

# Commentary

August 24, 2016

## IBM Extends Flash Storage Across All Primary Storage

*The first thought that pops into one's mind upon hearing the words "flash storage" is performance. And it is true that flash started in Tier 0 storage. However, flash is seen more and more as the wisest and most appropriate choice for the primary Tier 1 production storage that run businesses. But now flash is encroaching even further; into production storage where not only performance is important, but also scalable capacity, such as for big data analysis. So flash is becoming ubiquitous for production storage. Who would have thought it? But no single flash product is appropriate for all workloads and use cases. That requires vendors to develop an all-flash portfolio. How that all fits together will be illustrated by examining IBM's broad and expansive portfolio of all-flash products.*

### Why Flash Storage Is Becoming Ubiquitous

The three reasons for the success of flash as compared to hard disk drives (HDDs) are: 1) flash storage is better physically, 2) flash solutions offer a wider range of choices in four different key storage environments, and 3) flash storage can easily perform critical functions that HDDs find difficult, at best.

#### *Part I: Flash Is Better Physical Storage*

Hold thoughts about price, performance, and reliability aside for a minute. If a decision has to be made on flash storage versus HDDs purely on a physical basis, why would the choice not be for simpler, solid state flash architecture? Power, rack, stack, and floor space requirements all favor flash storage. A data center benefits from a smaller storage footprint and lower heating and cooling. A smaller footprint, not only for storage but also for

servers means less that can go wrong and, in general, reduces the operational administrative burden. None of these obvious checkmark characteristics favor HDDs.

Now back to price, performance, and reliability. Wear liability was once supposed to be an issue in flash, but no longer especially where IBM has a 7-year (an eternity in storage) wear guarantee for its flash products. Not only that, but natively flash is 3 to 4 times more reliable than spinning disks.

On performance, a particular flash solution can easily be 4 to 10 times faster than competing HDD-based arrays. If necessary, another flash product in a portfolio can be selected for dramatically faster performance if the additional benefit justifies the extra cost.

So that leaves price. But one does not just buy flash chips or HDDs

standalone, but rather as part of a storage solution. And that solution is subject to total cost of ownership (TCO) and return on investment (ROI) analyses. Already mentioned improved environmental factors (such as less power usage) is one small part of flash superiority. Data reduction techniques that are relatively easy to apply to flash storage, but not to HDDs, such as compression, get more usable capacity out of native capacity (i.e., 1 TB of physical flash can actually appear to be several times larger). And these benefits extend beyond the flash storage itself. For example, the number of server cores needed typically decreases substantially for an all-flash solution because flash performs IOPS management more efficiently. Overall, numerous studies have shown that flash storage is more than price competitive for primary storage.

## *Part II: Flash Delivers More Storage Environmental Platform Choices*

But HDD fans don't give up easily. Storage is not just about physical devices, but also the storage platforms on which they are built. What about the tried and true traditional arrays that do the job today in enterprise and mid-range markets? Architectures with HDD still play a useful role, but they are not the wave of the future in most storage environments. All-flash performs well in traditional storage environments, as well as new storage applications. IBM illustrates how this works. IBM defines four principal storage environments. The first is the virtual infrastructure in which traditional storage

systems play a major role and which require heterogeneous enterprise-class data services (such as snapshots). These data services are now available on selected all-flash IBM solutions so anywhere HDD plays flash storage can play, but with added performance and the other benefits. IBM's traditional array products that feature all-flash are IBM Storwize V7000F and IBM Storwize V5030F, as well as the IBM FlashSystem V9000.

A second environment is business critical storage, which is really mainframe systems, a major strength of IBM's. The IBM DS8888 is an all-flash array that can complement or substitute for HDD-based DS8000 series products with the acceleration that more and more applications demand.

A third environment is grid-scale cloud storage. Although Amazon and others have pioneered the use of commodity storage for public clouds, the cloud comes in many varieties, including private and hybrid, and, as we will explore later, flash should be the storage choice for flexibility and capability reasons that spinning disks cannot match. IBM all-flash products, namely, IBM FlashSystem A9000 and IBM FlashSystem A9000R, play in these use cases.

The fourth environment, and definitely one in which storage architectures, such as storage area networks (SANs) are persona non grata, is big data storage. This requires data lakes and oceans, along with high performance file storage for timely analysis. IBM's DeepFlash 150 is designed to serve the needs of these applications.

In other words, all-flash products can be focused, targeted, and optimized for all four environments with the necessary features and functions. HDD options fully support only traditional SAN and NAS environments, with some HDD-based commodity storage meeting other needs. Thus flash delivers more choices for more environments, offering flexibility and variety that HDD cannot match.

### *Part III: Flash Makes Easy What Hard Disks Find Hard*

Let's get physical again. HDDs have been a technological marvel especially as they represent a Rube Goldberg electro-mechanical mélange where moving arms and spinning platters (often spread across multiple devices) are involved for I/O actions. Flash storage uses much simpler direct mapping access approaches. Why does that matter?

The answer is that service level agreements (SLAs) often contain quality of service (QoS) objectives for running multiple applications, especially in multi-tenancy environments. Simply put, flash storage shines in meeting QoS objectives because of its physical nature and how it can be managed; meeting QoS with HDDs can be difficult, at best.

Software-defined storage (SDS) and flash storage are also a compatible couple. Different SDS products can be mixed with different flash storage products to target a wide range of applications and use cases. A case in point is

IBM's use of its Spectrum Scale software in conjunction with its DeepFlash 150 in big data environments.

Further development efforts in both SDS and flash storage would seem to benefit flash storage more than HDD storage as older technology typically cannot be reinvented to overcome a rising and disruptive technology such as flash storage.

### **Match Workloads to the Proper All-Flash Products**

But just because flash storage has many advantages over HDD alternatives does not mean that a single all-flash product can meet all workload and use case requirements for all storage environments and platforms. A portfolio of flash products is necessary. Why? The one size does not fit all rule applies. All workloads do not require the very fastest (and highest priced) flash performance. Different workloads require different features and functions, such as storage data services and price points are often a critical point of concern. Note though that performance varies for each all-flash product, each has a relatively substantial (as mentioned (say 4 to 10X) performance improvement over the HDD alternative with which it should be compared.

The all-flash storage product arena can be categorized into five different areas, each of which requires a flash solution targeted to its specific needs — high end, mid-range, entry level, capacity-driven/scale-out specialty storage, and converged infrastructure systems. Let's

examine how IBM's portfolio of all-flash products cover those bases.

## High End All-Flash Storage Arrays

### *Tier 0 Highest Flash Performance: FlashSystem 900*

When speed is everything, IBM's FlashSystem 900 provides Tier 0 application top acceleration performance. IBM states that the system can do up to 1.1 million 100% random read IOPS and 600,000 100% random write IOPS. This blazingly fast system turbocharges the performance of targeted workloads, like real-time database applications, for which the price of performance is justified by value-add and data services, such as snapshots, are not required.

### *Tier 1 Enterprise-Class Mainframe: DS8888*

The all-flash IBM DS8888 provides the critical performance demanded by certain workloads, such as high level online transaction processing, real-time analytics, and high performance data warehousing. But an all-flash mainframe array has to meet the other unyielding demands of any mainframe environment, such as robustness and reliability. The DS8888 achieves this via deep integration with z Systems environments. For example, it provides six 9's reliability for three or even four sites. High demand online transaction workloads require that level of business continuity.

### *Tier 1 Traditional Enterprise-Class: FlashSystem V9000*

FlashSystem V9000 is an all-flash array that supports data services, such as thin provisioning, snapshots replica-

tion, and high-availability configurations for disaster recovery, that traditional Microsoft, Linux, etc. data centers require for primary storage applications. FlashSystem V9000 scales to over 2 PB to meet large scale demands.

### *Tier 1 Cloud Enterprise-Class: FlashSystem A9000 and FlashSystem A9000R*

Many enterprises are moving to a cloud that has key characteristics, such as on-demand self-service, resource pooling, and rapid elasticity. Such cloud environments impose special demands that general purpose storage arrays do not meet. One is the need to support QoS for multitenancy environments to ensure that application resources are allocated according to service level agreements.

This is a major strength of IBM all-flash arrays. FlashSystem A9000 targets cloud service providers and enterprises with data-intensive requirements. FlashSystem A9000R targets large scale deployments in the very biggest enterprises and hybrid cloud environments.

## Mid-range All-flash Storage Array: Storwize V7000F

Existing Storwize customers who need extra performance as well as the other benefits (say environmental savings) of an all-flash array can continue to use the same capabilities they prefer for data services, such as external virtualization, real-time compression, snapshots, and remote mirroring. That familiarity translates into simplicity and ease of use when turning to the V7000F.

## **Entry-Level All-flash Storage Array: Storwize V5030F**

IBM does not ignore the entry-level enterprise's need for an all-flash array with this smaller brother of the V5000F. The V5030F brings the same features, efficiency, and dependability of its big brother, but with lower deployment cost with new flash drive options.

## **Capacity-Driven Production Storage: DeepFlash 150**

The products already discussed cover all aspects of Tier 1 primary production storage.

But there is more to the production storage world than just traditional views of Tier 1. It turns out that there is a capacity-driven production world as well that does not fit neatly into storage tiering strategies.

This is sometimes thought of as the scale-out NAS world which has typically been the province of specialty storage arrays. This world is populated with production workloads where both enhanced performance and very large-scale (multiple PBs) capacity are necessary. This includes high performance computing (HPC), media and entertainment (M&E), and big data among others.

But even though enhanced performance in these instances is desirable, the cost of all-flash storage was prohibitive. IBM's DeepFlash 150 is a product that overcomes that economic hurdle. Yes, the tradeoff is slightly less performance than a comparable traditional flash system, but the performance is

still 5X that of a comparable HDD array.

IBM's DeepFlash 150 brings general all-flash benefits such as 1/3 the rack space of HDD solutions, but also offers 10X the reliability of a hard disk array. For the PB+ world, fewer device failures is essential and provides a boon for storage operations professionals.

In addition, IBM's DeepFlash 150 is bundled with Spectrum Scale for managing a global file system at PB scale. This illustrates the value of SDS (Spectrum Scale) with flash hardware.

## **Converged Infrastructure: VersaStack**

Converged infrastructure solutions combine a pre-configured bundle of servers, storage, and networking. These integrated infrastructure offerings are becoming more popular because of technical and business efficiencies within data centers, such as through the pre-integration of technology components, IT resource pooling, and IT process automation.

VersaStack is the vehicle that Cisco and IBM use to support converged infrastructure solutions utilizing both companies' solutions. Cisco provides the compute element (UCS servers) and network elements (MDS or Nexus series switches). IBM provides the storage element. The Storwize V7000F or Storwize V5030 (all built with IBM Spectrum Virtualize), satisfy basic needs. For solutions that require high levels of performance for application acceleration at scale at the enterprise level, IBM provides it FlashSystem V9000.

---

## Conclusions

The penetration of all-flash storage into organizations for production storage is proceeding very rapidly even for an IT technology. In fact, all-flash storage might well be said to be well on its way to becoming ubiquitous for production storage.

Why is that the case? With all due respect to HDD technology, which has had a tremendous and well-deserved long run, flash is a superior technology. That superiority starts with the physical product that not only offers native performance and reliability advantages, but environmental/operational benefits through savings in power as well as rack and stacking. With the proper software and other hardware components, all-flash performs well in all storage environments. In addition, flash storage can support tasks, such as QoS, that are critical but are difficult or next to impossible for HDDs.

But all workloads are not created equal and also have different price sensitivities and needs for different feature sets, such as supporting data services. A portfolio of products is required to meet these demands.

IBM provides a comprehensive portfolio that includes at the high end the Tier 0 FlashSystem 900, the

mainframe-focused DS8888, the general enterprise-class FlashSystem V9000, and the enterprise-class cloud-ready FlashSystem A9000 and FlashSystem A9000R. The Storwize V7000F targets the mid-range while the Storwize V5030F does the same for entry-level workloads. DeepFlash 150 zooms in to meet the needs of capacity-driven applications and all-flash as part of the VersaStack bundle serve the needs for a converged infrastructure solution.

All in all, IBM has a far ranging portfolio of all-flash products in both range and breadth.

Flash product superiority coupled with a range of all-flash products to meet a wide set of workload and use case requirements are enabling flash storage to spread like wildfire for production storage.

To learn more about IBM's All Flash solutions, please visit this Website: <http://www.ibm.com/systems/storage/flash/flash-array.html>

---

## David Hill

---

Analyst Name: David Hill  
Topic Area: Storage

Mesabi Group LLC  
26 Country Lane  
Westwood, MA 02090  
[www.mesabigroup.com](http://www.mesabigroup.com)

*This document was developed with IBM funding. Although the document may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.*

Phone: (781) 326-0038  
email the author: [davidhill@mesabigroup.com](mailto:davidhill@mesabigroup.com)

---

The information contained in this publication has been obtained from sources Mesabi Group LLC believes to be reliable, but is not warranted by Mesabi Group LLC. Commentary opinions reflect the analyst's judgment at the time and are subject to change without notice. Unless otherwise noted, the entire contents of this publication are copyrighted by Mesabi Group LLC, and may not be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written consent by Mesabi Group LLC