

## IDC PERSPECTIVE

# IBM Steps into NVMe in a Big Way with the New FlashSystem 9100

Eric Burgener

## EXECUTIVE SNAPSHOT

---

### FIGURE 1

---

#### Executive Snapshot: IBM FlashSystem 9100 Announcement

In July 2018, IBM made a bold set of upgrades with the IBM FlashSystem 9100, its flagship dual controller enterprise array design targeted for consolidation of block-based workloads. The changes are significant as they represent a move toward more standardized designs in terms of both NVMe and storage array packaging. While other software-based products were also announced, changes in the FlashSystem 9100 stand out as a harbinger of the future direction of IBM all-flash array (AFA) evolution.

#### Key Takeaways

- IBM has announced its first true enterprise-class end-to-end NVMe-based general-purpose storage platform with the IBM FlashSystem 9100, delivering three times the performance, six times the storage density, 60% less energy consumption, and a 66% smaller rack space footprint than the IBM AFA that it will ultimately replace (the FlashSystem V9000).
- Additional enhancements to IBM's hybrid cloud integration capabilities drive additional value for customers.
- IBM also introduced the IBM FlashWatch program, which is focused around delivering a better customer experience by offering a number of up-front guarantees around the 9100 platform.

#### Recommended Actions

- Customers looking to share high-performance NVMe-based storage with very efficient utilization and comprehensive and proven enterprise-class data services should consider this system (which includes an NVMe over Fabric connection based on InfiniBand).
- Existing IBM customers can either continue to expand IBM FlashSystem V9000 configurations with that SCSI-based platform or, if NVMe performance is required, move to the new FlashSystem 9100 without having to learn new software (it runs the same Spectrum Virtualize software that the V9000 does).
- Customers should be thinking about how they will be integrating NVMe technology into their storage infrastructures over the next several years without having to disrupt application services.

Source: IDC, 2018

## SITUATION OVERVIEW

---

The all-flash array (AFA) market has grown rapidly over the past seven years. Initial shipments of enterprise-class all-flash storage appliances began around 2011, and by 2017 these appliances were available from all established enterprise storage vendors. In that same year, AFAs drove over 70% of all external primary storage revenue. Most of these systems were built around SCSI technologies, although there were several vendors shipping higher-performance storage platforms built around PCIe technology. As the storage demands of enterprises undergoing digital transformation became increasingly stringent, we saw NVMe technology begin to appear in general-purpose AFAs intended for mixed enterprise workload consolidation. Although there are a number of start-ups selling NVMe-based storage arrays, the major vendors have really started to hit their stride in 2018 with NVMe announcements around their flagship storage platforms.

IBM has had an interesting history in the AFA space. Its first entry in the market was with the IBM FlashSystem 800 back in 2013. This system leveraged custom PCIe hardware and was recognized for delivering what was arguably one of the lowest storage latencies in a shared storage, true enterprise-class offering. The system also included custom flash modules, which at the time delivered both higher performance and higher storage densities than were available with off-the-shelf SAS solid state disks (SSDs). PCIe was a storage protocol that was developed specifically for flash storage (not spinning disks) and therefore enabled the development of higher-performance and more efficient storage systems than those based around SCSI, but the PCIe industry standard was intended more for use in local rather than networked storage configurations. Because of this focus, building an enterprise-class array around PCIe technology required significant custom content. In the meantime, another flash-specific I/O protocol that was more geared toward enterprise storage systems use was developing (called NVMe), but it was not until recently that it became mature enough as a standard for the established storage providers to build systems around it. Interestingly, systems like IBM's FlashSystem, built as it was around IBM's enterprise-enhanced version of PCIe, delivered most of the performance benefits of NVMe-based systems but were based on custom rather than standard hardware.

For building high-performance storage systems that have to efficiently deliver consistently low latency and high throughput for mixed enterprise workloads, there is no doubt that NVMe is a better core technology than SCSI. NVMe is specifically developed for flash storage, allowing it to make much more efficient use of flash performance and capacity, and it will be the I/O protocol of choice to support new emerging memory technologies like storage-class memory going forward. NVMe supports generally at least 50% lower latencies on a per-device basis and up to an order of magnitude higher bandwidth and throughput. It offers three orders of magnitude higher parallelism as well, and this ability to handle up to 64,000 outstanding I/O requests makes it much better for building very large-scale systems that must handle a lot of random I/O. IDC believes that by 2021, NVMe-based storage platforms will drive over 50% of all external primary storage revenue, making it the heir apparent to SCSI for primary storage systems.

Arrays leveraging NVMe technology feature NVMe-based controllers, backplanes, and storage devices, but to make sure all this performance gets delivered to applications from networked storage requires NVMe over Fabric (NVMe-oF). NVMe-oF is a remote direct memory access (RDMA) protocol that allows NVMe to run over switched fabrics with basically the same latencies as local NVMe-based SSDs in a server. NVMe-oF is critical because it allows high-performance NVMe storage to be shared across literally hundreds of servers accessible across a switched fabric and allows the other benefits of enterprise-class networked storage (scalability, high availability, enterprise-class data services, and

the capacity utilization efficiencies of shared storage) to be applied to it as well (features typically not available in local storage configurations). The NVMe-oF specification allows for this RDMA protocol to run over Fibre Channel (FC), Ethernet, or InfiniBand host connections.

IBM has been vocal about its NVMe vision and product introductions. In May 2017, IBM announced that it would be leveraging NVMe technologies across both IBM's on-premise infrastructure and cloud-based offerings going forward. In December 2017, IBM showed the first NVMe-oF host connection connected to a true enterprise-class AFA, demonstrating an NVMe over InfiniBand connection between an IBM FlashSystem 900 and an IBM Power9 Systems server. In February 2018, IBM announced the general availability of this implementation, along with software support for NVMe-oF for additional IBM products, including the IBM Spectrum Accelerate and Virtualize families (and derivative products like the V9000 and A9000) as well as the Storwize product line. The latest announcement introduced the FlashSystem 9100, IBM's new AFA built around NVMe rather than PCIe technologies, available directly from the factory with NVMe-oF support (over InfiniBand) – a true enterprise-class, end-to-end NVMe-based storage solution.

## The IBM FlashSystem 9100

The July announcements covered systems as well as hardware and software products, but the FlashSystem 9100 launch is the cornerstone. In a 2U package, the 9100 supports up to 2.5 million IOPS, up to 460TB of raw capacity, up to 34Gbps of bandwidth, and within the array itself it is built entirely around NVMe rather than PCIe or SCSI. This is not an NVMe "turn" to the older FlashSystem V9000 (or FlashSystem 900) hardware – this is an entirely new, more efficient design that is optimized specifically for NVMe.

The FlashSystem 9100 is based on dual active/active controllers, and its enterprise-class data services are provided by Spectrum Virtualize (which now runs natively on the array). These are proven and mature, block-based enterprise-class data services that include inline data reduction (compression and deduplication), dual parity RAID, thin provisioning, space-efficient read/write snapshots, encryption, nondisruptive data migration, quality of service, and replication options (including asynchronous, synchronous, and stretch clusters, a feature IBM refers to as HyperSwap). The Spectrum Virtualize software supports storage tiering within the array, to other arrays (including non-IBM ones), and to multiple public clouds, and it supports 440 different types of heterogeneous storage arrays. Even with these features enabled, the FlashSystem 9100 is able to consistently deliver sub-100 microsecond latencies under load. IBM is also bundling additional software with the base purchase price, including IBM Spectrum Copy Data Management (a 5TB license), IBM Spectrum Protect Plus (an integrated modern data protection product), IBM Spectrum Virtualize for Public Cloud (a software-only version of the IBM Spectrum Virtualize product that can be run in the public cloud), and IBM Spectrum Connect (IBM's hybrid cloud deployment suite that supports container environments based on Docker and Kubernetes).

The packaging implementation is significant for IBM, driven by the maturing standards for high-performance AFAs. For years, IBM was able to deliver NVMe performance with its customized PCIe hardware, including its IBM FlashCore Modules, before other vendors could. These earlier systems used custom flash modules (the FlashCore Modules) instead of off-the-shelf SSDs, delivering higher performance, higher density, better reliability, dual-port capability, and lower cost/gigabyte (GB). The FlashCore Modules also support some additional onboard features, such as variable stripe hardware RAID, compression, encryption, and optimized wear leveling that allow those data services to be used

without impacting application performance. The FlashCore Modules are also certified as FIPS 140-2 compliant.

As the NVMe standard has matured, however, the differences between FlashCore Modules and SSDs have narrowed, and so with the 9100, IBM is offering customers the option to configure the system around either FlashCore Modules or off-the-shelf, dual-ported NVMe SSDs. Today, systems built around the FlashCore Modules are still likely to have higher performance, higher density, better reliability, and lower cost/GB, but that may not necessarily be true a year from now. The 2U24 packaging accepts 2.5in. form factor devices, giving it the ability to support either FlashCore Modules (available in 4.8TB, 9.6TB, or 19.2TB models) or dual-ported, commodity off-the-shelf NVMe SSDs (available in 2TB, 8TB, or 15.3TB models). This approach provides a smooth transition path to the NVMe standard, giving customers the option to select the storage media packaging that best meets their requirements.

FlashSystem 9100 nodes can be clustered in up to a four-way configuration, which supports 10 million IOPS, 136GBps of bandwidth, and up to 10PB of effective storage capacity (assuming 5:1 data reduction) in only 10U (8U for the cluster nodes, 2U for the intracluster network infrastructure). The system supports additional disk expansion cabinets that allow it to be scaled to support up to 32PB of effective storage capacity in a single 42U cabinet, although these disk shelves are connected over 12Gb SAS rather than NVMe; so this additional storage capacity will be composed of lower-performance storage (either SAS, SSDs, or HDDs). Cluster nodes are connected over redundant 10GbE networks, and the arrays support SCSI-based FC and Ethernet as well as NVMe over InfiniBand host connections.

IBM also introduced the FlashWatch program along with the FlashSystem 9100. Similar to other competitive programs like Dell EMC's Future-Proof Loyalty Program, it wraps a number of features and guarantees into an umbrella program that makes it easier for customers to understand the full value IBM is providing. Customer experience (CX) programs like IBM FlashWatch are targeted to drive additional value for customers in terms of a better customer experience. Bundled at no extra charge with all FlashSystem 9100 systems under a standard maintenance contract, it includes a data reduction guarantee (validated by IBM, can be up to 5:1), a 100% data availability guarantee (for systems configured with HyperSwap and deployed according to IBM best practices), a controller upgrade program, seven years of 24 x 7 support, a lifetime flash media endurance guarantee, and up to 45 days of nondisruptive data migration services (which cover both IBM and non-IBM source platforms). Note that there is an additional fee for the HyperSwap product. Cloud-based predictive analytics and comprehensive software bundling are also components of competitive programs of this nature, and IBM has done that as well for the FlashSystem 9100 with the inclusion of IBM Storage Insights (IBM's artificial intelligence-based storage analytics, storage resource management, and remote support platform) along with Spectrum Virtualize, as well as the additional Spectrum software components mentioned previously.

Other announcements around the FlashSystem 9100 included an additional set of cloud integration capabilities, including IBM-validated blueprints (called IBM Spectrum Access Blueprint) that make it easy to deploy hybrid cloud configurations with confidence; support for Docker and Kubernetes technology in private clouds (IBM Cloud Private and others); and modern data protection, secondary data reuse, and business continuity solutions. Although these capabilities were not new with this FlashSystem announcement, IBM also pointed out that it offers storage utility pricing models for many of its AFAs, including the new FlashSystem 9100, that give customers the option to change capex into opex when this best meets their needs.

## Analysis

The new hardware associated with this announcement represents a harbinger of what will come from IBM. This move to more standardized hardware puts the company on a path to ultimately benefit from cost decreases without giving up performance or scalability as NVMe and the new 2U24 designs hit higher volumes. IDC expects to see this new hardware become available in other FlashSystem designs going forward. In making this change, however, IBM is not discontinuing the V9000, making it easy for customers that have already qualified that hardware and want to expand existing systems to continue to do so without undergoing new qualification efforts. Offering customers the option to build the FlashSystem 9100 with either FlashCore Modules or SSDs provides customers the choice to select which of those approaches best meets their needs today while leaving the option open to move to more standardized hardware if they want to.

Throughout the history of the IT industry, custom hardware is often used when it provides performance, scalability, or functionality advantages over commodity hardware, but vendors that take that path must continue to reevaluate that decision with each new generation of systems. As standards evolve and commodity hardware "catches up" to custom designs, it is often in the best interest of vendors and customers alike to move to more standardized solutions. There are still examples of custom hardware among the leading AFA vendors – HPE uses a custom ASIC in its controller mesh, and Pure Storage is actually moving more in the direction of custom hardware over time. Hitachi Vantara has taken an approach similar to IBM in that it is moving away from custom PCIe-based hardware content for flash devices and disk shelving interconnects to more commodity hardware while leaving both options (custom flash modules and SSDs) open to customers. Other AFA vendors like Dell EMC, NetApp, and Kaminario, have largely moved to commodity hardware, placing most of their differentiating functionality in software.

This change also represents a further move in the direction of software-defined storage for IBM's most powerful AFAs. Much of the performance and scalability benefit of the prior generation of systems based around the FlashSystem 900 were dependent upon custom hardware, but technology advancements in CPUs, SSDs, and packaging allow IBM to move to more standardized hardware and move more functionality to software. Ultimately, that is where the industry is headed for general-purpose enterprise-class storage arrays, and IDC expects that this approach will lead to more IT configuration flexibility and lower costs over time. IBM already has a rich portfolio of software-defined storage solutions and, in fact, is an industry leader in this space in terms of revenue generation. But these changes with the FlashSystem 9100 strengthen that focus.

The hardware and software enhancements included with the new IBM FlashSystem 9100 deliver noticeably better total cost of ownership (TCO) relative to prior-generation systems. In comparison with the prior generation of V9000 systems, the 9100 offers three times the performance, six times the storage density, 60% less energy consumption, and requires 66% less hardware (for floor space savings) – financial metrics that are critical for both enterprises and cloud providers alike that want to build very dense IT infrastructure.

In conclusion, these announcements definitely move IBM in a positive direction, deliver more value for its customers, and further improve the advantages IBM offers in implementing and managing hybrid cloud environments. IBM has offered much of the same functionality in both guarantees and cloud integration as some of the industry leaders, but the lack of a unifying program to communicate the relevant features and benefits has made it more difficult for customers to see the value IBM is providing in these areas. With the introduction of the IBM FlashWatch program and by assembling

hybrid cloud integration capabilities under the unifying moniker of IBM multicloud solutions, IBM is making it easier for customers to understand just how it stacks up against the competition.

## ADVICE FOR THE TECHNOLOGY BUYER

---

- For current non-IBM customers, the IBM FlashSystem 9100 is a very competitive offering that supports end-to-end NVMe, including an NVMe-oF host connection based around InfiniBand; customers that want local storage latencies in a shared storage solution with comprehensive and proven enterprise-class data services should consider this system.
- For current IBM FlashSystem V9000 customers looking to expand existing configurations without having to qualify new systems, they can continue to buy the IBM FlashSystem V9000, but if they want NVMe performance, they can consider adding the 9100 while not having to learn new software (it runs the same Spectrum Virtualize software that the V9000 runs, but it runs it natively on the array).
- The new 2U24 hardware on the FlashSystem 9100 is an updated, more standardized design that IBM is likely to leverage for other all-flash systems as they continue to evolve; awareness of this may influence purchase decisions.

## LEARN MORE

---

### Related Research

- *A Deluge of NVMe-Based Announcements Gives Enterprise Storage Customers Attractive Higher-Performance Options* (IDC #US43940318, June 2018)
- *IDC TechBrief: NVMe Over Fabric* (IDC #US43854018, June 2018)
- *IDC MarketScape: Worldwide All-Flash Array 2017 Vendor Assessment* (IDC #US43310017, December 2017)

### Synopsis

This IDC Perspective briefly summarizes market trends that IBM's FlashSystem 9100 announcement addresses, provides an overview of the platform, and discusses the IBM platform enhancements and changes in terms of market significance. In July 2018, IBM made some key changes to its enterprise storage platform portfolio built around key new storage technologies like NVMe, artificial intelligence, software-defined storage, and multicloud. A major focus of the announcement was IBM's new FlashSystem 9100. A new CX offering, the IBM FlashWatch program, was also announced and should drive additional value for IBM AFA customers.

"Enterprises undergoing digital transformation are building and integrating real-time workloads that drive increasingly stringent performance requirements," said Eric Burgener, research vice president of Storage. "As a result, we are seeing established enterprise storage industry leaders like IBM integrate high-performance NVMe technologies into its flagship platform offerings."

## About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

## Global Headquarters

5 Speen Street  
Framingham, MA 01701  
USA  
508.872.8200  
Twitter: @IDC  
idc-community.com  
www.idc.com

---

### Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit [www.idc.com](http://www.idc.com) to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit [www.idc.com/offices](http://www.idc.com/offices). Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or [sales@idc.com](mailto:sales@idc.com) for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights.

Copyright 2018 IDC. Reproduction is forbidden unless authorized. All rights reserved.

