

Introducing the New IBM FlashSystem Family and SAN Volume Controller

Silverton Consulting, Inc. StorInt™ Briefing



RAYONSTORAGE.COM | GREYBEARDSONSTORAGE.COM
SILVERTONCONSULTING.COM | TWITTER.COM/RAYLUCCHESI
© 2020 SILVERTON CONSULTING, INC. ALL RIGHTS RESERVED

Introduction

IBM® has had a long and productive history with FlashSystem and SAN Volume Controller (SVC) storage solutions. Customers have benefited greatly from the high-end performance and advanced functionality that come with IBM storage hardware and Spectrum Virtualize storage software.

Recently, IBM refreshed the FlashSystem family to include entry-level and mid-range solutions suitable for any enterprise environment. IBM has also enhanced FlashSystem mid-range and high-end enterprise products with better performance and new storage media. In addition, IBM has introduced new higher performing and lower cost SAN Volume Controller engine hardware.

With the new FlashSystem family and SVC engines IBM has made storage simple to acquire, manage, use the multi-cloud/hybrid cloud, support high data availability, and provide cyber security resiliency.

Below, we will highlight the features of these enterprise storage solutions to demonstrate how IBM has simplified its portfolio as well as why these systems are the lowest cost and highest performing storage options on the market today.

IBM FlashSystem Storage

IBM FlashSystem storage has historically provided leading-edge, all-flash array (AFA) storage solutions that supply high IOPS, low latency and high-bandwidth performance with high reliability, high availability and advanced storage functionality. With the recent enhancements to the family, that tradition continues.

FlashSystem entry and mid-range enterprise storage solutions

IBM has added hybrid (flash-disk media) entry-level and mid-range enterprise storage solutions to their FlashSystem lineup. The FlashSystem entry enterprise storage systems remain unchanged and were previously known as IBM Storwize V5010E, V5030E and V5100.

FlashSystem entry and mid-range enterprise storage solutions include the following:

- **IBM FlashSystem 5010 AFA and 5010H hybrid entry enterprise storage**
– a 2U, dual-controller system that supports 16Gbps FC, 12Gbps SAS, 10Gbps iSCSI/FCoE and 1Gbps iSCSI.
With 16GB of cache



per system, it supports up to 392 drives that can include the following:

- Small Form Factor (SFF) SAS SSDs (400GB to 15.4TB);
 - SFF SAS disks – 15K RPM (300GB and 600GB), 10K RPM (900GB to 1.8TB) and 7.2K RPM (2TB); and
 - Large Form Factor (LFF) SAS disks – 15K RPM (300GB to 900GB), 10K RPM (900GB to 1.8TB) and 7.2K RPM (2TB to 10TB).
- **IBM FlashSystem 5030 AFA and 5030H hybrid entry enterprise storage** – a 2U, dual-controller system that supports 16Gbps FC, 12Gbps SAS, 10Gbps iSCSI/iSER (iSCSI Extensions for RDMA using iWARP or RoCE) and 1Gbps iSCSI. With 32 or 64GB of cache per system, it supports up to 504 drives per control enclosure or 1008 per two-way clustered system, including the same drives as 5010/5010H.
 - **IBM FlashSystem 5100 AFA and 5100H hybrid entry enterprise storage** – a 2U, dual-controller system that supports 16 or 32Gbps FC or NVMe over Fabric [NVMeoF]/FC, 10Gbps Ethernet (iSCSI) and 25Gbps [iSCSI/iSER, (using iWARP or RoCE)] with 64GB to 576GB of cache per system. It supports up to 24 NVMe devices in one control enclosure and can attach up to 736 SAS drives using expansion enclosures or 48 NVMe devices and can attach up to 1008 SAS drives using expansion enclosures per two-way clustered system, which can include the following:
 - SFF NVMe storage-class memory (SCM) drives (375GB to 1.6TB);
 - SFF NVMe FlashCore Modules (FCMs, 4.8TB to 38.4TB usable capacity, hardware compressed);
 - SFF NVMe SSDs (800GB to 15.4TB);
 - SFF SAS SSDs (800GB to 30.8TB);
 - SFF SAS disks – 15K RPM (300GB and 600GB), 10K RPM (900GB to 1.8TB) and 7.2K RPM (2TB); and
 - LFF SAS disks – 7.2K RPM (4TB to 14TB).
 - **IBM FlashSystem 7200 AFA and 7200H hybrid mid-range enterprise storage** – a 2U, dual-controller system that supports 16 or 32Gbps FC or NVMeoF/FC, 10Gbps Ethernet (iSCSI) and 25Gbps [iSCSI/iSER, (iWARP or RoCE)] with 256GB to 1.5TB of cache per system. It supports up to 24 NVMe devices in one control enclosure and can attach up to 736 SAS drives with expansion enclosures. or 96 NVMe devices and can attach 2944 SAS drives with expansion enclosures per four-way clustered system, including the following:
 - SFF NVMe SCM drives (375GB to 1.6TB);
 - SFF NVMe FCMs (4.8TB to 38.4TB usable capacity; hardware compressed);
 - SFF NVMe SSDs (800GB to 30.7TB);



- SFF SAS SSDs (800GB to 30.8TB);
- SFF SAS disks –10K RPM (900GB to 2.4TB) and 7.2K RPM (2TB); and
- LFF SAS disks – 7.2K RPM (4TB to 14TB).

FlashSystem high-end enterprise storage solutions

IBM offers FlashSystem 9200 and 9200R AFA storage for high-end enterprise storage:

- **IBM FlashSystem 9200 AFA high-end enterprise storage** – a 2U, dual-controller system that supports 16 or 32Gbps FC or NVMeoF/FC, 10Gbps Ethernet (iSCSI) and 25Gbps [iSCSI/iSER, (iWARP or RoCE)] with 256GB to 1.5TB of cache per system. It supports up to 24 NVMe devices in one control enclosure and can attach up to 736 SAS drives with expansion enclosures or 96 NVMe devices and can attach up to 2944 SAS drives with expansion enclosures per four-way clustered system, including the following:
 - SFF NVMe SCM drives (375GB to 1.6TB);
 - SFF NVMe FCMs (4.8TB to 38.4TB hardware compressed),
 - SFF NVMe SSDs (800GB to 30.7TB) and
 - SFF SAS SSDs (800GB to 30.8TB).



The **IBM FlashSystem 9200R** is a rack version of the FlashSystem 9200. IBM builds, tests, and installs the FlashSystem 9200R for configurations of two, three or four 9200 AFA storage system clusters, expansion enclosures, and fibre channel switches (for clustering) in an industry-standard rack.

NVMe Media Offerings for FlashSystem

The NVMe media offerings available on the FlashSystem mid-range and high-end enterprise storage systems require some explanation:

- **NVMe SCM** is an industry-leading, new storage media technology that offers extremely high endurance, very high IOPS and ultra-low latencies. IBM SCM is offered from two media vendors that use two distinct solid-state storage technologies. SCM is offered only on IBM FlashSystem 5100/5100H, 7200/7200H and 9200/9200R storage systems.



- **NVMe FCM** is an IBM-proprietary, NAND-based storage media that has been re-architected and re-designed to provide higher endurance than normal NAND SSDs; faster performance; and real-time, hardware-based data compression and encryption. IBM FCM media offers the highest density NAND-based storage available through the use of hardware compression. IBM FCMs are offered on FlashSystem 5100/5100H, 7200/7200H and 9200/9200R storage systems.

- **NVMe SSD** is an industry-standard NAND storage media that offers high IOPS and low-latency performance. Industry-standard NVMe SSDs are offered on FlashSystem 5100/5100H, 7200/7200H and 9200/9200R storage systems.

To perform at their best, NVMe SCM, FCM and SSD storage should be accessed via FC-NVMeoF host storage protocols. NVMeoF was designed to reduce IO overhead for SCM storage media but can also be used for NVMe FCMs and SSDs. NVMeoF/FC storage protocols are available on 16 or 32Gbps FC

IBM also offers iSER storage protocols on their entry and mid-range enterprise FlashSystem storage which provides another option to access NVMe storage faster than iSCSI alone. The iSER Ethernet protocol uses iWARP or RoCE hardware. While NVMeoF/FC is available on most current FC switch and director hardware, RoCE and iWARP for iSER may require upgrades to data center hardware, as not all Ethernet networking equipment supports RoCE or iWARP.

IBM FlashSystem storage options

The software at the heart of FlashSystem, **IBM Spectrum Virtualize**, offers a **clustering capability** that supports scale out from one to four Spectrum Virtualize [FlashSystem or SVC] storage systems, each of which can also scale up in capacity by adding more storage media, cache or host interfaces. It also offers **storage virtualization**, which allows external storage systems to be connected to and accessed through Spectrum Virtualize storage systems.

All IBM FlashSystem storage is designed to offer highly available “six nines” storage. For high availability **HyperSwap** systems installed by IBM Lab Services, IBM also offers an optional **High Availability Guarantee** which simplifies the support for high availability storage. FlashSystem data reduction levels (discussed below) are covered by the **IBM Data Reduction Guarantee**.¹

¹ Contact your IBM storage team to learn about additional IBM guarantees for FlashSystem storage.

IT has historically acquired storage by purchasing or leasing equipment. Lately, customers are realizing that a more flexible storage acquisition strategy can benefit the company's bottom line. As a result, IBM has made storage acquisition simple by adding the option to consume storage using an **IBM Storage Utility** or an **IBM Storage Subscription**. The IBM Storage Utility allows customers to pay for only the storage capacity used, while the IBM Storage Subscription plan is designed to reduce long-term, financial commitments for storage use.

IBM SAN Volume Controller engines

Like IBM FlashSystem, IBM SVC system nodes or engines run IBM Spectrum Virtualize software. SVC engines only support external storage. Other than boot drives, they do not support internal storage (unlike FlashSystem). IBM has introduced two new SVC nodes: the **SVC SA2 entry engine** and an **SVC SV2 high-performance engine**.

For availability purposes, the new engines are installed in pairs. Both new SVC engines offer 12 32Gbps FC/NVMe or 6 25Gbps Ethernet iSCSI host attachments and up to 768GB of cache per engine. The new engines can also be clustered with older SVC SV1 and DH8 engines, which is useful when non-disruptively upgrading older clusters without host interruption.

Both of the new SVC engines offer higher IO performance and, as a result, will support virtualization of higher performing enterprise storage than prior generation systems. The SVC SV2 high-performance engine offers up to 40% more IOPS at a similar price as the previous-generation system. The SVC SV2 entry engine offers up to a 25% lower cost of entry as compared to previous-generation SVC SV1 engines. (We will discuss storage system performance more fully below.)

IBM Spectrum Virtualize

As discussed above, IBM Spectrum Virtualize software offers a **clustered storage system** option. SVC and all IBM FlashSystem storage except FlashSystem 5010/5010H support clustering. While multiple generations of SVC engines or FlashSystem can be combined in the same cluster, IO performance may vary. Data centers that want more consistent IO performance should plan to use the same level of storage for all nodes in a Spectrum Virtualize storage cluster.

The performance of an all-IBM FlashSystem storage cluster should increase linearly as more equivalent storage nodes are added to a Spectrum Virtualize cluster. For instance, a two IBM FlashSystem 9200 system cluster should perform roughly twice as many IOPS as a single FlashSystem 9200 system cluster. (Note: Spectrum Virtualize storage clusters for FlashSystem 5030/5030H and 5100/5100H are limited to two FlashSystem storage systems. FlashSystem 7200/7200H,

FlashSystem 9200/9200R AFA and all SVC engines support up to four-way system clusters.)

As mentioned, Spectrum Virtualize offers **external storage virtualization**. FlashSystem storage from the 5100/5100H on up and all SVC nodes support this feature. Storage virtualization offers many advantages, not the least of which is that all external and internal FlashSystem storage can be migrated from one cluster node to another without host intervention. The ability to migrate data across the cluster without host intervention is a valuable feature for technical refresh activities.

Further, all storage virtualization customers can take advantage of Spectrum Virtualize enterprise class storage features such as advanced data reduction and storage tiering technologies (described below), as well as sophisticated caching to improve the functionality and performance of external storage. In addition, all internal and external storage capacity can be managed through the same console or single pane of glass. By doing so, IBM Spectrum Virtualize has simplified storage management for homogeneous and heterogeneous storage environments.

Spectrum Virtualize also offers a high level of integration with **VMware®**, including support for VMware vSphere® Virtual Volumes™ (vVols), vSphere APIs for Storage Awareness (VASA) I and II, vSphere Web Client, VMware vRealize® Orchestrator™ (vRO), vRealize Operations (vROPS) and VMware Site Recovery Manager™ (SRM). Similar levels of integration are available for Microsoft® Hyper-V™ environments.

Along with the extensive VMware integration, all IBM Spectrum Virtualize storage supports Kubernetes™ Container Storage Interface (CSI) for persistent volume snapshot offload and for use as container persistent volumes. Spectrum Virtualize will provide similar support for Red Hat® OpenShift™. Moreover, Spectrum Virtualize offers support for Red Hat Ansible® automation of IT infrastructure by managing hosts and volumes in a Red Hat OS environment.²

Spectrum Virtualize advanced data reduction and media features

IBM Spectrum Virtualize provides a number of data reduction technologies to reduce storage footprint, including the following:

- **SCSI UNMAP (flash delete)** command support sends deleted flash storage space back to the storage pool. IBM Spectrum Virtualize's SCSI UNMAP support will work for any host operating system that issues the command.
- **Controller software data compression** can be used to reduce the amount of space needed to store a block of data for non-FCM media.

² IBM Spectrum Virtualize storage software supports many other advanced, enterprise-class storage features not discussed here. Please contact your IBM Storage team or refer to Spectrum Virtualize documentation to learn more about its other capabilities.

- **Data deduplication** uses a block-based hash and specific pattern-matching functionality (e.g., all 0s/all 1s blocks) to reduce the physical size or footprint of stored data. IBM Spectrum Virtualize deduplication works cluster wide within pools, so any data reduction pool block can be used as a source for any other new block written to the pool.
- **Thin provisioning** consumes physical storage space only when data is written to a LUN or volume vs. when the storage is configured, which enables customers to share and pool unwritten storage space across LUNs within a pool.

Adding deduplication to IBM Spectrum Virtualize SCSI UNMAP, data compression (hardware for FCMs and controller software for non-FCM storage) and thin provisioning will enable customers to achieve up to a 5:1 reduction in the protected capacity required to store customer data.

In addition, Spectrum Virtualize **Easy Tier** offers automated, AI-based storage tiering across up to three levels of storage media to increase IO performance while also minimizing the cost of storage capacity. Easy Tier can be used to tier data across external storage as well as internal storage such as SFF or LFF SAS disk, SAS or NVMe SSD, and FCM and SCM storage media.

Easy Tier works by identifying the hottest (most frequently accessed) data in a storage pool and moving this hot data up to the highest performing storage. It also identifies the coldest (least frequently accessed) data and moves this cold data down to the lowest cost storage in a pool. Once configured and enabled, Easy Tier works automatically without host intervention.

On the other hand, Spectrum Virtualize enables customers to allocate LUNs or volumes to specific storage media, such as NVMe SSDs, FCM or SCM. Doing so can increase IO performance for applications using these volumes.

In sum, Spectrum Virtualize provides the flexibility needed to deliver the highest performance possible for enterprise storage: “set it and forget it” Easy Tier usage, automatic optimization of overall system performance and optimization of low-cost storage capacity or hand allocating application LUNs to highest performing storage media.

As mentioned, IBM FCMs support **onboard hardware, data-at-rest encryption**. For non-FCM media, Spectrum Virtualize also offers controller-based data encryption, which can secure internal, non-FCM FlashSystem storage as well as data on external virtualized storage.

Advanced data replication and data copy services

For customers that need the highest levels of business continuity/disaster recovery (BC/DR) resiliency, IBM Spectrum Virtualize offers two- and three-site advanced replication/mirroring support, including the following:

- **IBM Metro Mirror synchronous storage replication** – customers can configure two FlashSystem storage systems in a Metro Mirror connection over two sites to synchronously replicate data across metro-area-wide distances. Any write to a local Metro Mirror FlashSystem would have its data sent to the remote FlashSystem before the IO operation is acknowledged at the issuing host.
- **IBM Global Mirror asynchronous storage replication** – customers can configure two FlashSystem storage systems in a Global Mirror connection over two clusters a great distance apart in order to asynchronously replicate data across regions. Data is written to the local Global Mirror FlashSystem and the IO is completed on the local system before that data is sent to the remote system.

Both Metro Mirror and Global Mirror provide two-site replication, but they can also be combined for **three-site** replication. A **Global Mirror/Metro Mirror** replication solution has two systems in a Metro Mirror connection and another remote system at a third site in a Global Mirror connection with one of the Metro Mirror systems. This solution provides both within-region (Metro Mirror) and out-of-region (Global Mirror) BC/DR options to meet ultra-resilient customer needs.

In addition to advanced replication services, Spectrum Virtualize supports **FlashCopy data services**, which provide support for local (on the storage system), full or space-efficient copies of LUNs or volume data. FlashCopy can be used to take rapid copies of production data for use as backups or dev-test activities and can be deleted just as quickly when no longer needed. a

For multi-cloud and hybrid cloud services, IBM offers **Spectrum Virtualize for Public Cloud** in Amazon® Web Services™ (AWS™) which can be used to replicate or migrate data from on prem Spectrum Virtualize storage to the AWS public cloud. By providing Spectrum Virtualize for Public Cloud IBM has simplified the use of multi-/hybrid-cloud environments.

With data in IBM Spectrum Virtualize for Public Cloud, customers can snapshot replicated data to Amazon S3 to provide an “air gap” between their on prem and replicated cloud data for simpler cyber resiliency.

IBM Storage Insights

IBM Storage Insights is an optional cloud-based storage monitoring, trending and support solution that offers a single, unified view of IBM FlashSystems, SVC and

other storage system health. It uses new AI functionality to analyze and improve data center storage operations and provide proactive support for IBM customers throughout the world.

To deploy IBM Storage Insights, register for Storage Insights services, download and install the data collector and add the data center's storage configuration. The system immediately begins to monitor data center storage and provide insights into how customers can improve storage performance and operations.

IBM Storage Insights uses telemetry data sent by storage systems to IBM. For IBM FlashSystem storage, a wide range of system indicators and sensors is collected on an ongoing basis for both the storage system and its media. This data is sent to IBM. Similar data is collected for IBM SVC engines. Storage Insights uses this data to help ensure that IBM FlashSystem and SVC storage remain in optimal operational condition and to maximize system performance and data availability.

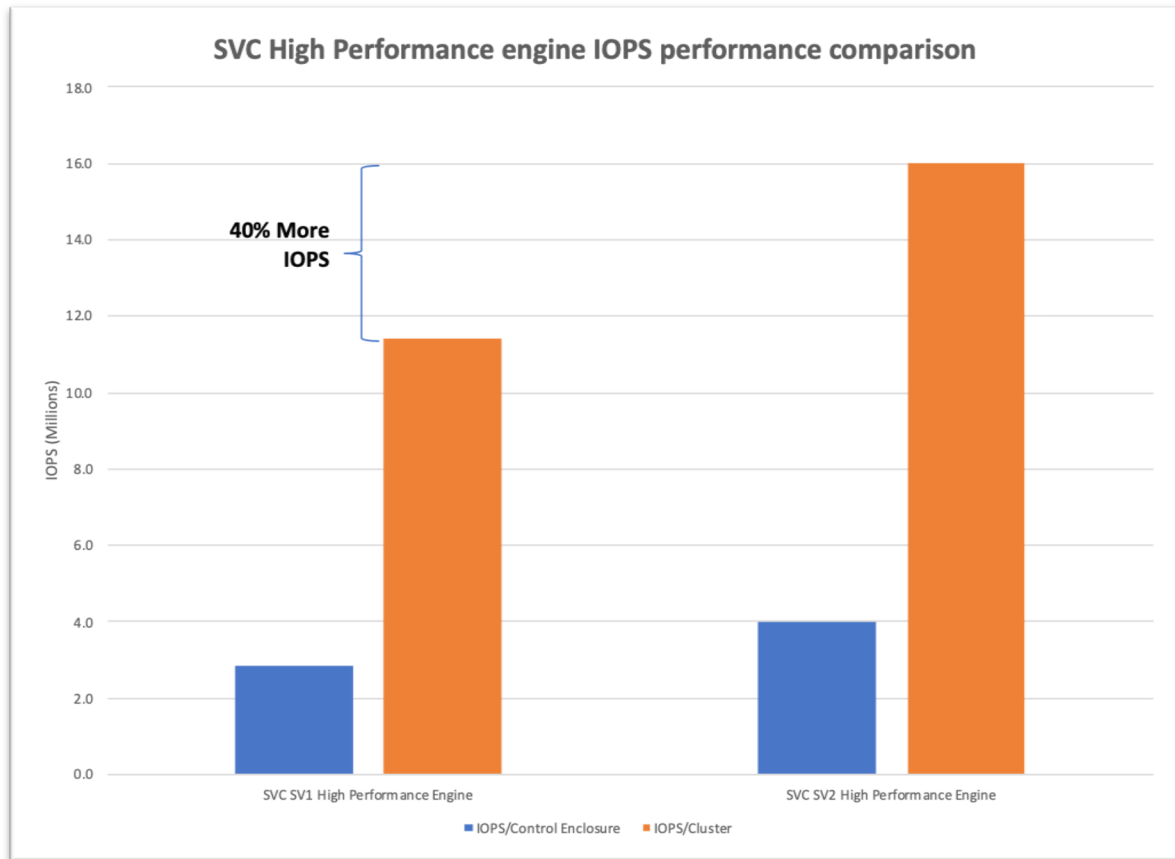
Storage systems use a one-way data transmission to the Storage Insights cloud, which is secured via HTTPS. All IBM Storage Insights data is also encrypted at rest with AES-256 cryptography and stored in the customer's own unique data store on the Storage Insights cloud. Only customer and IBM personnel have visibility into that customer's storage information in order to monitor storage health and operations and track system IO performance and capacity growth.

Furthermore, separate telemetry data and metadata are extracted from the customer's Storage Insights cloud for IBM internal use to facilitate field-wide analysis. This information is used to develop new storage system best practices and to proactively identify and resolve problems that could impact IBM storage customers worldwide.

IBM Storage Insights also offers support for non-IBM block storage, such as Dell EMC Unity and Unity XT storage, NetApp FAS and AFF storage and Hitachi Vantara VSP G Series storage systems.

IBM Storage Insights is available to IBM FlashSystem and SVC storage systems customers as both a free and a paid, licensed service. The free version covers all IBM block storage; the licensed version adds IBM object and cluster file storage, along with currently supported non-IBM block storage systems.

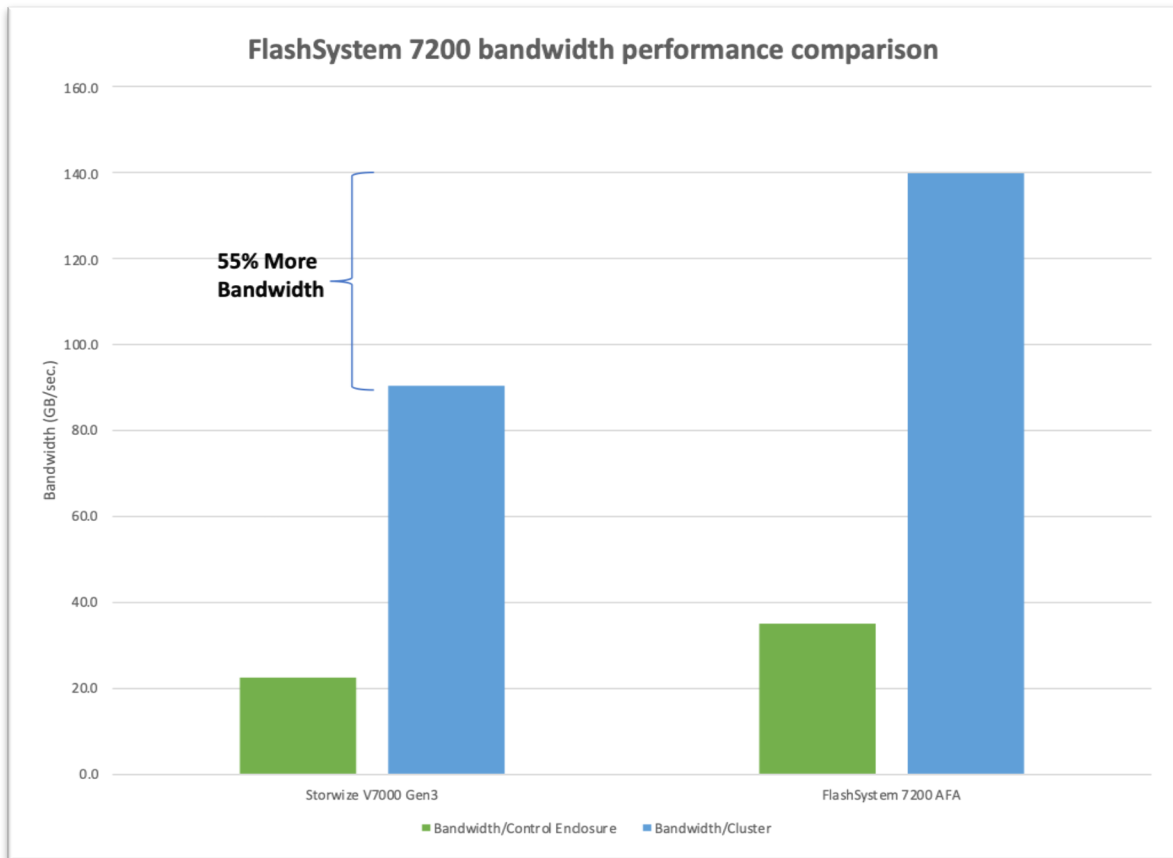
IBM SVC engine and FlashSystem performance



As discussed earlier, the new SVC SV2 high-performance engine offers 40% more IOPS than the previous-generation SVC SV1 high-performance engine at the same price. It offers up to 4M IOPS per control enclosure, or 16M IOPS per cluster with up to 48.8 GB per sec per control enclosure or 195 GB per sec per cluster of data bandwidth.

The new SVC SA2 entry engine offers up to 2.4M IOPS per control enclosure or 9.6M IOPS per cluster and up to 33 GB per sec per control enclosure or 132 GB per sec per cluster of data bandwidth.

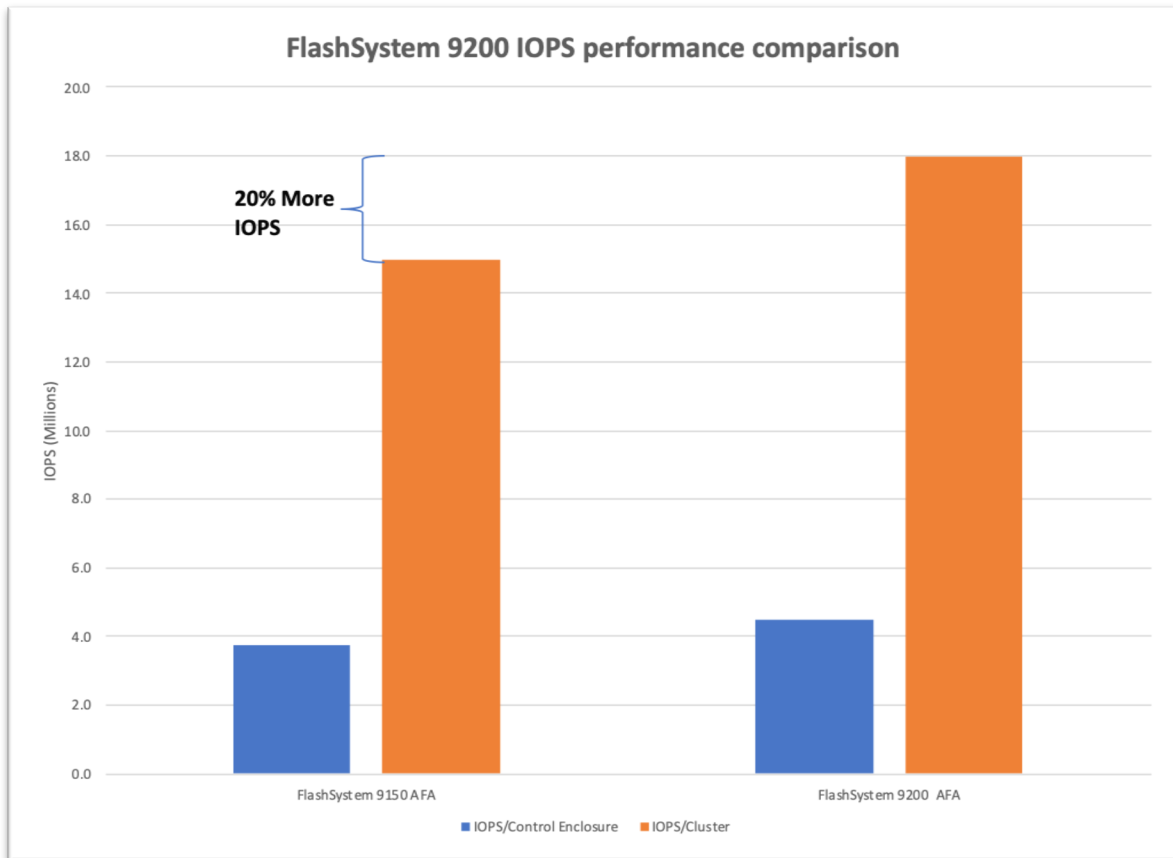
These SVC performance numbers are based on using external storage with FCM media and 32Gbps FC host access.



The FlashSystem 7200 AFA offers 43% more IOPS than the previous-generation system. It offers up to 2.4M IOPS per control enclosure or 9.2M IOPS per cluster. On top of that, the new FlashSystem 7200 AFA offers 55% better throughput, or 35 GB per sec per control enclosure or 140 GB per sec per cluster of data bandwidth.³

The FlashSystem 5100/5100H, new FlashSystem 7200/7200H and FlashSystem 9200/9200R support extremely low latency IO, i.e., down to a **70µsec** response time using NVMeoF/FC protocols with FCM media.

³ All FlashSystem performance numbers are based on using internal FCM storage media and 32Gbps FC.



The new FlashSystem 9200/9200R AFA offers up to 20% more IOPS at 4.5M IOPS per control enclosure or 18M IOPS per cluster, along with 20% more bandwidth at 45 GB per sec per control enclosure or 180 GB per sec per cluster, than the prior-generation (FlashSystem 9150) system.

Summary

IBM's new FlashSystem storage and SVC engines provide more economical storage as well as higher performance than previous-generation systems.

For mid-range and high-end enterprise storage, the introduction of new SCM media, higher density FCM storage and industry standard NVMe SSDs together with NVMeoF/FC host protocol gives customers the ultimate choice in storage performance, footprint and capacity cost. Indeed, IBM offers both the highest density flash (FCM) and the highest performing media (SCM) in the same storage system. The inclusion of NVMe SSDs makes high performing, FlashSystem storage even more affordable.

Easy Tier automation makes all these media types very easy to use by offering completely automatic performance tiering across media. Easy Tier also supports

FlashSystem hybrid storage ease of use, performance and economical storage capacity.

Furthermore, the enhancements to the new generation SVC engines provide a lower entry cost for IBM's new SVC engine and higher performing storage.

IBM continues to introduce higher performing FlashSystem storage and SVC engine hardware, while Spectrum Virtualize software continues to set the bar with its enterprise-class storage functionality. Indeed, Easy Tier, sophisticated data reduction, multi-site replication and other advanced Spectrum Virtualize capabilities provide FlashSystem storage and SVC engines with the high-end capabilities that we have come to expect from IBM enterprise storage.

If you are adding storage to your data center environment and want enterprise-class functionality, you can't go wrong with IBM Spectrum Virtualize FlashSystem or SVC engine storage, the simplest, lowest entry cost and highest performing storage options on the market today.

Silverton Consulting, Inc., is a U.S.-based Storage, Strategy & Systems consulting firm offering products and services to the data storage community.



Disclaimer: This document was developed with International Business Machines Corporation (IBM) funding. Although the document uses publicly available material from various sources, including IBM, it does not necessarily reflect the positions of such sources on the issues addressed.