



IBM FlashSystem 9100 Storage Family

Silverton Consulting, Inc. StorInt™ Briefing

Introduction

The new IBM® FlashSystem® 9100 storage adds another all flash array to the IBM FlashSystem product family. IBM FlashSystem 9100 storage takes advantage of the latest drive technology, packaging and Intel® processors to improve performance and make data storage more cost effective.

One key storage technology added to IBM FlashSystem 9100 is the NVMe Express™ (NVMe™) drive interface. IBM FlashSystem 9100 can make use of a new generation of FlashCore storage modules (FCM) with hardware data compression, that's been re-designed for NVMe and reconfigured into a small form factor envelope or small form factor, industry standard, NVMe SSDs. IBM FlashSystem 9100 exploits NVMe for both IBM FCMs and SSDs to reduce IO latency and improve performance. Moreover, the new NVMe interface enables IBM FlashSystem 9100 to be ready for Storage Class Memory (SCM) SSDs once they become more prominent.

IBM FlashSystem 9100

IBM FlashSystem V9000 was a 6U-tall storage system. In contrast, the new IBM FlashSystem 9100 comes in a 2U-configuration with optional storage expansion. The new FlashSystem supports up to 24 small form factor (SFF) 2.5" drives and from 128 GB to 1,536 GB of memory per controller. The new hardware has been optimized to support higher memory bandwidth, faster IO performance and IBM Spectrum Virtualize™ data services.



The new IBM FlashSystem 9100 includes four onboard 10GbE ports and comes in two system versions:

- **IBM FlashSystem 9110** supports dual 8-core Intel Skylake CPUs per canister, and
- **IBM FlashSystem 9150** supports dual 14-core Intel Skylake CPUs per canister.

IBM FlashSystem 9110 supports the same capacity and other configuration options as IBM FlashSystem 9150 but with less compute power available to sustain IO operations and data services. With more compute cores, IBM FlashSystem 9150 will support greater IO performance and faster data services for diverse customer requirements.

As discussed earlier, both new FlashSystems have adopted the latest-generation NVMe drive interface for storage drives. That is, all internal SFF drive slots use

NVMe for both the new IBM-engineered FCMs as well as the industry-standard SSDs. IBM FlashSystem 9100 NVMe drives available for internal slot use include:

- **IBM FCMs** designed for a 2.5" SFF using the NVMe drive interface with on-the-FCM, performance-neutral hardware compression and FIPS 140-2 Level 1 encryption at 4.8TB, 9.6TB or 19.2TB of raw flash storage; and
- **Industry-standard SSDs** in a 2.5" SFF using the NVMe drive interface with optional, on-the-drive, FIPS-140-2 Level 1 encryption at 1.92TB, 3.84TB, 7.68TB or 15.36TB of raw flash storage.



IBM FlashSystem 9100 encryption key management can make use of IBM Security Key Lifecycle Manager as well as USB-based secret keys.

The new FCMs are the first FlashCore storage designed to use an SFF drive envelope and an NVMe drive interface. As with all FlashCore storage, the new FCMs incorporate advanced, proprietary IBM engineering that reduces IO latency, increases flash capacity/durability and enhances storage availability. Proprietary enhancements include FCM hardware-based accelerated architecture, IBM microlatency technology and advanced flash management.

Within the internal 24 SFF drive slots, customers can intermix NVMe FCMs and NVMe SSDs in the same system. This way, customers can tailor their IBM FlashSystem 9100 media to support the performance and capacity needs of their data center environment.

The new IBM FlashSystem 9100s offer up to 460TB (24 19.2TB FCM drives) of raw capacity in the 2U appliance. Moreover, IBM FlashSystem 9100 supports additional storage capacity expansion cabinets using Serial-Attached SCSI (SAS) SSDs. Two storage expansion cabinet shelf configurations are available:

- **2U 24 SAS SSD** expansion shelf, or
- **5U 92 SAS SSD** expansion shelf.

As IBM FlashSystem 9100 offers NVMe FCMs, NVMe SSDs and SAS SSDs (with optional expansion), different levels of media performance can be supplied in the same system. IBM has made available **Easy Tier**[®] automated storage tiering on its block storage systems for years now. Easy Tier can automatically move active data to faster performing storage media and less active data to slower performing storage. (The fastest and most expensive media should be reserved for data with the highest access frequency.)

In addition, IBM FlashSystem 9100 supports Spectrum Virtualize's new data reduction pools that incorporate thin provisioning, compression and data deduplication. Alternatively, IBM FlashSystem 9100 customers with FCM's automatically receive hardware compression for maximum IO performance. With FCM media, IBM FlashSystem 9100 is designed to be IO performance equivalent to the FlashSystem V9000.

Recall that each IBM FlashSystem 9100 controller comes with eight onboard 10GbE



ports. IBM FlashSystem 9100 appliances also offer three IO slots per canister, any of which can be configured as follows:

- **4-port 16Gbs FC** with a maximum of 24 ports per controller,
- **2-port 25GbE** with iWARP or RoCE compatibility with a maximum of 12 ports per controller, and
- **2-port 12Gbs SAS** connections for expansion storage.

With the optional three IO slots per canister, customers have a great deal of flexibility in how they connect their host servers to IBM FlashSystem 9100 appliances.

Why NVMe

NVMe is a new drive interface designed by storage industry heavyweights to lower IO overhead and thereby reduce IO latency for next generation, ultra-fast devices. One challenge with SAS and SATA SCSI drive protocols is the amount of chatter and number of interface transactions required to perform an IO operation. This excess chatter takes time to complete and delays IO execution for really fast devices such as FCMs.

The new NVMe device protocol was specifically designed to reduce IO operation protocol chatter. Moreover, the NVMe protocol also includes up to 4KB of data in a write command request. That is, writes require fewer interface transactions to move up to 4K of data from the storage system to the NVMe device.

Similar optimizations reduce the overhead required to perform an NVMe read operation. Even the status returned at the end of a drive IO operation shows reduced latency for NVMe devices.

Furthermore, on IBM FlashSystem 9100, NVMe FCMs and SSDs attach directly to the server PCI Express™ (PCIe™) bus, which itself can roughly double the bandwidth of SAS drives. With these and other speed optimizations, some flash drive vendors show an almost 50 μsec or 33% reduction in response time (IO latency) for NVMe SSDs as compared to SAS SSDs

Another advantage of NVMe is that the industry is starting to define a new host-to-storage-system protocol that extends NVMe device latency reductions all the way back to the host. The industry term for this new protocol is NVMe over Fabric (NVMeoF). While the new NVMeoF protocol is still being defined, it has the potential to radically reduce the NVMe storage IO latencies for the already lightning fast IBM FlashSystem 9100.

The NVMeoF protocol will operate with RDMA-compatible Ethernet (RoCE or iWARP), current and previous-generation FC and InfiniBand fabric interfaces. IBM FlashSystem 9100 FC and Ethernet IO ports are already **NVMeoF ready**. That is, once customer servers can support NVMeoF, IBM FlashSystem 9100 will be ready to use NVMeoF for data access and provide even faster response times.

Spectrum Virtualize functionality

IBM FlashSystem 9100, like the FlashSystem V9000 and Storwize products uses Spectrum Virtualize functionality to provide data services. Spectrum Virtualize offers a number of advanced, enterprise-class storage features and capabilities, some of which are discussed below.

While IBM FlashSystem V9000 offered RAID-5 data protection, the FlashSystem 9100 offers DRAID-6 for customers that want to protect data from loss due to a second SSD or FCM drive failure.

IBM Spectrum Virtualize recently added SCSI UNMAP (flash delete) command support that frees up deleted drive flash storage space back to the drive flash storage pool. IBM Spectrum Virtualize's SCSI UNMAP support will work for any host operating system that issues the command. SCSI UNMAP adds to IBM Spectrum Virtualize's other data reduction capabilities. Coincident with the release of IBM FlashSystem 9100, Spectrum Virtualize now also offers data deduplication. For volumes/LUNS residing in IBM Spectrum Virtualize data reduction pools, data deduplication uses a block-based hash as well as specific pattern-matching functionality (e.g., looking for all 0s/1s blocks) to reduce the size of data stored.

With deduplication, IBM Spectrum Virtualize employs metadata to recognize if a block-based hash matches a previously stored data block and discards any duplicate data. IBM Spectrum Virtualize deduplication works cluster wide, so any data reduction pool block can be used as a source for any other new block written to the pool. Adding deduplication to the already present SCSI UNMAP, data compression and thin provisioning in IBM Spectrum Virtualize will enable some customers to achieve up to a 5:1 reduction in the protected capacity required to store a data load.

IBM Spectrum Virtualize is a clustered storage system. As such, any IBM FlashSystem 9100 storage can combine with other FlashSystem 9100 and Storwize V7000 Gen 2 and Gen2+ storage systems to create a single cluster of storage. Performance of an IBM FlashSystem 9100 storage cluster should increase linearly as customers add nodes to the cluster. For instance, a two-controller IBM FlashSystem

9100 cluster should perform roughly twice as many IOs per second as a single controller FlashSystem 9100 cluster.

Spectrum Virtualize has many other valuable features, such as:

- **Storage virtualization** – support for external storage arrays connected to the host through and managed by Spectrum Virtualize storage systems.
- **VMware interoperability** – support for a high level of VMware integration including VMware Virtual Volumes (VVols), vSphere APIs for Storage Awareness (VASA) I & II, vSphere Web Client, vRealize Orchestrator (vRO)-Operations (vROPS) manager and Site Recovery Manager (SRM).
- **Global and Metro Mirror replication** – support for disaster recovery by replicating Spectrum Virtualize cluster volumes to other Spectrum Virtualize systems located at remote sites.
- **Data at rest security** – support for encryption of data at rest using AES cryptography to insure customer data remains secured.

IBM Spectrum Storage Insights

IBM Spectrum Storage Insights is an optional cloud-based storage monitoring, trending and support solution that offers a single, unified view of IBM storage system health. It uses new cognitive functionality (IBM Watson™ and other AI methods) to analyze and improve data center storage operations and provide proactive support for IBM storage customers throughout the world.

Deploying IBM Spectrum Storage Insights is as easy as downloading, registering for Storage Insight services, installing the data collector software and adding the appropriate storage configuration. The system immediately begins to monitor storage operations and provide insights into how customers can improve storage performance and operations.

IBM Spectrum Storage Insights uses one-way data transmission to the Storage Insights cloud and is secured via HTTPS. All IBM Spectrum Storage Insights data is encrypted with AES-256 cryptography and stored in the customer's own, unique data store on the Storage Insights cloud. This way, only customer and IBM personnel have access and visibility to that customer's storage data in order to monitor storage health and operations and track performance and capacity growth.

Furthermore, separate data and metadata is also extracted for IBM internal use to facilitate field-wide analysis. This data is used to develop new storage system best practices, proactively identify and resolve problems that could impact other customers and provide support for IBM storage customers worldwide.

IBM Spectrum Storage Insights uses IBM storage system's telemetry data sent by storage systems to IBM. For IBM FlashSystem 9100, a wide range of system indicators and sensors are collected on an ongoing basis and sent to IBM for Storage

Insights analysis. This analysis helps to ensure that IBM FlashSystem 9100 storage remains in optimal operational condition to maximize system performance and data availability.

IBM Spectrum Storage Insights is available to IBM FlashSystem 9100 customers called **IBM Spectrum Storage Insights**, and a paid, licensed service. The free version covers all IBM block storage, and the licensed version adds IBM object and cluster file storage, along with some non-IBM-vendor block storage systems.

Summary

The new IBM FlashSystem 9100 represents the next evolution in flash storage performance and functionality, using the newest industry-standard NVMe drive interface, latest-generation media and latest Intel processors to increase raw capacity, reduce latency and provide high throughput/IOPS performance.

The new FCMs, with their highly available and extremely fast FlashCore technology, add the latest NVMe protocol in an SFF envelope, along with always-on, performance-neutral hardware compression. This combination of capabilities and features makes IBM's new FCM among the highest performing, most reliable flash drives available today.

Moreover, merging NVMe FCMs and industry-standard NVMe SSDs with SAS SSDs using Easy Tier functionality gives customers significant choice as how to configure their IBM FlashSystem 9100 media to match data center requirements. Adding the cloud-based, Spectrum Storage Insights solution makes for an optimal IBM FlashSystem 9100 storage operation.

Finally, combining IBM FlashSystem 9100 controller hardware, NVMe FCM, NVMe SSD and SAS SSD media with Spectrum Virtualize enterprise-class functionality in one system presents an attractive, small-footprint, advanced storage capability that any customer could use in their storage environment.

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.