

REPORT REPRINT

IBM updates storage, refracts a spectrum of systems into a single-platform brand

MARCH 2 2020

By Tim Stammers

The company has updated much of its primary storage systems lineup and has rebranded its Storwize systems as FlashSystem devices. That rebranding highlights the existence of a wide-ranging family of IBM storage systems, all based on the same Spectrum Virtualize software platform.

THIS REPORT, LICENSED TO IBM, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



Introduction

IBM has updated much of its primary, block-level storage lineup, and in the process has also rebranded its Storwize low-end to midrange systems as part of its existing FlashSystem family. The Storwize lineup included hybrid disk-and-flash systems as well as all-flash arrays, all of which have now been given the FlashSystem name previously applied only to IBM's high-end all-flash storage systems. The updates include three FlashSystem systems that were previously at the low end of the Storwize lineup alongside new top-end FlashSystem systems, and new features for all of the arrays in the newly enlarged FlashSystem family. Those features include the use of QLC flash in IBM's custom NVMe drives or modules; options to use Intel Optane and Samsung's Z-NAND NVMe storage-class memory drives as an additional storage tier, and the addition of three-site replication to IBM's Spectrum Virtualize software. IBM says that its updated top-end FlashSystem now delivers 18m IOPS, and provides latencies as low as 70 microseconds, while other FlashSystem arrays can pack 4PB of data in just 2U rack space with data reduction. The updated and new systems will begin shipping in early March.

451 TAKE

These updates and new systems introduce competitive features that are either unique or very unusual among major suppliers' primary storage offerings. Aside from the new features, absorbing the Storwize systems into the FlashSystem family makes so much sense that we wonder why it wasn't done before. Now, the new FlashSystem lineup is much wider and includes a line of systems that is powered by the Spectrum Virtualize software and ranges from low end to high end, and out into public clouds. These systems existed before, but the rebranding drives the point home about the common software platform, and it's an important point. It means that data can be replicated, migrated and tiered across all of those systems, which are all managed via the same interface and APIs. Spectrum Virtualize can also be hosted in public clouds, allowing any on-premises FlashSystem device to replicate data directly to the clouds for disaster recovery or application migration. The software can also do all of the above for a very wide range of other IBM and third-party storage systems. Taken as a whole, it's an unusual proposition, and it exists because of the strength of the Spectrum Virtualize software.

Spectrum Virtualize as uber controller

Spectrum Virtualize originated as the software that powered IBM's SAN Volume Controller (SVC) appliance, which was launched in 2003. IBM later launched Spectrum Virtualize as a software-only alternative to the appliance. The two products have a strong reputation and, according to IBM, have scored thousands of deployments.

When the SVC appliance was launched, it was a category-busting product that helped introduce 'storage virtualization' into the industry lexicon. There are still few equivalents to it, or to the Spectrum Virtualize software. In either form – as SVC appliance or as Spectrum Virtualize software – the software runs on one or more clustered servers and can virtualize over 500 types of IBM and third-party storage system, pooling and tiering their capacity behind a single control system and management interface. That can heavily streamline storage administration, as well as simplify the migration of data between storage systems. IBM also claims that Spectrum Virtualize can boost performance, and in the past has cited SPC benchmark results to back this claim (although the usual caveat applies: benchmarks only attempt to indicate real-world performance.)

A single platform from low end to high end, and public clouds

IBM's latest product launches are not about the use of Spectrum Virtualize as a stand-alone virtualization platform, but about the use of Spectrum Virtualize as an integral part of IBM storage systems. When IBM bought Storwize for its data compression technology in 2010 and Texas Memory Systems (TMS) for its all-flash arrays in 2012, the company faced the problem that neither of those companies' products offered the enterprise data services that are table stakes for datacenter storage. The solution was simple: IBM integrated Storwize technology into Spectrum Virtualize, creating the storage systems that were branded as Storwize. The ex-TMS products were combined with Spectrum Virtualize software and branded as FlashSystem devices.

Now that Storwize systems have been rebranded as FlashSystem devices, the FlashSystem brand ranges from entry-level hybrid disk-and-flash storage systems that cost less than \$20,000 to high-end all-flash systems carrying seven-figure list prices. IBM also offers versions of Spectrum Virtualize for hosting in Amazon and IBM public clouds.

The FlashSystem lineup includes other products not based on Spectrum Virtualize. These are the all-flash A9000 and A900R, which are based on Spectrum Accelerate (formerly XIV) software, and the all-flash tier-zero FlashSystem 900. However, in IBM's own words, these products have been 'de-emphasized.' IBM's block-level primary storage also includes its DS8900F storage system, which is primarily aimed at mainframe environments and is not branded as a FlashSystem.

FlashSystem updates

In addition to absorbing Storwize into the FlashSystem family, IBM has introduced new systems at the middle and top end of what it is now the FlashSystem range. The new boxes boost performance by using faster controllers and updated NVMe flash drives, and by offering new options to use Intel Optane or Samsung Z-NAND NVMe storage-class memory drives as a storage tier. All the systems were already powered by NVMe all-flash drives and supported NVMe externally via NVMe-oF.

Some of the previous Storwize arrays and all of the FlashSystem arrays were powered by IBM-made custom NVMe flash drives, or in Big Blue's nomenclature, FlashCore Modules (FCMs.) These FCMs are offered in a range of capacities, and they have been reworked so that the largest FCM now packs 38TB of usable capacity into the 2.5" industry-standard form factor. That is an extremely large capacity for a 2.5" drive, and IBM claims it is the biggest on the market. With data compression and dedupe, this allows a single FlashSystem array fitted with 38TB FCMs to provide 4PB of effective capacity in just 2U rack space, assuming a not-unreasonable 4.5:1 data reduction ratio.

The compression is completed by FPGAs mounted within the FCMs to boost overall system performance by offloading the system main controllers. IBM says that per GB of raw capacity, its FCMs carry comparable list prices to third-party NVMe flash drives, even though the latter do not offer onboard compression. The new FCMs are powered by the latest QLC variant of NAND flash, which presumably is a major reason for their very high maximum capacity. This use of QLC puts IBM at the cutting edge, as 451 Research knows of only two other vendors that are doing so in primary storage systems (see Competition section below.)

QLC stores more data per chip and is, therefore, cheaper than the TLC flash that currently dominates datacenter storage, but it is also slower and suffers a shorter write-life or endurance. However, IBM says its new FCMs carry the same seven-year warranties as their predecessors because of what it says is its efficient use of wear-leveling, dynamic allocation, sparing and other flash management techniques.

REPORT REPRINT

Big Blue also claims the FlashSystem lineup tops out at a whopping 18m IOPS (4kB blocks, 100% read) and 180GB/s throughput per four-way cluster, with latencies as low as 70 microseconds. As always, 451 advises that simplistic marketing claims only provide approximate indications of real-world performance.

Although IBM is not the first major vendor to offer Intel Optane NVMe drives as an option for storage systems, it is the first to offer Samsung's Z-NAND drives in any role within a primary storage system. Z-NAND is an SLC variant of flash that sits between Optane and TLC flash in terms of performance and cost. IBM says production configurations of FlashSystems might include as much as 5TB of Z-NAND, allowing multiple database instances to be held in persistent storage that is faster than regular flash but significantly cheaper than Optane. The Optane and Z-NAND drives are NVMe drives made by Intel and Samsung, respectively, rather than FCMs powered by those two memories.

Spectrum Virtualize and SVC updates

The latest updates to the Spectrum Virtualize software include new support for three-way replication and Red Hat Ansible. As DR and resiliency requirements rise, three-site replication is expanding out from its traditional base of banks and other deep-pocketed organizations into wider usage. The first version of Spectrum Virtualize for use in public IaaS clouds shipped in 2018, and currently two clouds are supported: the IBM Cloud and AWS.

Spectrum Virtualize running in AWS has also gained the ability to compress and dedupe the data it manages in EBS, and the number of clustered EC2 instances of Spectrum Virtualize has been doubled from two to four. That brings the AWS version of Spectrum Virtualize level with the IBM Cloud version, which had all those features before – except the dedupe, which is also new to the IBM Cloud version.

The AWS and IBM Cloud instances of Spectrum Virtualize are now supported by IBM's analytics and 'AI-infused' cloud-based Storage Insights Pro management service, as are EMC Unity XT, NetApp FAS and AFF, and Hitachi Vantara VSP G-series storage systems (Spectrum Virtualize is not required to provide that third-party support.)

Spectrum Virtualize implemented as the SVC appliance has also been updated; IBM has begun offering two versions of the appliance with different hardware configurations, replacing the single version previously offered.

Competition

IBM's biggest competitors in primary storage are Dell EMC, Hitachi Vantara, HPE, NetApp and Pure Storage. Dell EMC is the biggest of this pack. While IBM is highlighting its use of a single software platform to underpin almost all of its FlashSystem lineup, Dell EMC is moving toward a similar goal; it is in the process of rationalizing and consolidating its two EMC-originated and Dell-originated midrange to high-end primary storage portfolios.

There are only two other major suppliers that match IBM's ability to allow on-premises storage systems to replicate or migrate directly to instances of their operating software hosted in public clouds. Those two companies are NetApp and Pure Storage, via NetApp's Cloud Volumes ONTAP, and Pure's Cloud Block Store. However, neither the NetApp nor the Pure Storage system lineups stretch as far down into the low end as the FlashSystem range does, nor have they been enlarged with the former Storwize products. At the other end of the range, we don't know whether the Pure or NetApp systems can match the performance that IBM is claiming for a four-way clustered FlashSystem array.

REPORT REPRINT

As stated above, 451 is aware of two other suppliers of datacenter primary storage powered by QLC flash. One is Pure Storage, which unlike IBM is not using QLC as a direct replacement for TLC flash, but is instead offering QLC-powered storage systems as slower but less expensive alternatives to its existing TLC-powered systems. The other is startup VAST Data, which is mixing QLC flash with Optane in an unusual scale-out storage system.

The ability of Spectrum Virtualize to virtualize third-party storage is unusual, but not unique. NetApp and Hitachi Vantara both offer the ability to virtualize and apply data services to third-party storage systems by using the operating systems that power their flagship products – NetApp’s ONTAP and Hitachi’s SVOS. However, unlike IBM, neither NetApp nor Hitachi offer stand-alone storage virtualization products like Spectrum Virtualize or the SVC appliance.

SWOT Analysis

STRENGTHS

IBM offers a range of primary storage systems that are based on the same software, and span from entry level to high end.

WEAKNESSES

IBM’s portfolio contains no unified storage systems offering both file and block-level access to data.

OPPORTUNITIES

The ability to host Spectrum Virtualize in public clouds is of value to customers implementing or planning to implement hybrid clouds.

THREATS

Spending on on-premises storage systems is shrinking, and is set to continue doing so for a number of reasons that include the growing usage of public clouds.