

July 22, 2016

A Storage Solution for a Hybrid World

FalconStor's Infrastructure-Agnostic Storage Platform

Stratecast Analysis by
Lynda Stadtmueller



Stratecast Perspectives & Insight for
Executives (SPIE)

Volume 16, Number 26

A Storage Solution for a Hybrid World

FalconStor's Infrastructure-Agnostic Storage Platform

Introduction¹

Data is enjoying its rightful place in the sun, as businesses recognize and capitalize on the value to be derived from collecting and analyzing metrics about every aspect of business operations, consumer behavior, and environmental situations. And yet, even as data worship grows, storage—the way data is retained for use—has remained mired in the past. Although technologies such as virtualization have transformed the way compute and network functions are delivered, storage remains largely tethered to purpose-built physical machines running proprietary software.

To be sure, storage vendors and service providers continue to enhance their products to address escalating enterprise needs for greater speed, capacity, and simplicity in their storage solutions. In recent years, leading storage vendors, including EMC, HPE, IBM, Dell, and NetApp have introduced substantial innovations to their product portfolios, such as flash technologies, converged and hyperconverged infrastructures, and even software-defined functionality. In addition, cloud service providers have capitalized on storage growth with cloud-based database and object storage services.

To take advantage of most innovative storage solutions, enterprises are forced to discard their legacy environments and start over, performing the dreaded “lift and shift.”

However, to take advantage of these innovative solutions, enterprises are forced to discard their legacy environments and start over, performing the dreaded “lift and shift.” As a result, enterprise storage professionals are more likely to deploy next-generation products for their *next* storage workload, while leaving existing storage as is. Each new stand-alone system increases the overall complexity of the storage environment, increasing the administrative burden and preventing efficiency and consistency across storage solutions.

In response, software providers are beginning to develop software-based management solutions that offer tools for managing a heterogeneous storage environment. Among these, FalconStor's FreeStor platform shows promise as a sophisticated software-defined storage platform that enables businesses to derive maximum value from their storage.

The Problem with Storage

The basic problem with data storage comes down to physics. Back in 2010, Dave McCrory (now CTO of distributed database provider Basho Technologies) introduced the concept of “data gravity.” McCrory wrote,

Consider Data as if it were a planet or other object with sufficient mass. As Data accumulates (builds mass), there is greater likelihood that additional Services and Applications will be attracted to this data...Service

¹ In preparing this report, Stratecast conducted interviews with:

- FalconStor – Tim Sheets, Vice President, Marketing

Please note that the insights and opinions expressed in this assessment are those of Stratecast, and have been developed through the Stratecast research and analysis process. These expressed insights and opinions do not necessarily reflect the views of the company executives interviewed.

and Applications can have their own Gravity, but Data is the most massive and dense, therefore it has the most gravity. Data if large enough can be virtually impossible to move.²

McCrory's apt metaphor supported an argument for placing data close to the applications that access it, rather than in a centralized cloud database.³ But data gravity also explains the challenges enterprises face in migrating and replicating stored data—and why they are reluctant to make changes to their existing storage, even if a new system could offer greater benefits. Data gravity can explain why enterprises are less likely to move storage to the cloud than other business workloads: 57% of IT decision-makers surveyed by Stratecast cite “data migration challenges” as the primary restraint to cloud storage.

Moving stored data takes a lot of effort, capacity, and bandwidth—and is almost sure to disrupt operations. As data volumes increase, the challenges exacerbate. Therefore, the place where data is created or initially stored tends to be where it remains.

Enterprises' reluctance to move existing storage, and the dearth of adequate migration solutions, creates additional challenges for the IT department and the enterprise as a whole. From a tactical perspective, storage administrators face:

- Administrative complexity in deploying and managing multiple, discrete storage systems
- Inefficiency in utilizing storage resources
- Inconsistent visibility and performance reporting across storage components

Of greater concern, as businesses prepare to compete in the fast-paced digital economy, is that the storage environment can be a hurdle to achieving strategic business goals. For example, the heterogeneous, inflexible storage environment may:

- Inadequately support the requirements for high-availability, always-on workloads
- Impede application performance (measured in input-output operations per second and throughput)
- Introduce delay into timelines for application deployment and scaling
- Introduce business risk associated with insufficient backup and recovery
- Hinder achievement of a hybrid IT environment that comprises physical and virtual, premises-based and cloud applications

Limited Leverage of Virtualization and Software-Defined Storage Technologies

As noted, storage may be the last bastion of old-school IT, dominated by purpose-built physical devices with proprietary software, which do not easily integrate with other systems and models. Most storage vendors have made limited use of technologies and models that are revolutionizing other IT components, such as servers and networks.

² Dave McCrory, [Data Gravity – in the Clouds | McCrory's Blog](#).

³ Stratecast explored architectures for placing compute close to data sources in SPIE 2016-22, [The Fog Rolls In: Network Architectures for IoT and Edge Computing](#)

That's not to say that storage vendors have stopped investing in innovation: in fact, they are actively introducing enhancements to their product portfolios that are designed to make storage workloads faster, more efficient, cost-effective, and flexible. Included among the enhancements are elements of *storage virtualization* and *software-defined storage*.

Over the past few years, storage vendors have introduced virtualization technologies to shift more functional control from the hardware to software platforms. With storage virtualization, a software layer (functioning like a hypervisor on a virtualized server) abstracts the storage control plane from the data plane, enabling networked storage devices to be pooled and centrally managed.

More recently, a number of vendors have started describing their products as “software-defined storage.” Although the term is sometimes used interchangeably with storage virtualization, software-defined storage usually represents a higher level of functionality available through the central controller, across all network-connected clusters. Services such as deduplication and replication, which are generally tied to the hardware, may be software-based in software-defined storage

solutions. Software-defined storage is often associated with hyperconverged solutions, in which storage, computing, and network components are pre-packaged together; the software-defined components are integrated to operate and scale together.

The ironic thing about most software-defined storage systems is that they are tied to vendor-specific hardware. The vendors are not fully

The ironic thing about most software-defined storage systems is that they are tied to vendor-specific hardware. The vendors are not fully decoupling the logic from the hardware; instead, they're merely enabling the software controller to pool and manage the vendor's own hardware devices.

decoupling the logic from the hardware and freeing users to use any hardware of their choice; instead, they're merely enabling the software controller to pool and manage the vendor's own hardware devices. This limits its use to new installations or sometimes previous-generation hardware that the vendor has made backward-compatible.

FreeStor: An Infrastructure-Neutral Software-Defined Storage Platform

A contrasting solution is offered by FalconStor, a long-time provider of software-based storage services. The company's FreeStor platform is a sophisticated, software-defined storage solution; that is, it abstracts the control logic from the underlying storage hardware. In addition, the platform also acts as an intelligent “data services” layer.

The first important difference between FreeStor and other vendors' software-defined storage solutions is that it is hardware-agnostic. That feature enables enterprises to retain their existing multi-vendor, heterogeneous storage environment, and tie it together via the “intelligent abstraction platform.”

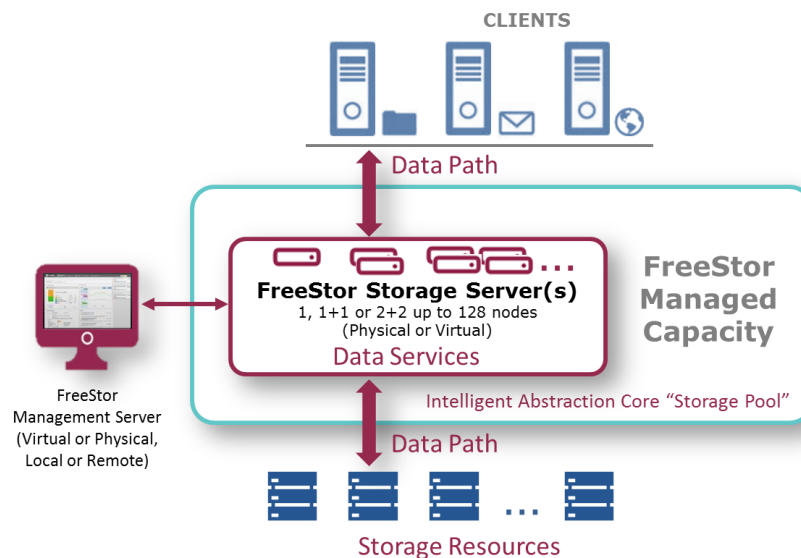
Because it is infrastructure-independent, the platform is uniquely designed to serve the hybrid environment that IT organizations are building. More than a data repository, FreeStor facilitates data flow and services between applications and platforms (cloud and premises-based; open and proprietary; physical, virtual, and hybrid) and storage resources (cloud and premises-based; flash, HDD, and tape; physical, virtual, and hybrid). As organizations increase the deployment models and environments from which they collect, share, and utilize data, a fragmented, vendor-specific storage management approach will be unsustainable.

FreeStor Architecture

The FreeStor solution works by inserting intelligence into the data path between the application servers and the storage hardware. As shown in Figure 1, the solution is deployed using two types of servers:

- FreeStor Storage Servers (available as a virtualized or physical server) are installed in front of the storage systems. These servers perform the data abstraction (virtualization) functions, and apply data services (e.g., snapshots, replication) to the storage systems.
- The FreeStor Management Server manages the resource pooling, and controls data services. The Management Server connects to up to 128 FreeStor Storage Servers, and runs the centralized FreeStor management portal.

Figure 1: FreeStor Topology



Source: FalconStor

The insertion of an additional component into the data path does not significantly increase delay; the company says the latency burden is approximately 130-150 microseconds. While the solution requires enterprises to deploy additional hardware components into the storage environment, the ability to utilize virtualized versions of the platform on commodity servers minimizes the investment. Furthermore, the virtualized option enables enterprises to deploy a cloud-based FreeStor Storage Server to “front” cloud storage services such as AWS S3 and Azure StorSimple, as well as OpenStack-based services.

Defying Data Gravity: Simple, Disruption-Free Migration and Replication

FreeStor’s architecture and powerful software platform mitigate many issues associated with data gravity. Using the unified FreeStor console, abstracted (virtualized) data can be easily migrated, backed up, replicated, or recovered across locations or deployment models or hardware, without the risks associated with data movement. As a result, enterprises have tremendous flexibility in deploying and managing storage resources across multiple storage platforms and locations, all from a single pane-of-glass.

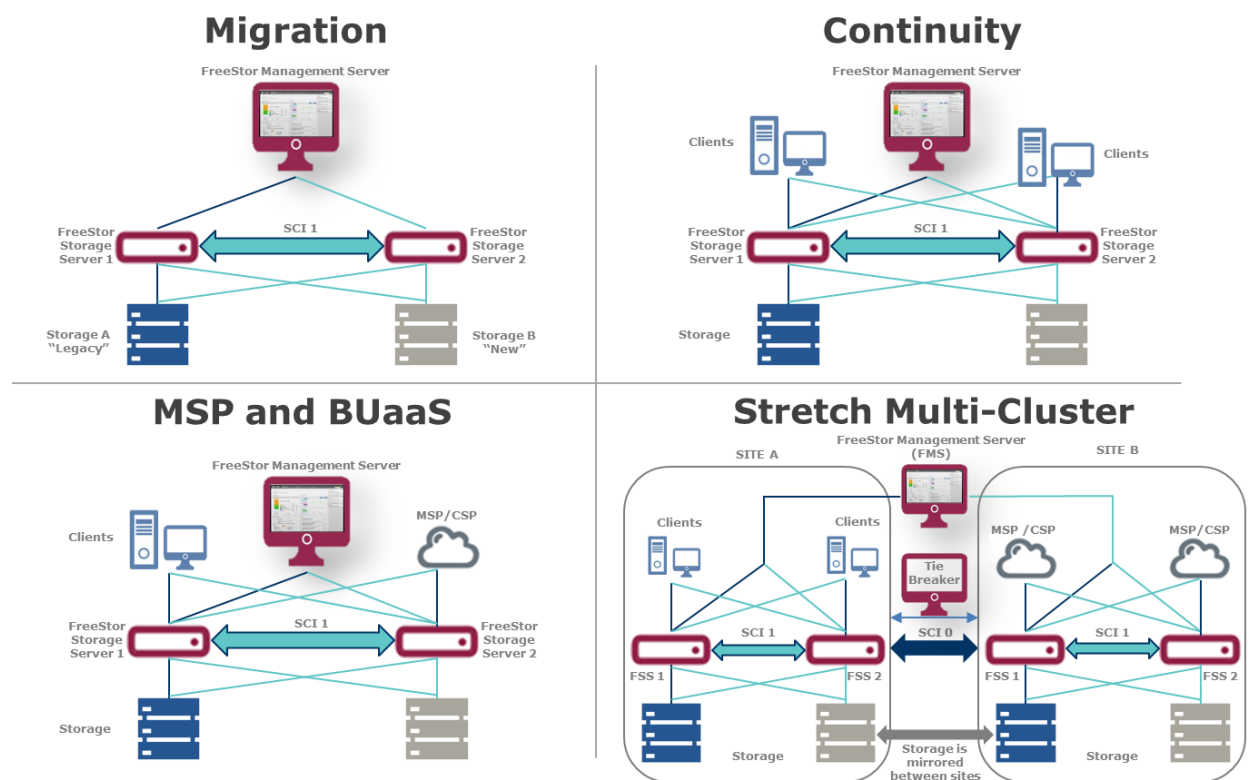
Central to FreeStor's value is its "active-active" high availability engine that works on disparate, heterogeneous hardware. The company deploys the engine in a 1+1 cluster pair (either locally or in a stretched configuration across different geographies). The engine can also support 4-way clusters in a 2+2 configuration, with a local, metro or stretched configuration (illustrated in the lower-right example in Figure 2). Integrated WAN optimization functionality minimizes network-induced latency into the replications. In addition, users have the option to deploy in-line deduplication for storage workloads that can benefit from that functionality.

The high-availability capability allows users to easily decommission or add hardware; migrate to cloud storage; create high-availability or disaster recovery environments, all with a recovery time objective (RTO) of zero.

Figure 2 shows four scenarios of how users might configure their storage environments for high availability:

- **Migration** from legacy to new storage hardware.
- **Continuity scenario**, in which both the storage and the application is replicated.
- **Managed Service Provider (MSP) and Backup as a Service (BUaaS)**, in which a service provider utilizes the FreeStor platform to offer high-availability services to clients.
- **Stretch Multi-Cluster**, enabling mirroring between physical locations, each of which supports multiple clusters.

Figure 2: FreeStor Replication Scenarios



Source: FalconStor

A Storage-Wide Pool of Data Services

The high-availability engine is “behind the curtain,” built into the FreeStor platform. What is visible to users is a rich pool of data services or functions that are easy to deploy and manage across all storage infrastructures, through a single console.

Services supported by the platform include:

- Migration services
- Backup
- Disaster recovery
- Archiving

Most recently, FalconStor introduced a “cloud connector” capability specifically designed to migrate premises-based workloads and data to Amazon Web Services and Microsoft Azure clouds, as well as OpenStack environments. The functionality supports the use of the cloud for backup and recovery. It also facilitates “cloud-hopping” (that is, the ability for customers to easily move workloads among cloud service providers).

While the platform is shipped with all services included, FalconStor offers users the budget-friendly option of paying for only the services that they activate.

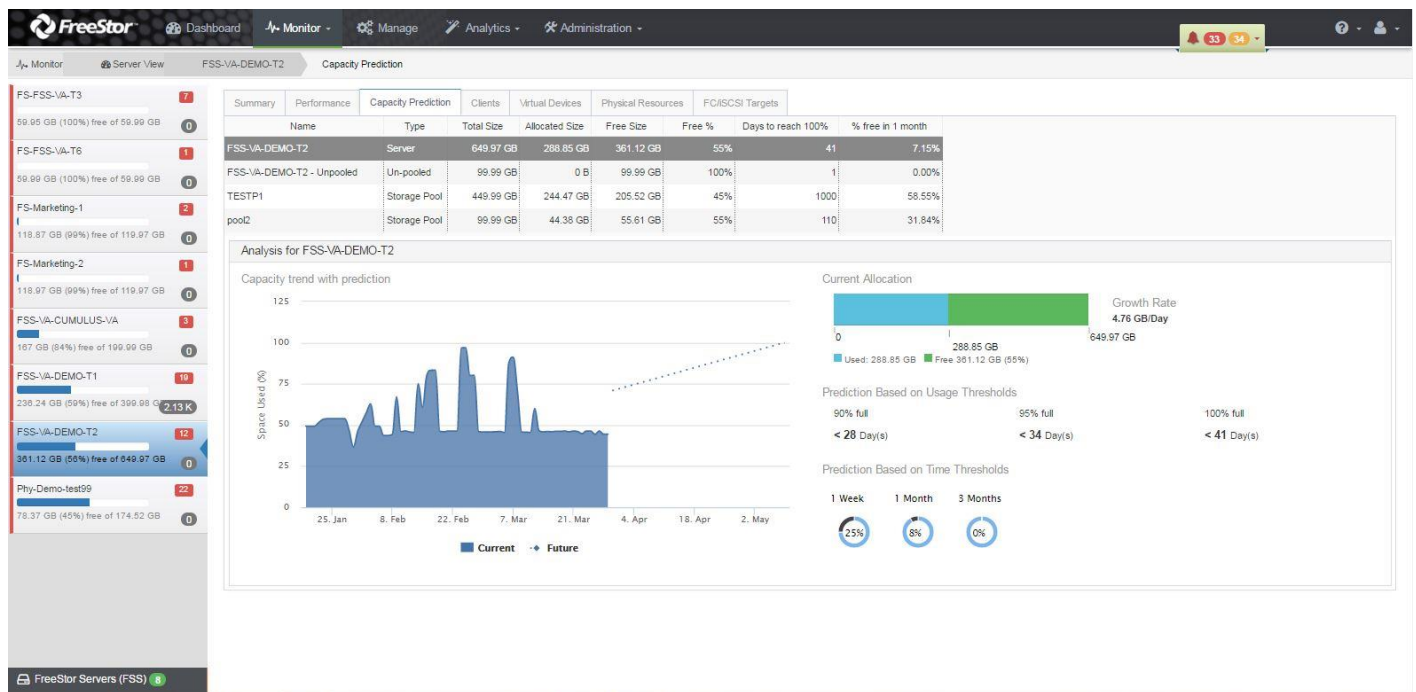
Applying Intelligence and Predictive Analytics Across the Storage Environment

For IT leaders, the most valuable capability of the FreeStor platform may be its analytics functionality. The platform offers visibility and reporting across the entire storage environment, via a single console. Granular, device-level performance and health-monitoring data enables users to:

- Ensure capacity is optimally allocated across resources—i.e., capture, reclaim, reprovision, forecast future capacity consumption, based on real time trend data
- Scale storage as needed
- Monitor devices for maintenance issues
- Ensure that storage performance (measured in Input/Output Operations per Second; latency; throughput) is appropriate for the associated applications
- Deliver and measure attainment of service level agreements for internal clients (enterprises) or external clients (MSPs)

Recently, FalconStor has added predictive analytics capabilities to its intelligent platform. Utilizing historical and trending data, these capabilities enable users to predict capacity utilization across the entire storage infrastructure environment. Figure 3, below, is an example of a report.

Figure 3: FreeStore Predictive Analytics - Sample Report



Source: FalconStor

The real value to enterprises is that the insights revealed via the monitoring and analytics functions are immediately actionable. Consider:

- If the management dashboard reveals a sudden decrease in throughput for a particular device, the storage can be activated on the mirrored device, while local teams perform maintenance.
- If storage managers want to see whether the enterprise might benefit from newer, flash-enabled technology, they can try out the technology on one or more workloads.
- If managed service providers want to create new value-added services that align storage performance with application requirements, they have the tools to monitor and manage the outcome.

The intuitive management portal—which is available in a mobile as well as Web version—makes it easy for storage managers to quickly enact changes to the storage environment based on dashboard data or alerts. Perhaps equally important in a fast-paced business environment, the managers can try out various options, making one decision today and another tomorrow, without wasting time or disrupting the environment.

Stratecast The Last Word

With the FreeStor platform, FalconStor enables businesses and service providers to simplify and optimize the storage environment, and, at the same time, to more easily derive value from their storage. Utilizing the functionality integrated into the platform, an enterprise can:

- Unify its storage environment, centrally managing resource pools, and consistently applying services, regardless of underlying hardware or deployment model
- Deploy, migrate, and replicate storage without disrupting operations or losing data
- Optimize storage and application performance, by taking action based on granular performance data and predictive analytics
- Introduce new storage technologies and deployment models, such as flash and cloud, into the storage environment without requiring a disruptive “lift-and-shift” change
- Protect data and applications through high-availability options
- Control capex costs via optimal storage capacity utilization
- Ease the storage administration burden, by minimizing the number of vendor-specific storage specialists required, and leveraging the easy-to-use management and services portal

As hybrid environments become the norm, businesses will have little time and patience for inflexible, stand-alone components in the data center. While storage, in general, has lagged behind components such as servers and networking in embracing software-defined solutions and infrastructure-neutral management tools, FalconStor has more than made up the gap with its FreeStor platform.

Lynda Stadtmueller

Vice President – Cloud Computing

Stratecast | Frost & Sullivan

lstadtmueller@stratecast.com

About Stratecast

Stratecast collaborates with our clients to reach smart business decisions in the rapidly evolving and hyper-competitive Information and Communications Technology markets. Leveraging a mix of action-oriented subscription research and customized consulting engagements, Stratecast delivers knowledge and perspective that is only attainable through years of real-world experience in an industry where customers are collaborators; today's partners are tomorrow's competitors; and agility and innovation are essential elements for success. Contact your Stratecast Account Executive to engage our experience to assist you in attaining your growth objectives.

About Frost & Sullivan

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies? For more information about Frost & Sullivan's Growth Partnership Services, visit <http://www.frost.com>.

CONTACT US

For more information, visit www.stratecast.com, dial 877-463-7678, or email inquiries@stratecast.com.