2020-21 ENTERPRISE ALL-FLASH ARRAY BUYER'S GUIDE

The Insider's Guide to Evaluating Enterprise All-Flash Arrays

The End-to-end NVMe Opportunity

The first-generation all-flash arrays (AFAs) dramatically improved the performance density of enterprise data centers. For many enterprises, these AFAs provided a 10x performance improvement compared to legacy HDD-based storage systems. AFAs also eliminated many storage performance problems and the data center management overhead associated with avoiding or addressing performance issues in legacy arrays.

Most of the first-generation AFAs connected to flash memory via disk-oriented SATA or SAS protocols. These protocol stacks include many features that do not apply to flash memory. Further, they do not take advantage of the parallelism and other capabilities of flash memory and other non-volatile memories.

Storage industry experts recognized the deficiencies of existing protocols. They formed the NVM Express consortium to create the NVMe and NVMe-oF standards that would fully expose the benefits of non-volatile memory in computing environments. NVMe is designed from the ground up to deliver high bandwidth and low latency storage access for current and future NVM technologies.

In 2020, the NVMe and NVMe-oF ecosystems are large and mature enough for vendors to incorporate these technologies into enterprise-class storage systems. Vendors do so in three primary ways:

- By using NVMe SSDs in their arrays
- NVMe-optimized storage operating systems
- By using the NVMe over Fabrics (NVMe-oF) protocol to transfer data between storage controllers and application hosts

NVMe SSDs. The NVMe protocol supports up to 64,000 queues per PCIe-attached SSD, versus just one queue with SAS and SATA. This new parallelism means each application or CPU thread can have its own queue, eliminating the need for I/O locking.

NVMe enables arrays to deliver more IOPS and throughput. Because the NVMe protocol itself is streamlined, it can deliver double the number of IOPS per CPU core compared to SATA or SAS data transfer protocols, and do so with less latency. (Figure 1)

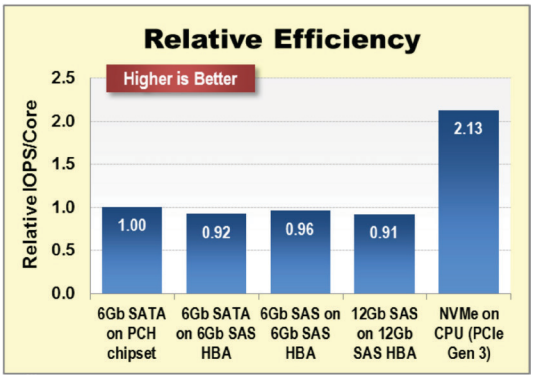

Figure 1 *Source: https://www.nvmexpress.org/wp-content/uploads/NVMe_Overview.pdf*

NVMe-optimized storage operating systems. NVMe allows storage vendors to eliminate overhead associated with I/O locking from the data path. The process of incorporating NVMe SSDs into enterprise storage systems also revealed other opportunities to reduce CPU overhead and latency by streamlining code.

NVMe over Fabrics. NVMe-oF extends the benefits of NVMe across the data center. With NVMe-oF, application servers get more IOPS at lower latency while using fewer CPU cycles for storage-related functions.

AFAs that incorporate all three of these advances enable end-to-end NVMe in the data center, delivering another order-of-magnitude improvement in performance. Replacing legacy HDD-based arrays or first-generation all-flash arrays with an array supporting end-to-end NVMe will deliver more data center performance with fewer servers, in less space, using less power. The economic return on investment can be substantial.

Enterprise storage vendors are at different points in the NVMe journey. More than 40% of the AFAs DCIG researched for the 2020-21 Enterprise All-Flash Array Buyer's Guide incorporate NVMe SSDs into their designs. All of these have taken steps to optimize their storage operating systems. However, less than a fourth of the AFAs support NVMe-oF.

IBM FlashSystem 9200

Source: IBM

The FlashSytem 9200 is based on a 2RU, dual-controller enclosure that supports up to 24 industry standard NVMe drives or IBM FlashCore Modules (FCM). It also supports up to four NVMe storage class memory (SCM) drives per enclosure. These SCM drives are based on Intel Optane or Samsung zSSD in capacities ranging from 375 GB to 1.6 TB each.

IOPS per Controller / System	2.2M / 18M IOPS
Bandwidth per Controller / System	22.5 / 180 GB/s
Latency (microseconds)	< 70

The IBM FlashSystem 9200 provides end-to-end NVMe data transfers through support for both NVMe SSDs and NVMe/FC. IBM claims the FlashSystem 9200 can deliver 18 million IOPS, 180 GB/s bandwidth, and latencies under 70 microseconds. IBM also claims 2.2 million IOPS and 22.5 GB/s bandwidth per rack unit, providing greater performance density than any other product in the Buyer's Guide, and among the highest performance claims in the industry.

Raw Capacity per System	32,000 TB
Media: NVMe / 12 Gb SAS	✓ / ✓
Largest Capacity: SSD / SCM Module	38.4 TB / 1.6 TB
Raw Density	461 TB/RU

IBM FlashCore Modules provide inline-hardware compression, data protection, and specialized flash management features. IBM's 38.4 TB FCM is the largest SSD supported by any AFA in this Buyer's Guide;

which helps explain the exceptional raw storage density (461 TB/RU) the array can provide.

A fully scaled-out FlashSystem 9200 storage system can provide nearly 7.2 PB* of compressed NVMe all-flash storage in just 8 units of rack space. Deduplication may double that to 16 PB. Storage capacity may be further expanded through 12Gb SAS-attached expansion enclosures for a total of 32 PB of raw flash capacity.

IBM FlashSystem 9200 runs the IBM Spectrum Virtualize storage operating system. Its storage virtualization capabilities enable it to bring legacy storage into the FlashSystem management domain, extending its comprehensive enterprise class data services to more than 500 third-party storage systems. This capability also facilitates data migration to the FlashSystem. Longer term, IBM's integrated Easy Tier AI-driven automated tiering can use legacy arrays as a tier of storage.

IBM offers much flexibility in acquisition and deployment of these systems, including storage-as-a-service on-premises, in IBM's cloud, and in the public cloud.

IBM's combination of hybrid cloud capabilities and external storage virtualization of 500+ arrays is unique among products in the Buyer's Guide. This enables the FlashSystem to bring hybrid cloud capabilities to nearly any existing storage array.

IBM Storage Insights predictive analytics and proactive support mechanisms contribute to IBM's six 9's of measured availability for the FlashSytem 9200, and guarantees 100% availability through the IBM FlashWatch program.

IBM offers all-inclusive and "bundles plus a la carte" licensing options on the FlashSystem 9200.

DCIG View of the AFA Marketplace and IBM

Any organization that has yet to adopt an all-flash storage infrastructure for all active workloads is operating at a competitive disadvantage. Deploying a current generation enterprise all-flash array:

- Helps organizations move faster
- Makes existing applications run faster even as data sets grow
- Accelerates application development
- Enables IT departments to say, "Yes" to new workloads and quickly get those new workloads into production
- Increases the performance density of data centers
- Drives down data center operating costs

The IBM FlashSystem 9200 exemplifies these attributes and enables all of these benefits. Beyond the technology, IBM's proactive support, plus its flexible acquisition and deployment options, demonstrates that the company continues to listen to their enterprise customers and then deliver solutions that meet their changing needs.

Note: *The data for "All AFAs" is based on DCIG's independent research into enterprise storage arrays and includes the 99 arrays that satisfy DCIG's criteria for Enterprise All-Flash Arrays.* ■

**Based on Flash Core Module always-on 2:1 compression.*

Acquisition & Deployment Options

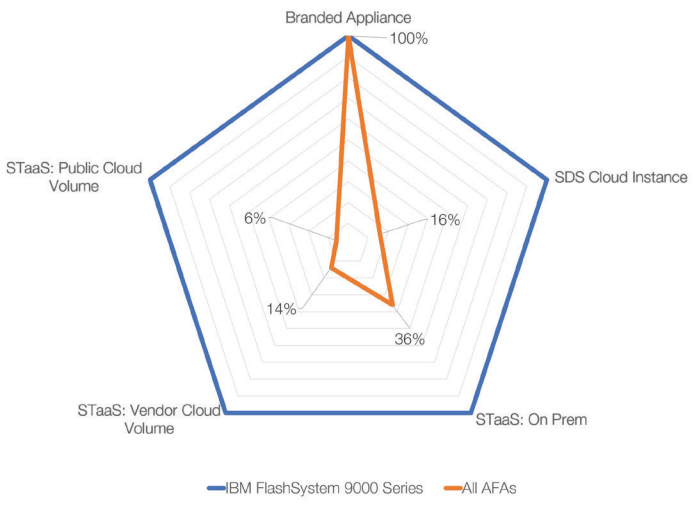

Source: DCIG

Predictive Analytics and Proactive Support

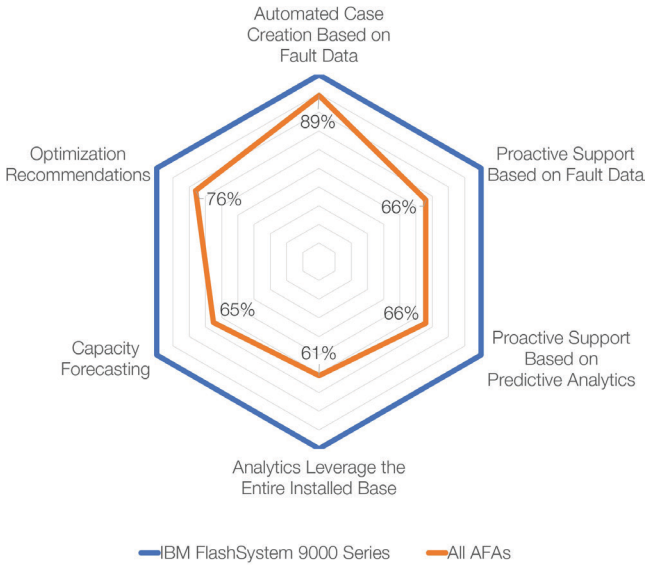

Source: DCIG

About DCIG

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analysis. DCIG analysts provide informed third-party analysis of various cloud, data protection, and data storage technologies. Learn more at **www.dcig.com.**