

IBM Releases a Comprehensive, Software-Defined Storage Family

Situation

Data creation is growing at unprecedented rates with the amount of information stored doubling every year. The rapid expansion is being fueled by a new wave of application types, driven by the mobilization of data. Since the speed at which these applications are introduced has increased, individual organizations sometimes utilize “shadow IT” development that may not conform to best IT practices or support capabilities. Data is now more active throughout the lifecycle, with analytics driving competitive advantage and compliance needed to satisfy legal obligations as well as governance. Service-level requirements have increased, being driven by the user’s need to access data quickly and from anywhere 24 hours a day, 365 days a year. This confluence of events has placed additional strain on an already burdened storage administration staff that is limited in financial resources.

Additionally, much of the traditional storage infrastructure currently in place was installed for contemporary internet and client-based applications that drove specialized, mixed workloads. The variety of workloads required different physical storage products, technologies, protocols, APIs, and data types. This diversity of storage, which was successful in meeting the needs of contemporary applications, ended up creating independent storage silos that cannot share resources, services, or administrative processes. This hierarchical and proprietary storage is too expensive, underutilized, and requires too much manual effort to administer for the new dynamic IT environment. These storage islands also hinder data analytics, limiting the breadth of information used in an analysis. This situation drives IT to add additional capacity to an individual silo while overall utilization sits well below 50%. It is clear that a new way of procuring, provisioning, and managing storage is needed to minimize escalating costs while meeting user expectations.

In support of these needs, software-defined storage (SDS) is an approach that aims at making data and storage application aware and hardware agnostic, enabling server administrators, application managers, and developers to provision storage and services in a policy-driven, self-service manner. SDS enables this through the federation of heterogeneous storage, presenting a single, unified front from which storage resources and services can be provisioned and utilized. Not only can different storage vendors and architectures be aggregated into a single pool through SDS, but lower cost commodity hardware can be utilized for net new capacity or when legacy storage systems are retired. There is a fair amount of diversity across first generation SDS solutions, in large part because of the lack of a common industry definition. There are a number of capabilities that an SDS solution should have to effectively meet the needs of new and legacy applications while addressing rising CAPEX and OPEX costs.

- **Heterogeneous Storage** – As IT organizations take advantage of new hardware and software storage technology, it is not economically feasible to replace all of the existing storage infrastructure consisting of different vendors, protocols, and data types. An SDS solution needs to federate all storage, enabling capacity and services to be provisioned and consumed more easily and efficiently.
- **Commodity Hardware** – By having the storage intelligence reside in software instead of proprietary hardware, commodity hardware can be used to lower cost, simplify deployment, and enable increased flexibility. Commodity hardware should co-exist with legacy heterogeneous hardware, replacing the legacy storage when that hardware’s economic life is exhausted.
- **Common Management Framework and Services** – Using a single administrative tool to manage all storage eliminates the needs for specialized staff, reduces the risk of error and, improves efficiency in utilization and data movement. Automated provisioning prevents application errors and reduces cost by removing the need to over allocate storage.

- **Open Architectures** – Support for industry standards including OpenStack and Hadoop lowers costs and leverages the strength of open source communities.
- **Data Tiering** – Automatic data placement across storage systems, media, and cloud meets service levels and reduces cost. By having data stored on the right technology given where it is in the lifecycle, both performance and cost savings can be realized.

IBM Announces Spectrum Storage Family

On February 17, 2015, IBM announced the Spectrum Storage family of storage products. This family represents IBM's solution for a software-defined storage solution that addresses

a wide range of environments. The Spectrum Storage products can be broken down into two categories. The first category is Data Access. Spectrum Virtualize, Accelerate, and Scale represent the three Data Access architectures that IBM supports with its SDS offering. The second category is storage management consisting of Spectrum Control, Protect, and Archive. These products address provisioning, managing, protecting, and moving data as easily and seamlessly as possible. In addition to offering the storage solutions as software to run on in-house or distributed infrastructure, public cloud such as IBM's SoftLayer Cloud service can be utilized as a data target.



Spectrum Storage Data Access

Spectrum Accelerate, Virtualize and Scale span three unique data access environments. These software-defined storage products are delivered via software and based on existing IBM technologies, which increases deployment flexibility, lowers risk, and enables the leveraging of existing investments.

- **Spectrum Accelerate** – Based on proven enterprise class XIV technology, Spectrum Accelerate delivers block data storage and services across existing special purpose appliances or commodity based infrastructure. What is new is that Spectrum Accelerate is being offered as a software-only solution, giving customers greater flexibility in deployment with existing or commodity infrastructure. Spectrum Accelerate is targeted at lower cost environments such as cloud services which need to respond quickly to changes. Running as a virtual machine, a server-based storage area network can be built from commodity x86 servers, Ethernet switches, solid state, and high density disk drives. Accelerate operates on premises, off premises, or in a hybrid deployment model. Public cloud infrastructure such as SoftLayer self-service infrastructure as a service (IaaS) can be used to extend the on-premises infrastructure and is procured in a simple per TB pricing model.
- **Spectrum Virtualize** – Based on SAN Volume Controller, **Spectrum Virtualize** creates a single storage pool and management framework across heterogeneous SAN-attached storage systems. Spectrum Virtualize transforms what were previously isolated storage silos from different vendors, enabling provisioning to cross system boundaries and greatly increase storage utilization while significantly simplifying administration. Real-time compression further increases efficiency, and enables up to five times the data to be stored on the existing hardware.
- **Spectrum Scale** – Based on IBM's distributed file system or Elastic Storage, Spectrum Scale delivers high performance scale out file and object storage. Spectrum Scale is targeted at big data and Hadoop-based clients and HPC and OpenStack users. The distributed file system provides a single, global name space across multiple systems and sites. Spectrum Scale features automated tiering and lifecycle management, utilizing storage technologies from flash to physical tape.

Spectrum Storage Management

Spectrum Control, Protect, and Archive provide a single framework for administration, protection, and archiving. These products enable greater efficiency, control, and agility across in-house, distributed, cloud, and hybrid storage deployments.

- **Spectrum Control** – Based on Virtual Storage Center, Spectrum Control is a single point of management control for all types of data on heterogeneous disk, flash, file, and object storage systems. Spectrum Control utilizes analytics to automate provisioning and move data to the most cost-effective tier.
- **Spectrum Protect** – Based on Tivoli Storage Manager, Spectrum Protect is a multi-dimensional data protection program that can protect virtual, physical, and cloud data. Backup, snapshot, archive, recovery, and space management services are all available. Deduplication data reduction and incremental forever backup options greatly reduce the amount of space needed to store backups.
- **Spectrum Archive** – Utilizes native tape as a low cost archive tier. The linear tape file system (LTFS) makes accessing files stored on LTFS-formatted media similar to accessing files stored on other forms of storage media such as disk or flash drives. This enables the tape tier to inhabit the same global namespace as the primary disk storage. This simplifies moving data to and from the lower cost tape.

	Virtualize	Accelerate	Scale
Storage	SAN-attached storage	Direct-attached storage, Softlayer cloud	SAN & Direct-attached storage,
Need	Simplify management & improve efficiency	Reduce cost, improve flexibility & speed in deploying storage for cloud & other	High performance & scalable file/object storage. Multiple locations or countries
Target	Multiple heterogeneous, storage systems	Efficient, flexible storage for clouds & remote/branch office	Big data & analytics; unstructured data; HPC; OpenStack users

EMA Perspective

The announcement of the Spectrum Storage family by IBM supports a shift from hardware-focused, proprietary storage hardware to an approach where the storage intelligence and services are delivered in software. Software-defined storage is the next logical next step in the evolution to the software-defined data center, following on the success of server virtualization. With the Spectrum storage family, IBM is leveraging existing and proven software technology that has been deployed in hundreds or thousands of sites, embellishing it as required to add incremental capabilities or tighten integration with other elements of the Spectrum Storage family. The storage can be located within the data center, distributed environments, and even the cloud, supporting a “data anywhere” approach that meets the needs of new applications.

By delivering a comprehensive family of products, IBM Spectrum Storage can mix and match Spectrum Storage products to address the application and data requirements of each individual enterprise. With an aggressive commitment of \$1 billion in Spectrum Storage development over the next five years, IBM is making a significant investment in its software-defined storage future. The wait for new capabilities will not be long, as the multi-cloud gateway is already in the alpha phase, and will be available later this year. The multi-cloud gateway enables multiple cloud service providers to be used for applications such as backup, disaster recovery, tiering, archive, and data sharing.

EMA applauds the approach taken by IBM. Spectrum Storage delivers the benefits of software-defined storage to current and future users of IBM’s storage hardware platforms, as well as customers that wish to deploy a commodity based storage infrastructure. With the steady and inevitable evolution to a software-defined data center, IBM is well positioned to lead their customers through the transition.

About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA’s clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at www.enterprisemanagement.com or blogs.enterprisemanagement.com. You can also follow EMA on [Twitter](#), [Facebook](#) or [LinkedIn](#).

3107.032315