- ▶ Supported on Model AFF only (maximum of 24):
 - (#AH2A) 1.92 TB 12 Gbps SAS 2.5-inch flash drive
 - (#AH2B) 3.84 TB 12 Gbps SAS 2.5-inch flash drive
 - (#AH2C) 7.68 TB 12 Gbps SAS 2.5-inch flash drive
 - (#AH2D) 15.36 TB 12 Gbps SAS 2.5-inch flash drive
 - (#AH2E) 30.72 TB 2 Gbps SAS 2.5-inch flash drive
- ▶ Supported on Model A9F only (maximum of 92):
 - (#AH7J) 1.92 TB 12 Gbps SAS 3.5-inch flash drive
 - (#AH7K) 3.84 TB 12 Gbps SAS 3.5-inch flash drive
 - (#AH7L) 7.68 TB 12 Gbps SAS 3.5-inch flash drive
 - (#AH7M) 15.36 TB 12 Gbps SAS 3.5-inch flash drive
 - (#AH7N) 30.72 TB 12 Gbps SAS 3.5-inch flash drive (requires EOS 8.0)

Physical and electrical specifications

This section describes the specifications for the 9500 FlashSystem control and expansion enclosures.

IBM FlashSystem 9500 Control Enclosure for 4666 AH8 / 4666 UH8 / 4983 AH8

The control enclosure features the following specifications:

- ▶ Physical specifications:
 - Height: 17.43 cm (6.8 in.)

- Width: 44.6 cm (17.6 in.)
- Depth: 82.6 cm (32.6 in.)
- Approximate weight:
 - Empty: 25.18 kg (55.51 lb.)
 - Fully configured: 69.31 kg (152.82 lb.)
- ▶ Air temperature:
 - Operating: 5–35° C (41–95° F) 0–3048 m (0–10,000 ft.). Above 900 m, the allowable maximum air temperature decreases by 1° C per 300 m.
 - Nonoperating: 1–45° C (34–113° F).
- ▶ Relative humidity:
 - Operating: 8%–80% noncondensing
 - Nonoperating: 8%–80% noncondensing
- ▶ Electrical power:
 - Voltage range: 200–240 V AC
 - Frequency: 50–60 Hz
 - Maximum power: 3300 W (power usage varies with configuration and system usage)
 - Heat dissipation (BTU per hour): 11258 Btu/h
 - Acoustical noise emission: 8.2 bels (idling), 8.8 bels (operating)

IBM FlashSystem 9000 SFF Expansion Enclosure for 4666 and 4983 Model AFF

The Expansion Enclosure AFF features the following specifications:

- ▶ Physical specifications:
 - Height: 8.7 cm (3.4 in.)
 - Width: 48.3 cm (19.0 in.)
 - Depth: 55.6 cm (21.9 in.)
 - Approximate weight:
 - Empty: 16.7 kg (36.8 lb)
 - Fully configured: 25.0 kg (55.1 lb)
- ▶ Air temperature:
 - Operating: 5–35° C (41–95° F) 0–3048 m (0–10,000 ft.). Above 900 m, the allowable maximum air temperature decreases by 1° C per 300 m.
 - Nonoperating: 1–50° C (34–122° F).
- ▶ Relative humidity:
 - Operating: 8%–80% noncondensing
 - Nonoperating: 8%–80% noncondensing
- ▶ Electrical power:
 - Voltage range: 100–240 V AC
 - Frequency: 50–60 Hz
 - Power: 800 W
 - Heat dissipation (BTU per hour): 1,037
 - Acoustical noise emission: 6.2 bels (idling), 6.2 bels (operating)

IBM FlashSystem 9000 LFF High-Density Expansion Enclosure for 4666 and 4983 Model A9F

The High-Density Expansion Enclosure A9F features the following specifications:

- ▶ Physical specifications:
 - Height: 22.2 cm (8.75 in.)
 - Width: 48.3 cm (19.0 in.)
 - Depth: 96.8 cm (38.1 in.)
 - Approximate weight:
 - Empty: 67.0 kg (147.7 lb)
 - Fully configured: 135.0 kg (297.0 lb)
- ▶ Air temperature:
 - Operating: 5–35° C (41–95° F) 0–3048 m (0–10,000 ft). Above 900 m, the allowable maximum air temperature decreases by 1° C per 300 m.
 - Nonoperating: 1–50° C (34–122 degrees F).
- ▶ Relative humidity:
 - Operating: 8%–80% noncondensing
 - Nonoperating: 8%–80% noncondensing
- ▶ Electrical power:
 - Voltage range: 180–264 V AC
 - Frequency: 47–63 Hz
 - Power: 2400 W
 - Heat dissipation (BTU per hour): 8,189
 - Acoustical noise emission: 8.5 bels (idling), 8.5 bels (operating)

Note: All noise emission levels are the declared upper limit sound power level in bels for a random sample of machines. All measurements are made in accordance with ISO 7779 and reported in conformance with ISO 9296.

Software and licensing

All IBM FlashSystem 9500 and IBM FlashSystem 9500R systems include preinstalled IBM Storage Virtualize Software V8.7 or later. The following product numbers are for the base licenses for each model of the IBM FlashSystem 9500 system:

- ▶ IBM Storage Virtualize Software for IBM FlashSystem 9500 Controller (5639-011) requires at least one IBM FlashSystem 9500 Control Enclosure 4666-AH8 or 4666-UH8 for installation.
- ▶ IBM FlashSystem 9500 Control Enclosure 4983-AH8 is LIC, and does not require a separate IBM Storage Virtualize Software license

All features are inclusive except for encryption, which is a feature code that is enabled for those countries that allow it, and external virtualization. Any connected storage that is not an IBM FlashSystem 9500 control enclosure requires the External Virtualization license per storage capacity unit (SCU) that is based on the tier of storage that is available on the external storage system. In addition, if you use FlashCopy and Remote Mirroring on an external storage system, you must purchase a per-tebabyte license to use these functions.

Consider the following points:

- ▶ FlashSystem 9500 control enclosures support external storage virtualization. Use of the external virtualization capability is entitled through the acquisition of IBM Storage Virtualize Software for SAN Volume Controller (PID 5641-VC8 in AAS and PID 5725-M19 in Passport Advantage®). These licenses are the same licenses as for IBM SAN Volume Controller. Therefore, IBM SAN Volume Controller licenses can be used for the IBM FlashSystem 9500 system for these features.
- ▶ With the IBM FlashSystem 9500 system, a license exists for the hardware-assisted encryption if it is purchased (Encryption Enablement Pack [#ACE9 and #ACEE]). This feature code is needed if you want to use USB-Key encryption, IBM Security Guardium Key Lifecycle Manager-based encryption, or both on the control enclosure.
- ▶ As a best practice, use the Encryption USB Drive Pack feature (#ACEA). This feature provides four USB flash drives for storing the encryption master access key. A total of three USB flash drives are required per FlashSystem cluster when encryption is enabled in the cluster, regardless of the number of systems in the cluster.

If encryption is to be used in a cluster, this feature must be ordered on one IBM FlashSystem, which results in a shipment of four USB flash drives.
- ▶ A storage system that is used only as a quorum device does not need a software license.

To set these licenses, use the Licensed Function page in the System Setup wizard.

If you are adding these licenses to a system, select **Settings** → **System** → **Licensed Function** in the management GUI. You can also use the `chlicense` CLI command to update current license settings on the system.

For more information about IBM Storage Virtualize licensing, see [Licensed functions](#).

IBM Storage Virtualize external licensing

The IBM FlashSystem 9500 system supports differential, capacity, and key-based licensing. For example, with external virtualization, differential licensing charges different rates for different types of storage. This difference provides cost-effective management of capacity across multiple tiers of storage. Licensing for these functions is based on the number of SCUs that is purchased.

With other functions, such as remote mirroring and FlashCopy, the license grants a specific number of terabytes of capacity for that function.

Key-based licenses, such as encryption, require an authorization code and key to be downloaded to the system before the function can be used.

Note: SCUs are needed for only virtualized storage that does not have the 5639-011 base license.

The SCU is defined in terms of the category of the storage capacity, as listed in Table 10.

Note: Any storage use case that is not listed in Table 10 is classified as Category 1.

Table 10 SCU category definitions

License	Drive class	SCU Ratio
SCM	Storage Class Memory (SCM) devices	SCU equates to 1.00 TB of usable of Category 1 storage
Flash	All flash devices, other than SCM drives	SCU equates to 1.18 TB of usable of Category 1 storage
Enterprise	10 K or 15 K RPM drives	SCU equates to 2 TB of usable of Category 2 storage
Nearline	Nearline SATA drives	SCU equates to 4.00 TB of usable of Category 3 storage

Table 11 shows an example of calculating SCUs (storage capacity unit). The example is a customer who virtualizes external disk arrays with 5 TB SCM devices, 30 TB SSD flash drives, 400 TB Enterprise drives, and 800 TB Nearline capacity.

Table 11 Example of calculating SCUs

Category	Type	Capacity (TB)	TB per SCU	# of SCUs
Category 1	SCM	5	1.00	5
Category 1	SSD	30	1.18	26
Category 2	Enterprise	400	2.00	200
Category 3	Nearline	800	4.00	200
Total		1235		431

A total of 431 SCUs are required for the example that is shown in Table 11. When you calculate the number of SCUs per category, fractions must be rounded up to the next higher integer number.

For more information about IBM Storage Virtualize Differential Licensing, see [Licensed functions](#).

IBM Global Financing

IBM Global Financing offers competitive financing to credit-qualified customers and IBM Business Partners to assist them in acquiring IT solutions. The offerings include financing for IT acquisition, including hardware, software, and services, from IBM and other manufacturers or vendors. Also offered is commercial financing (revolving lines of credit, term loans, acquisition facilities, and inventory financing credit lines) for IBM Business Partners.

Offerings for small, medium, and large enterprise customer segments, rates, terms, and availability can vary by country. For more information, contact your local IBM Global Financing organization or see [IBM Global Financing website](#).

Ordering information

For more information about ordering IBM FlashSystem 9500 expansions, see “Scaling up and scaling out” on page 47.

For more information about ordering hardware features, see “Options and feature codes” on page 57.

For more information about ordering software licenses, see “Software and licensing” on page 63.

Related information

For more information, see the following resources:

- ▶ IBM Documentation for IBM FlashSystem:
<https://www.ibm.com/docs/en/flashsystem-7x00/8.7.x>
- ▶ IBM FlashSystem 9500 product page:
<https://www.ibm.com/products/flashsystem-9500>
- ▶ IBM Offering Information page (announcement letters and sales manuals):
http://www.ibm.com/common/ssi/index.wss?request_locale=en
- ▶ IBM FlashSystems & SAN Volume Controller FAQ
<https://www.ibm.com/downloads/cas/2DWAMWRB>
- ▶ IBM FlashSystem Family FAQ
Overview of the IBM FlashSystem family with guidance on how to select the product that is right for you:
<https://www.ibm.com/downloads/cas/9OGKVV2R>
- ▶ IBM FlashWatch FAQ
Guidance about the IBM FlashWatch programs:
- ▶ IBM Storage Virtualize FAQ
Details about the IBM Storage Virtualize products, covering the IBM FlashSystem family and SAN Volume Controller:
<https://www.ibm.com/downloads/cas/2DWAMWRB>

Authors

Hartmut Lonzer brings 46 years of technical and sales expertise at IBM to his role as a Storage Advisory Partner Technical Specialist for DACH. Previously, he served as an OEM Alliance Manager for Lenovo at IBM Germany. Based at the company's headquarters in Ehningen, his focus lies on the IBM FlashSystem Family and IBM SAN Volume Controller products, with experience dating back to their introduction.

Jon Herd is an IBM Executive Technical Advocate working for the European Storage Competency Center (ESCC), Germany. He covers the United Kingdom, Ireland, and Europe, advising customers on a portfolio of IBM storage products, including IBM FlashSystem products. He also leads special projects for senior and executive management and is the SME lead for new product introduction in the ESCC. Jon has been with IBM for more than 47 years, and has held various technical roles, including Europe, Middle East, and Africa (EMEA) level support on mainframe servers and technical education development. He has written many IBM Redbooks publications about IBM FlashSystem products and is an IBM Redbooks Platinum level author. He holds IBM certifications in Product Services at a Thought Leader L3 level, and Technical Specialist at an experienced L1 level. He is also a certified Chartered Member of the British Computer Society (MBCS - CITP), a Certified Member of the Institution of Engineering and Technology (MIET), and a Certified Technical Specialist of the Open Group (TOG).

Vasfi Gucer works as the Storage Team Leader on the IBM Redbooks Team. He has more than 30 years of experience in the areas of systems management, networking hardware, and software. He writes extensively and teaches IBM classes worldwide about IBM products. His focus has been primarily on storage, cloud computing, and cloud storage technologies for the last 8 years. Vasfi is also an IBM Certified Senior IT Specialist, Project Management Professional (PMP), IT Infrastructure Library (ITIL) V2 Manager, and ITIL V3 Expert.

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

Evelyn Perez
James Whitaker
Paul Edmonds
IBM UK

Luis F Lopez
Edgar Daniel Gutierrez Haro
IBM Mexico

The team also would like to thank the authors of the previous edition of this publication: **Corne Lottering, Douwe van Terwisga, Kendall Williams, Konrad Trojok, and Shu Mookerjee**

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

Stay connected to IBM Redbooks

- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>

Notices

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

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

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

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

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

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

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

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo)  ®	IBM®	Passport Advantage®
AIX®	IBM Cloud®	PowerPC®
DS8000®	IBM FlashCore®	PowerVM®
Easy Tier®	IBM FlashSystem®	QRadar®
FlashCopy®	IBM Security®	Redbooks®
Guardium®	InfoSphere®	

The following terms are trademarks of other companies:

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

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

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

VMware, VMware vSphere, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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



REDP-5742-00

ISBN

Printed in U.S.A.

Get connected

