

IBM FlashSystem 9500

Overview

The IBM FlashSystem 9500 is an end-to-end NVMe storage system, meaning it supports NVMe over Fabric with Fibre Channel and Ethernet as well as attaching storage devices in the form of SSDs and FlashCore Modules using PCIe with NVMe protocol. The FlashSystem is targeted with the performance and capacity to cover requirements for high-end enterprises. With Spectrum Virtualize software in the FlashSystem 9500, it is capable of virtualizing other attached storage systems providing common administration, advanced features such as replication and snapshots, and the ability to seamlessly migrate data. The FlashSystem 9500 is a scale-out storage system with the basic unit being a pair of controllers that can be scaled to four pairs or eight controllers in a cluster.

Storage devices used in the FlashSystem 9500 include NVMe protocol SSDs and FlashCore Modules along with SCM devices. The FlashCore modules were re-engineered from the original FlashSystem 900 to attach to the controllers using PCIe and support the NVMe protocol. With the introduction of the FlashSystem 9500, PCIe Gen4 is used. FlashCore modules are in a 2.5" form factor. Up to 12 SCM devices can be supported as a storage tier. Both 24 and 96 2.5" device enclosures are supported with up to 184 expansion solid state devices per controller pair.

The functionality of the FlashSystem is the same as the IBM SAN Volume Controller (SVC). The management is a web-based GUI that is highly advanced, targeted at an IT generalist. Different from most other storage systems, the advanced features are pre-installed on the FlashSystem 9500 with many being included in the base price of the system. Additional cost features are enabled with license keys.

For high availability, FlashSystem 9500 runs in a paired configuration with custom clustering software. Clustering for a scale out configuration is supported for up to four pairs of controllers. (Note – IBM uses the term 'node' for a controller

Highlights

- End-to-End NVMe high performance storage system
- Fibre Channel and Ethernet attach – NVMe or block access protocols
- Virtualized external storage with 32PB max capacity
- Advanced features – snapshot and remote replication, inline data reduction, active-active stretched clusters
- Thin provisioning
- Easy-Tier tiering between device levels and storage systems
- VAAI, SRM & VASA support
- Encryption for internal disks and external attached storage
- Transparent Cloud Tiering
- Storage Insights analytics & monitoring
- Safeguarded Copy for data resiliency

rather than a pair of controllers, which is different from most other vendors).

IBM Storage Insights is included with the FlashSystem purchase, providing storage analytics and monitoring. Storage Insights improves resiliency, manages problem resolution, and has storage resource management reporting.

For the FlashSystem 9500, IBM has an evergreen program where the controllers will be replaced with purchase of a three-year hardware and software maintenance agreement.

Usage

The IBM FlashSystem 9500 is targeted as a high-end enterprise general purpose, block storage system with the ability to virtualize existing storage systems as a consolidation solution. End-to-End NVMe support on the 9500 gives the system performance capabilities to address demanding applications and for consolidation of multiple workloads to a single platform.

- Characteristics
 - Performance – IBM has released performance data stating 2.5M Read Miss IOPS from a single system with a pair of controllers and 10M IOPs from a cluster of four pairs. Bandwidth performance quoted as 106GB/s.
 - Availability – The FlashSystem 9500 is dual controller nodes (a pair) as a system with the option for clustered nodes. It is designed as a high availability system with the ability for failover between controllers.
 - Replication for BC/DR – Remote mirroring, both synchronous (FC) and asynchronous (FC and IP), is supported along with consistency groups. File replication is asynchronous over Ethernet.
 - Active-active stretched clusters are supported between two sites with HyperSwap
 - Encryption for internal and external storage
 - Data reduction – deduplication, compression, and compaction
- Applications
 - The FlashSystem 9500 is targeted at any block storage application. As a general-purpose storage system, the 9500 fits a very wide range of application applicability.
- System environments
 - The FlashSystem 9500 has been integrated with VMware Site Recovery Manager and has support for VAAI, VASA 2.0 and VVOL.
 - With the ability to virtualize other storage systems and provide both passthru LUN access and non-disruptive migration of LUNs, FlashSystem 9500 is well suited as a consolidation solution.
- Deployment and Administration
 - The FlashSystem 9500 has an intuitive GUI that allows it to be installed and managed by a customer with moderate storage skills.

- The management is integrated into IBM management products but can also be completely managed from the GUI.
- A mobile application for management is available as a download from the Apple App Store.

Key Capabilities

Hardware Architecture

FlashSystem 9500 is an Intel-based dual controller system. Each controller pair is packaged with up to 48 NVMe devices in 4U. FlashSystem 9500 supports both standard NVMe SSDs as well as custom devices called FlashCore Modules and storage class memory. The system supports up to a max of 48 front end host ports with support for FC, NVMe-oF/FC, iSCSI, and NVMe-oF/RDMA. Additionally, FlashSystem 9500 supports up to 4-way clustering for increased scalability.



Figure 1: IBM FlashSystem 9500 Enclosure for SSD & FlashCore Modules




Figure 2: FlashSystem 9500 Physical View

IBM's FlashCore Modules (FCM) are custom NVMe devices that provide inline-hardware compression and data protection. FCMs include an onboard processor, QLC NAND to provide storage density and SLC NAND to provide performance.



Figure 3: FlashCore Modules

Storage Virtualization

The virtualization of external storage systems from a variety of vendors has two forms that can be utilized. One is to explicitly map the LUNs (volumes) from the external storage for host access through the FlashSystem and the other is to migrate the data through the FlashSystem to internal or external storage to take advantage of the block-level virtualization with the advanced features for thin provisioning, wide striping, and the EasyTier automated data tiering. Since all the accesses are through the FlashSystem controllers, this type of virtualization has been called an “in-band” type of virtualization approach.

Storage virtualization utilizes mapping tables to transform underlying storage into storage presented to hosts. The underlying physical storage volumes known as MDisks (Managed Disks) within the FlashSystem 9500 are grouped and then mapped into VDisks (Virtual Disks). These VDisks are presented as LUN’s to hosts from an HA pair of nodes within an IO group. The graphic below illustrates the mapping of resources.

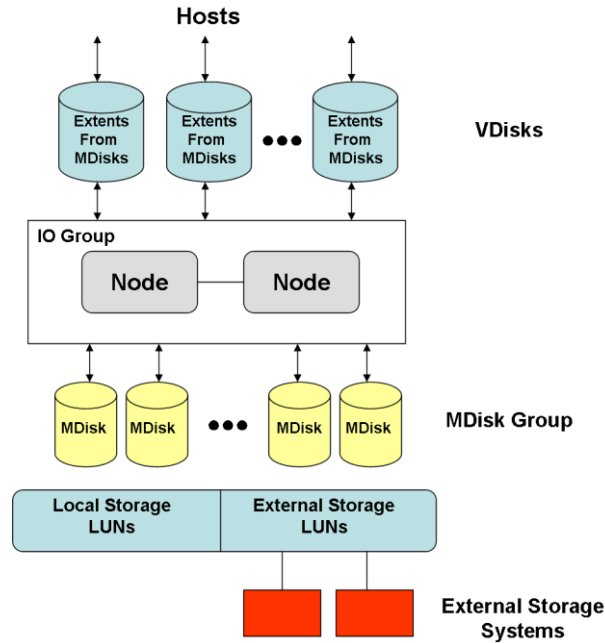


Figure 1: IBM FlashSystem 9500 Virtualization Implementation

Data Reduction

Included in the system is data reduction via Data Reduction Pools. Data is deduplicated, compressed, and compacted as it is received in the 9500 controllers and stored in reduced form. Data may also be compressed in the FlashCore Modules. Using Spectrum Virtualize software, data in the virtualized external storage systems may also be reduced.

HyperSwap Active-Active Stretched Clusters

A stretched cluster in Spectrum Virtualize systems is where clusters may be in two different data centers. A volume can be moved between clusters using mirroring or migration tools such as VMware VMotion. The stretched clusters take advantage of the clustering software and the network interconnects between nodes. I/O Groups are isolated by cluster for a stretched cluster. HyperSwap uses Metro Mirror, Global Mirror with Change Volumes, and non-disruptive Volume Move capabilities. Two sites are supported with automatic failover between volumes in the two controller pairs, one at each site. The Quorum disks which are the intelligence for making a decision on which of the surviving systems should be primary may be connected to another system over Fibre Channel or as a server in a third location connected over IP.

Data Protection

IBM FlashSystem includes several data protection functionalities including FlashCopy, Metro Mirror, and Global Mirror.

The FlashCopy feature provides an instant, point in time copy of a source VDisk to a target VDisk. The term “FlashCopy” is IBM’s term for point in time copy (snapshot) software on various platforms. The implementation of FlashCopy allows point-in-time copies of VDIs (LUNs) to be created. The FlashCopy is host independent and can be created across internal storage and different external storage systems. Beginning in October 2020, the snapshot implementation was changed from copy-on-write to redirect-on-write. FlashCopy is included with the base price of the system.

FlashSystem 9500 provides two products for remote replication. Metro Mirror provides synchronous replication with Global Mirror providing asynchronous replication.

SafeGuarded Copy

IBM FlashSystem supports Safeguarded Copy functionality similar to that found in IBM DS8900 systems. Safeguarded Copy provides immutable, logically air gapped copies of data to protect against logical data corruption and ransomware attacks. The Safeguarded Copies are point in time copies that are logically separated from other volumes and cannot be mapped to a host. They may also not be altered or deleted before a set retention time. Along with immutability and the logical air gap of the copies, Safeguarded copies provide further protection from bad actors by separating duties between administrators of data and administrators of safeguarded copies. FlashSystem Safeguarded Copies can be automatically created, deleted, or restored using the IBM Copy Services Manager.

Transparent Cloud Tiering

Transparent Cloud Tiering allows volumes to be tiered to clouds including Amazon S3, OpenStack Swift, and IBM Cloud. IBM Cloud Object Storage is also a tiering target.

Significant Announcements

- 2020 – IBM refresh of FlashSystem portfolio and consolidation of Storwize arrays.
- 2021 – FlashSystem adds SafeGuarded Copy
- 2022
 - FlashSystem 9500 replaces 9200
 - IceLake Processors
 - Updated FCM devices

The FlashSystem 9500 system was introduced to replace the previous FlashSystem 9200. Included in this update was a processor upgrade to Intel IceLake processors as well as the introduction of FCM v3 devices utilizing QLC media. Additionally, Safeguarded Copy functionality has been added to the

FlashSystem portfolio to provide cyber resiliency and ransomware protection. Evaluator Group views all of these as positive enhancements to the FlashSystem offering.

Consistent hardware and processor upgrades indicate development and investment in the product by IBM. Meanwhile the continued development and functionality of FCM devices is a differentiating factor for FlashSystem. The cyber resiliency functionality added with SafeGuarded copy is additionally important for addressing the increasing threat of ransomware and cyber-attacks.

Evaluator Group EvaluScale: IBM FlashSystem 9500 - SAN Storage

Evaluator Group product review methodology “EvaluScale” assesses each product within a specific technology area. The definitions of the criteria and explanations of how products are reviewed can be found in the [Evaluation Guides](#).

	Criteria	Description	Requirement	EG View of FlashSystem 9500	Explanation for FlashSystem 9500
1	Capacity	Current capacity of system to meet demand	Must have enough capacity to meet current demand and have ability to scale-up – adding more capacity up to a practical limit.	Exceeds requirements	Scale-out system can scale very large and support for storage virtualization provides massive scale.
2	Price – including data reduction	Cost of system. This includes data reduction effect – compression/ deduplication	Must be competitive with other leading solutions in this space meaning prices have no more than 20% variance from an average of the other solutions. This includes the effect of data reduction according to the Evaluator Group Data Reduction Estimator tool.	Meets requirements	Inline data reduction reduces cost to store data. More than a single model for different price points.
3	Performance	Latency IOPs Bandwidth	The performance requirement can vary based on high-end, mid-tier or entry-level usage. Requirements are that the system be within 10% of the top numbers for products in this class using SPC or IOMARK benchmarks.	Exceeds requirements	End-to-End NVMe provides tremendous boost in performance. Low latency with NVMe. Significant performance increases – reported data.
4	Scaling – performance and capacity	Ability to increase to meet future demands	Scale-out means scaling both performance and capacity to meet demands up to a practical limit – more capacity without sacrificing performance by crippling the access density.	Exceeds requirements	Has scale-out capability (adding controllers and storage) up to 8 controllers. Added with storage virtualization, this provides advantages.
5	Security – encryption	Data at rest encryption and key management	Data encryption is a perceived requirement across different usage categories for many industries and applications. The	Exceeds requirements	Have encryption with SKLM managed keys as well as internal.

			high-end enterprise requires an external key manager. A key manager also is a positive factor for the mid—tier. Multi-factor authentication support.		Safeguarded copy for data resiliency.
6	Data protection	Snapshots – R/W, large number Synchronous replication Asynchronous replication Tiering to clouds	Read/write snapshots are a requirement with a number that roughly equals the number of LUNs supported. The high-end enterprise requirement for remote replication is for both synchronous and asynchronous technology while mid-tier and entry usage require asynchronous. Tiering to clouds may be a benefit to all segments but would only be a requirement in the high-end enterprise.	Exceeds requirements	Full complement of advanced features and tiering to cloud for data protection or migration.
7	Business continuity	Active-active stretched clusters	The high-end enterprise generally requires active-active stretched clusters. Some mid-tier environments will need that as well but would not be a current requirement but could be a differentiator.	Meets requirements	Business continuity with HyperSwap feature. Limited number of volumes supported.
8	Economic considerations	Warranty Evergreen updating Environmentals – power & space Simplicity for admin	The overall environmental footprint being roughly on par with other leading systems in this area is the requirement measure. An extended warranty period for devices and an evergreen program for the controllers in the case of an all-flash system is now a requirement given the competitive nature.	Exceeds requirements	Advanced GUI - simplicity for use. FlashSystem has low power, space requirements. Evergreen program for controllers.
9	Storage technology	Use of solid state technology as storage devices and for caching	The requirement is for the system to support solid-state technology for at-rest storage and to have an extension cache using solid-state storage technology.	Exceeds requirements	Uses latest technology NVMe SSDs and custom FlashCore modules connected over PCIe. SCM devices are also available.
10	Application/system integration	VMware VAAI, VASA, SRM, etc. and Microsoft	Given the pervasiveness of hypervisor use in server virtualization, the integration for	Exceeds requirements.	Spectrum Virtualize software has full integration with major application / system



ODX, SCOM
integration

storage primitives and
management is a requirement.

software. Supports CSI
for containers.

Evaluator Group Opinion and Outlook for IBM FlashSystem 9500

Evaluator Group believes the FlashSystem 9500 system to be a leading high-end block storage system with advanced features, scale-out capability, storage virtualization, and data reduction.

As a high-end SAN system Evaluator Group believes IBM FlashSystem either meets or exceeds requirements in most areas identified. FlashSystem is a high-performance, scale-out block storage system that builds upon a mature software base in Spectrum Virtualize. FlashSystem's scale-out capability combined with custom FCM devices provides a system capable of meeting the capacity and performance needs of most environments.

Performance is further boosted by end-to-end NVMe support, an area that Evaluator Group believes crucial to modern high end SAN systems. While FlashSystem offers NVMe-oF options in FC and RoCE, adding support for NVMe-oF/TCP would be a key area of improvement. Evaluator Group has identified NVMe-oF/TCP as being a leading protocol for NVMe fabrics and believes this addition would help FlashSystem remain competitive in this area.

FlashSystem 9500 additionally has a robust set of advanced features to meet most enterprise's needs including data protection, business continuity, encryption, and cloud tiering. The addition of IBM's SafeGuarded Copy functionality adds crucial cyber resiliency functionality. FlashSystem's virtualization feature is a key differentiator for the product that is not available in many of the competing solutions. Storage virtualization allows FlashSystem to virtualize external systems and expand its capacity up to 32 PB.

IBM's custom FlashCore Modules are an additional key differentiator for the product. The combination of QLC and SLC NAND provides both high capacity and performance, while the devices also provide built in data reduction and protection functionality. Evaluator Group also acknowledges FCM devices to have had reports of very high reliability in the field.

Evaluator Group projects the future outlook of IBM FlashSystem to continue as a strong competitor in the SAN market. The system is built off mature software and supported by a large vendor with a long history in the storage market in IBM. Evaluator Group believes that the consolidation of Storwize and FlashSystem products into a single FlashSystem line was a positive move that provided clarity around IBM's focus on FlashSystem moving forward. Additionally, the systems have continued to see both hardware and software updates on a regular basis, indicating that development of FlashSystem is ongoing. While IBM has restricted some of the advanced features from lower-level models in the FlashSystem line, FlashSystem 9500 does not have this issue as IBM's high-end system and should continue to receive the full set of any new advanced features moving forward.

Information that is more detailed is available at <http://evaluatorgroup.com>

Copyright 2023 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.