IBM FLASHSYSTEM – TECHNICAL REVIEW AND EVALUATION 2020

CATEGORY: TECHNICAL EVALUATION

Release 1.1
22 FEBRUARY 2021

IBM FlashSystem Review 2020 - Second Edition

Published by Brookend Limited.

Document reference number BRKSW0131-1.

Second edition (1.1) updated to reflect upgrades to the IBM portfolio.

CONTENTS

Introduction

Shared storage in the form of SANs or storage area networks has been part of the enterprise landscape for almost 30 years. Modern storage systems offer great value for money, flexibility, resiliency, and efficiency. The market for these products is extremely competitive and as a result, vendors continue to evolve their solutions to keep up. IBM recently aligned the FlashSystem product range onto a single software platform and codebase, enabling customers to scale from small business all the way to large enterprise using a single solution.

IBM commissioned Architecting IT to review and evaluate the capabilities of FlashSystem. This document combines the results of hands-on testing and research, which is also available online as a series of three blog posts.

In addition, IBM recorded a Storage Unpacked podcast that highlights the latest FlashSystem enhancements and provides more details on how customers are using the FlashSystem technology.

This document is divided into three sections. The first covers FlashSystem hardware and in particular the use of IBM proprietary NAND flash storage known as FlashCore modules. The second section of the document looks at FlashSystem software. IBM has standardised on storage software derived from the SAN Volume Controller, now called Spectrum Virtualise (PDF). We look at how this standardisation enables customers to scale from entry-level to high-end with the same platform, interface and management tools.

The third section of this document looks at the ease of use of FlashSystem. In particular, this covers the process of installation, management and operations. This section also covers purchasing models and IBM's adoption of transparent pricing.

Hands-On

To gain a better understanding of FlashSystem, IBM provided Architecting IT with an evaluation system over a period of two weeks. The evaluation FlashSystem 5030 was installed in the Architecting IT Lab and used as storage for virtual workloads on VMware vSphere. This connectivity was established with both Fibre Channel and Ethernet protocols.

The aim of the hands-on testing was not to provide a point-by-point breakdown of the platform, but to validate claims made by IBM on the simplicity and usability of the FlashSystem platform.

Hardware

Modern storage arrays have mostly commoditised onto standard server-based form factors from the monolithic (single rack) solutions of the late 1990s onwards. IBM FlashSystem is no exception, with an evolutionary transition towards standardised hardware and media. Here we look at the FlashSystem hardware options and unique features including FlashCore Modules.

The FlashSystem Family

The FlashSystem family is now divided into three main groups:

- **Entry** – Models 5000 and 5200
- **Midrange** – 7200 and 9200
- **High-end** – 9000 and 9200R

Hardware is consistent across all models and built on a dual-controller/node (or canister) architecture that scales from a 2U server form factor to a single 19" rack solution. The following table *(Table 1 - FlashSystem Models*) summarises aspects of the hardware configuration.

Table 1 - FlashSystem Models

Model	Form Factor	Processor	Memory	Connectivity	Drives
5015 (H)	2U, Dual Canister	2x 2.2Ghz, 2 core Intel Broadwell	32-64GB	1 GbE (iSCSI) & 12Gb SAS, optional 4x 16Gb FC, 10/25GbE iSCSI, 12Gb SAS	12x 3.5" or 24x 2.5" SAS, max 392 with expansion
5035 (H)	2U, Dual Canister	2x 2.2Ghz, 6 core Intel Broadwell	32-64GB	10 GbE (iSCSI) & 12Gb SAS, optional 4x 16Gb FC, 10/25GbE iSCSI, 12Gb SAS	12x 3.5" or 24x 2.5" SAS, max 504 with expansion
5200 (H)	1U, Dual Canister	2x 2.3Ghz 9 core Intel Skylake	64-512GB	10GbE (iSCSI), optional 16/32Gb FC, 25GbE iSCSI/iSER	24x 2.5" NVMe SSD, optional 4x PM drives, max 748
7200 (H)	2U, Dual canister	4x 8-core 2.1GHz Intel Cascade Lake Processors with hardware compression assist.	128GB – 1.5TB	10GbE iSCSI onboard, with optional 16/32Gb FC and FC-NVMe, 25GbE iSCSI/iWARP/NVMe-oF	12x 2.5" NVMe SSDs, 12x 3.5" SAS HDD or 24x 2.5" SAS HDD/SSD, plus expansion
9100	2U, dual canister	Up to 4x 14-core 2.2Ghz Skylake Processors with compression assist	128Gb – 1.5TB	10GbE onboard (iSCSI), with 25GbE and 16Gb FC optional.	24x 2.5" NVMe Drives
9200	2U, dual canister	4x 16-core 2.3Ghz Cascade Lake Processors with compression assist	128GB – 1.5TB	10GbE iSCSI onboard, with optional 16/32Gb FC and FC-NVMe, 25GbE iSCSI/iWARP/NVMe-oF	24x 2.5" NVMe or FCM drives (760 per system), 12Gb SAS expansion

The 5015 replaces the previous 5010, while the 5035 replaces the 5030, both with slight performance upgrades. The 5200 is a replacement for the 5100 in a 1U chassis.

The current models replace the previously branded Storwize products and despite the name, are offered as hybrid and all-flash solutions as indicated with the "H" suffix. The 5015 and 5035 models have onboard iSCSI (either 1/10GbE) and the option to add SAS expansion shelves. Host connectivity can also be extended with 16Gb Fibre Channel, 10/25GbE iSCSI or SAS.

The 5200 offers connectivity with 10GbE onboards and expansion cards for 16/32Gb Fibre Channel, 25GbE iSCSI or iSER.

Hard drives are offered across five categories:

- 2.5" Persistent Memory (SCM) drives
- 2.5" Tier 1 Flash – 1.9TB – 30.72TB capacity
- 2.5" High performance HDD (10,000RPM) – 900GB – 2.4TB capacity
- 2.5" Capacity HDD (7,200RPM) – 2TB only
- 3.5" HDD (7,200RPM) – 4TB – 16TB capacity

The 5200 models upwards offer IBM FlashCore Modules (FCM), a custom solid-state drive that uses IBM MicroLatency technology from the acquisition of Texas Memory Systems (more on this in a moment).

Figure 1 - FlashSystem 5100 rear view

Where the 5000 and 5200 series systems use dual controllers that are mounted side-by-side, the 7200 models use a top/bottom design that provides for greater network port expansion.

The FlashSystem 9200R combines two, three or four FlashSystem 9200 systems in a single rack acting as a cluster, with pre-cabled configuration and Fibre Channel networking.
The features of each FlashSystem option are summarised in the following table (*Table 2 - FlashSystem Features by Model).*

The 5200 is possibly unique in the market to be the only 1U storage platform supporting two controllers (canisters in IBM terms).

Table 2 - FlashSystem Features by Model

	FlashSystem 5015	FlashSystem 5035	FlashSystem 5200	FlashSystem 7200	FlashSystem 9200
IBM Spectrum Virtualise Software	⊗	⊗	⊗	⊗	⊗
IBM Storage Insights	⊗	⊗	⊗	⊗	⊗
VMware and Red Hat OpenShift Container Integration	⊗	⊗	⊗	⊗	⊗
3-Site Replication	⊗	⊗	⊗	⊗	⊗
Local & Remote Replication	⊗	⊗	⊗	⊗	⊗
IBM Easy Tier	⊗	⊗	⊗	⊗	⊗
Transparent Data Migration	⊗	⊗	⊗	⊗	⊗
Data Reduction Pools		⊗	⊗	⊗	⊗
Scale-out Clustering		⊗	⊗	⊗	⊗
HyperSwap High Availability		⊗	⊗	⊗	⊗
Encryption		⊗	⊗	⊗	⊗
NVMe Flash and FC-NVMe host connectivity			⊗	⊗	⊗
High-performance compression and encryption in FCM			⊗	⊗	⊗
External Storage Virtualisation			⊗	⊗	⊗
Storage-Class Memory			⊗	⊗	⊗

Scalability

IBM offers a standard set of features and software functionality across the entire FlashSystem range, through the use of a single storage operating system (which will be discussed in more detail later in this document).

The range of offerings from 5015 upwards provides an entry-point for SMB/SME customers, right through to the largest enterprise requirements. Customers can choose to cluster FlashSystem enclosures together to create increased levels of availability, which is offered as a pre-packaged solution in the form of the 9200R.

Having a single product range based on a consistent hardware and software architecture is important because:

- Customers can choose their entry point and scale up or down on the basis of their requirements.
- Customers can pick the right model and solution to expand existing infrastructure, either increasing capacity (more disk shelves) or throughput (more controllers). FlashSystem aids that strategy by providing clustering and array-based replication.
- A single look and feel provides consistent operations, management and reporting. Scripting and automation that works on one platform will be guaranteed to work on another.
- A consistent hardware model provides more predictability in terms of performance and availability.
- A single platform type reduces the amount of training and skills development needed for IT teams.

FlashCore Modules

One exciting innovation that has developed from IBM's TMS acquisition in 2012 is the FlashCore Module (FCM). Most storage vendors use off-the-shelf commodity SSDs or have built systems that are entirely bespoke in design. FCM bridges the two options by offering a custom-designed solution which fits into a standard FlashSystem appliance using the 2.5" form factor and NVMe interface.

FlashCore Modules have gone through a number of evolutions since the technology was acquired from TMS. Today, the FCM 2 hardware uses custom FPGAs, Micron QLC NAND (96, layer, previously TLC in the FCM 1) and Everspin STT-MRAM to eliminate the need for super-capacitors. This allows FCM to support capacities up to 38.4TB per device (uncompressed).

SSD vendors are notoriously secretive about the algorithms used to manage the NAND within their devices, so using a custom design allows IBM to improve the capabilities of FCM, which includes built-in encryption and compression. The encryption feature ensures no data can be read from a FlashCore Module removed from a FlashSystem appliance. Encryption keys are managed internally by the system.

Data compression within FlashCore modules produces a typical 2:1 reduction, with no performance impact. Where data reduction results in higher compression ratios, the following maximum effective capacities are achievable:

- 4.8TB FCM – 21.99TB maximum effective capacity
- 9.6TB FCM – 21.99TB maximum effective capacity
- 19.2TB FCM – 43.98TB maximum effective capacity
- 38.4TB FCM – 87.96TB maximum effective capacity

The maximum effective capacity ratio is determined by the size of metadata available on the drive used to store compressed data. At a typical 2:1 ratio, a single 24-drive 2U system supports an effective capacity (before RAID overhead) of over 1.8PB and maximum effective capacity of over 2PB.

FlashCore Modules provide IBM with a significant differentiation over many competitors using traditional NVMe SSDs, including improved performance while using cheaper, higher-capacity NAND. FCM 2 delivers 2 DWPD (device writes per day), around double the equivalent TLC-based drive. Features including heat segregation (placing active data on healthy flash cells) are an enabler to meeting the DPWD capability.

*(**Side note:** FlashSystem does have restrictions on combining FCM and non-FCM drives in the same chassis, including the layout of RAID groups, so it's worth checking the configuration documentation for the full details).*

Software

This section looks at the FlashSystem from the perspective of the software used to deliver storage services.

Storage Operating System

Like many vendors, IBM uses a "storage operating system" to manage the operations of FlashSystem and in the latest models, that software is based on IBM's SAN Volume Controller or SVC. From 2015 onwards, IBM rebranded SVC as Spectrum Virtualise to align with other solutions in the portfolio although we will use the names interchangeably here (as does IBM in documentation).

The initial design of SVC was to act as a storage virtualisation gateway between Fibre Channel storage and hosts. The virtualisation or abstraction functions allow dynamic reconfiguration of host-presented storage without impacting availability to hosts. For large enterprises, this capability enables physical storage to be added, removed or otherwise reconfigured without affecting host operations. In many instances SVC was used as a tool to restructure or reconfigure existing physical resources as part of a rationalisation or upgrade program.

The first implementation of SVC acted purely as a gateway, however in 2010 SVC was updated with the capability to use local disks and released as the Storwize V7000 platform. In 2018, IBM released the FlashSystem 9100 as the first to use SVC software with FlashCore modules (other FlashSystem solutions at the time used XIV code). In February 2020, IBM consolidated the Storwize and XIV-based FlashSystem products into a single family based on SVC. This now comprises the 5000, 7000 and 9000 models available today.

Standardisation

The storage industry has long debated the issue of standardisation on a single platform storage operating system. Many enterprise vendors have built up portfolios of products from multiple acquisitions or development streams. Each product has a specific niche or unique design aspect that fits part of the enterprise market.

While this approach has many benefits (including faster access to customers or a potential market), the use of multiple architectures can result in both confusion and additional management for end users. Standardisation on a single storage operating system provides the following benefits:

- **System administrators are offered a consistent look and feel** that spans the GUI, API and CLI modes of operation. Processes, procedures and scripts developed for one product will work seamlessly on others in the family. Administrators don't need to learn how to work with multiple solutions.
- **System design and layout is standardised.** The design of storage platforms (where RAID and data protection remains important) is achieved in a consistent manner, whether installing from the entry-level systems to enterprise-class products. This

means customers don't have to redesign for branch or edge cases where the deployed solution may be smaller.
- **Standardisation offers predictability.** As environments scale, customers want predictable performance and to know that migration to a larger piece of infrastructure will not introduce a whole new learning curve on platform management.

IBM has made the move to implement a single storage operating system that supports entry-level SMB through midrange and enterprise customers.

As we discussed earlier, the hardware options allow for a range of deployment scenarios. Some features of SVC are only available where the hardware is capable of implementing it (for example encryption, or data reduction pools). In all cases, SVC offers a consistent look and feel across the product range.

Storage Virtualisation

Before we dig into the details of Spectrum Virtualise, it's worth remembering where the software came from. In the early 2000s, shared storage gained massive popularity, with Fibre Channel SANs leading the way in the enterprise. iSCSI was more successful in smaller organisations where the cost of FC couldn't be justified. One big challenge for many customers was the ability to fully exploit assets in their data centre.

Without rigorous management processes, it was easy for IT organisations to end up with a fragmented expanse of storage hardware that ultimately resulted in poor utilisation and uneven performance. SVC provided a virtualisation layer to abstract the details of the underlying hardware, either to reorganise resources or to use SVC as a method of avoiding forklift upgrades.

Many of the terms and definitions within SVC still reflect the ability to use external storage, even if that feature isn't utilised (external virtualisation, for example, is available on the 5200 upwards). Customers still have the option to use external virtualisation as a consolidation tool. This feature can be particularly useful when moving from a mix of current hardware that is due for decommissioning.

SVC software runs on each node within a FlashSystem array enclosure. Changes to the configuration of a system are synchronised in memory between nodes as commands are executed through the GUI or CLI. Multiple enclosures may be joined together to form a cluster, using either IBM HyperSwap or standard FC clustering. We'll return to this in a moment.

Feature Set

Storage features are implemented within Spectrum Virtualise, with each release introducing new or amended functionality. In the first section of this document, we showed the features of FlashSystem in a summary table. Here we look at these features in more detail.

Storage Insights

Storage Insights is IBM's solution for cloud-based monitoring and management. The SaaS platform collects data from IBM FlashSystem deployments to provide a single rolled-up management view across all installations. Customers register for Storage Insights and permit IBM to upload telemetry and metadata into the Insights platform.

Through the use of AI, IBM can provide customers pro-active management of their estate including the collective wisdom across the entire customer base. You can read more about the benefits of centralised telemetry in this post.

VMware and Red Hat OpenShift Integration

FlashSystem provides integration into common application and virtual infrastructure frameworks, including, for example, VMware vSphere through VAAI and vVOLs technology. Some configuration work is needed to exploit these integration points, including the deployment of local agents or proxy software.

3-Site Replication

Data replication is typically achieved between two locations, either synchronously within short distances or asynchronously over a wider area. A common solution for enhanced replication protection involves using 3-site replication where data is replicated to two remote sites that have sufficient metadata to determine the differences between the two sites if the primary array is lost. This makes it easy to re-establish a consistent replication relationship between the surviving FlashSystem arrays with minimal data replication.

Local and Remote Replication

FlashSystem supports local replication through FlashCopy, which uses a copy-on-write (CoW) mechanism to protect replicated data. FlashCopy images can be either snapshots, where only metadata is duplicated until data is written, or clones, where data is physically copied as a background process. Clones provide the benefit of isolated performance and protection against data loss in RAID group failures (but not system loss).

Metro Mirror provides synchronous replication between FlashSystem arrays and is generally used for short-distance protection, such as two data centres up to 300km apart. Global Mirror provides asynchronous replication over distance, which is effectively unlimited but doesn't offer full RPO=0 (recovery point objective) resiliency. Both solutions use either Fibre Channel or IP for inter-system connectivity and can be configured in a range of different scenarios.

Easy Tier

Easy Tier delivers data placement across multiple tiers of physical storage. Hybrid systems initially used a limited amount of flash storage for faster workloads and hard drives for capacity. Easy Tier provides the capability to balance workloads across a mix of all storage types within FlashSystem, either NVMe SSDs (including SCM), HDDs or SAS drives. Extent size is determined at storage pool creation time and can be between 16-8192MiB.

Transparent Data Migration

TDM is a feature of Spectrum Virtualise that exploits the storage virtualisation functionality to enable the import of data from external storage and within internal storage capacity. Internal disks are grouped into MDisks, a RAID set of similar disk characteristics. MDisks and external storage can then be grouped into a storage pool, either for presentation to hosts or as part of a migration process. The ability to mix external and internal storage is a powerful tool in managing the re-organisation of storage resources and delivering outage-free migrations.

Data Reduction Pools

A Data Reduction Pool or DRP is a storage pool against which specific data reduction technologies such as compression or deduplication may be applied. Data Reduction Pools can be built from external storage LUNs, which provides a way to implement data reduction techniques on storage systems that don't offer those features natively.

HyperSwap

HyperSwap is the capability to use multiple FlashSystem arrays as a single logical image using ALUA, either locally within the same data centre or across a campus location. Both systems in a HyperSwap high availability (HA) pair store an independent copy of data, updating each copy synchronously before confirming I/O success to a connected host.

Ease of Use

Storage hardware has become a commodity, and the significant areas of differentiation for vendors are with software, features and the consumption business model. In the previous sections, we looked at SVC (aka Spectrum Virtualise) and how IBM has consolidated the FlashSystem family onto this single storage operating system. FlashSystem now offers a wide range of mature features, but how easy is the system to operate and manage?

IBM provided Architecting IT with an evaluation FlashSystem 5030, configured as follows:

- Dual Controllers
- 16Gb Fibre Channel
- 10Gb Ethernet
- Four 2TB 12Gb SAS SSDs
- Seventeen 2.4TB 10K HDDs

The testing aimed to validate the ease-of-use claims for FlashSystem and SVC. This process doesn't include specific performance testing but does look at performance graphs when validating features within the GUI. Over the course of two weeks, we configured and used the FlashSystem with Fibre Channel and iSCSI connectivity against both dedicated servers and virtual server (VMware) environments.

Installation

The initial installation and configuration process itself is straightforward, although slightly more challenging when only remote hands support is available (as under the current COVID-19 regulations). The FS5030 requires direct connection using a laptop or local computer to perform the initial basic steps of installation. Direct connectivity is needed, as the array acts as a local DHCP server and assigns a temporary IP address to the connecting device.

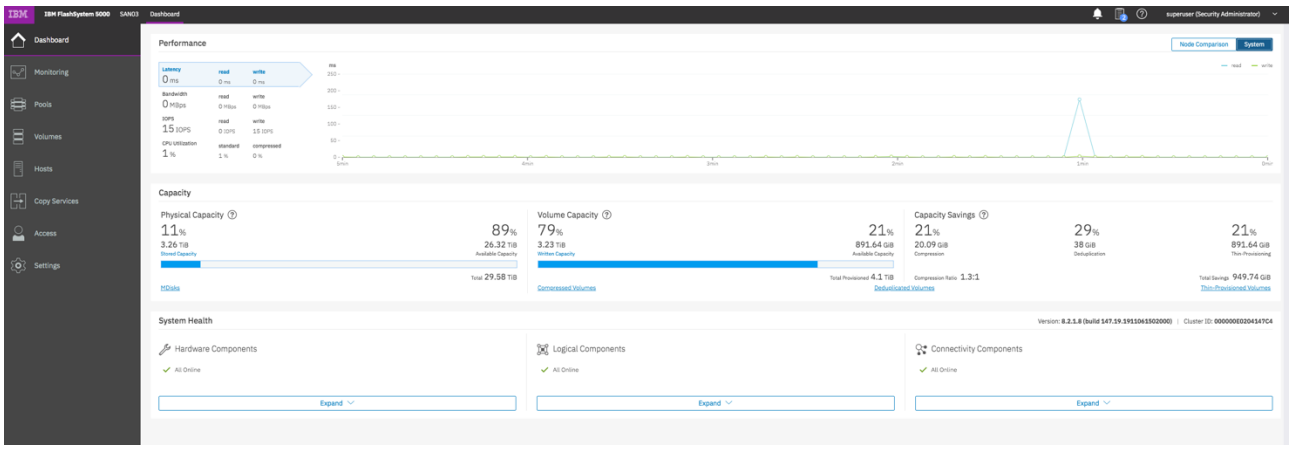

From this point onwards, the remainder of the installation proceeds from a web GUI accessible across the network. Here are some of the key operational benefits of FlashSystem.

Visualisation

One of the most useful features of any storage GUI is a visual representation of the appliance itself. *Figure 2 - FlashSystem Visualisation* shows the FS5030 as a schematic, both front and rear. These views make it easy to identify occupied and vacant drive slots, active Fibre Channel ports and active network ports. In addition, each item is clickable, displaying further details on each component.

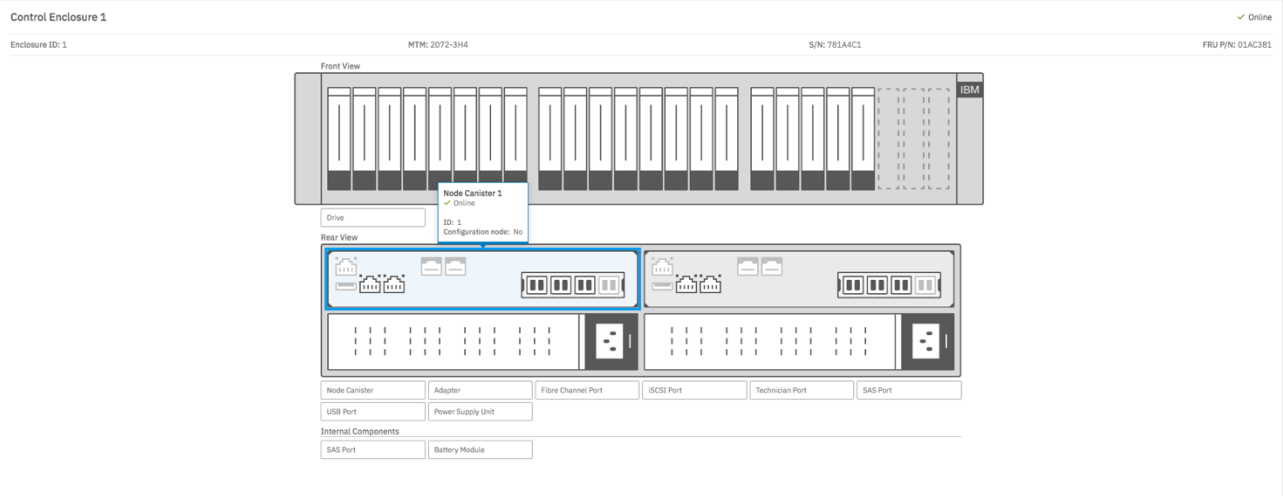

Figure 2 - FlashSystem Visualisation

Although many platforms offer visual representations of storage hardware, the benefits shouldn't be underestimated. Visual cues make it easier to determine the right ports to plug in networking or drives and to see whether those components are active or not.

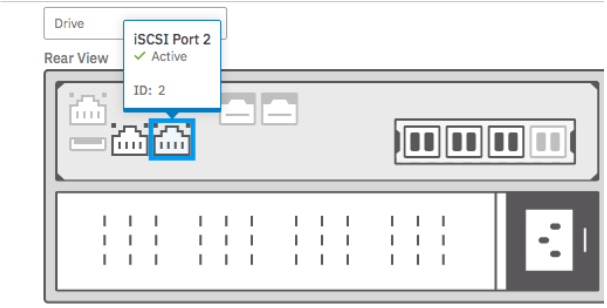

Figure 3 - Annotation on a single component

Command Line Interface

Graphical Interfaces are great, but power users want access to platforms via API or CLI. A lot of storage administration is repetitive and potentially error-prone, so automating with scripts reduces the chances of error or other typos. FlashSystem offers a comprehensive GUI with dozens of commands covering all aspects of the platform. You can find a list here, as part of the FlashSystem documentation.

The implementation of a CLI is vital to how systems operate in a multi-tenant environment. FlashSystem GUI commands are essentially CLI commands, enabling one single point of serialisation for all system operations. This design means that users can be administering through both the GUI and CLI at the same time, without overwriting or undoing each other's work.

API

Storage appliances have offered GUIs and CLIs since inception. Mature enterprise environments have been automating storage for almost 20 years with CLI integration. As data centres transform to be more automated, the CLI isn't always the best ongoing solution, as syntax and data formats can change over time. Many scripting processes rely on "scraping" the results from the CLI, which can introduce errors as software updates are applied. APIs provide a much better interface for automated management, as they are more efficient and secure.

FlashSystem provides API access with almost sixty separate function calls, which are internally translated into CLI calls by the FlashSystem API server. IBM offers online examples in Perl and with CURL. A PowerShell plug-in is also available through IBM Spectrum Connect.

Spectrum Connect

Spectrum Connect software also provides connectivity into a range of other platforms for automation and virtualisation. This includes:

- IBM Storage Enabler for Kubernetes (automation for dynamic storage provisioning)
- VMware vSphere VASA (storage awareness APIs)
- VMware vSphere Web Client
- VMware vRealize Orchestrator
- VMware vRealize Operations Manager
- Microsoft PowerShell Plugin

Spectrum Connect functionality is delivered through a separate server platform that offers scalability and centralisation for customers with many FlashSystem installations.

Spectrum Control & Insights

For customers with many FlashSystem appliances, IBM offers Spectrum Control and Spectrum Insights. Spectrum Control provides an on-premises dashboard that centralises the visualisation of many appliances into a single screen. These views encompass storage systems and storage networking, as well as application-based information from common hypervisors and operating systems.

Modern storage solutions demand ease-of-use and features to ensure 100% availability (or as close to it as possible). The "last mile" in achieving high availability comes from the shared knowledge of the thousands of deployed solutions in the field.

Spectrum Insights is IBM's analytics-based monitoring solution that runs as a SaaS platform in the cloud. Insights offers alerts and problem resolution across FlashSystem infrastructure and applications.

Note: *Spectrum Insights requires a separate chargeable licence.*

Pricing

Our final aspect of usability is in the pricing of storage solutions. For years, storage pricing has been shrouded in secrecy, making it difficult to compare one vendor to another or mix and match different configuration options. The public cloud has made the cost of IT more transparent than ever before. This transition forces on-premises infrastructure vendors to reassess their selling strategies, including the online publication of prices.

IBM now offers customers the ability to configure and build entry-level FlashSystem 5000 series models that include pricing at the component level. The customer can include software into this model and quickly gain an idea of base pricing before discounts.

Transparent pricing is available in the US and Europe and only for FlashSystem 5000 models. It would be great to see IBM extend this facility to other products in the FlashSystem family, so customers can see how the capabilities of each model scales with their associated cost.

The Architect's View™

Shared storage appliances have evolved significantly over the past two decades. The storage monoliths of the last decade have given way to high-density solutions that pack an enormous amount of value into a 2U chassis. Today's enterprise customers want simplicity, reliability, efficiency and value for money.

FlashSystem has evolved into a unified solution that scales from small business requirements to the high-end enterprise. As vendors standardise on the 2U server form factor, features such as FlashCore provide IBM with differentiation in an increasingly competitive market.

The capability of the storage operating system continues to be essential in defining the success of shared storage. Although we shouldn't need to know, implementation details are still important because media and hardware have specific characteristics that are echoed in the capabilities of software. Consistency is a "must have" so it's pleasing to see IBM finally standardising on a single storage O/S and single product range (except for mainframe). SVC is a mature technology that fits the requirements well.

IBM continues to deliver solutions for a highly competitive storage market, with features and functionality that matches the leaders in this field.

References and Further Reading

FlashSystem podcast and blog posts.

- #175 – IBM FlashSystem Deep Dive (podcast)
- IBM FlashSystem Review – Part 1 - Hardware (blog post)
- IBM FlashSystem Review – Part 2 - Software (blog post)
- IBM FlashSystem Review – Part 3 – Ease of Use (blog post)
- IBM FlashSystem 7200 Product Guide (pdf)
- IBM FlashSystem 9100 Product Guide (pdf)
- IBM FlashSystem 9200 Product Guide (pdf)
- Implementing IBM FlashSystem 5010 and FlashSystem 5030 with IBM Spectrum Virtualise V8.3.1 (pdf)
- FlashSystem 5000 – online interactive demonstration (web)
- FlashSystem 7000 – online interactive demonstration (web)
- FlashSystem 9000 – online interactive demonstration (web)

More Information

Architecting IT is a brand name of Brookend Ltd, an independent consultancy, working to explain technology and business value to the end customer. This report was commissioned by IBM, however editorial control of the final published document remains with Brookend Ltd.

Email: architectingit@brookend.com
Twitter: @architectingit

The Author

Chris M Evans has worked in the technology industry since 1987, starting as a systems programmer on the IBM mainframe platform. After working abroad, he co-founded an Internet-based music distribution company during the .com era, returning to consultancy in the new millennium. Chris writes a popular blog at https://www.architecting.it/blog, co-hosts the Storage Unpacked podcast, attends many conferences and invitation-only events and can be found providing regular industry contributions through Twitter (@chrismevans) and other social media outlets.