

White Paper

# Architecting the Hybrid Cloud

## Unifying Data Access with a Cross-Cloud Scale-Out File System

By Scott Sinclair, ESG Senior Analyst

November 2016

This ESG White Paper was commissioned by Elastifile and is distributed under license from ESG.



---

## Contents

Introduction .....	3
The Challenges of the Old and the New Storage Silos .....	5
Architecting Modern Hybrid Data Center and Agile Workflows with a Scale-out File System.....	6
The Bigger Truth.....	7

## Introduction

As modern industry reaps the productivity benefits that come with a digitally enabled world, one counterbalance has caused a seemingly exponential increase in rate of market change. As quickly as innovations emerge, both in business and in technology, customer demands and market dynamics shift just as quickly, requiring new business models and new ways to interact with customers. In these volatile times, the strength of a business is often defined by its ability to adapt to changing market conditions. And in this digital age, the limiting (or enabling) factor in a business' level of agility is how the business interacts with its information, i.e., with its data. In this respect, few, if any, industries have been immune.

For the IT organizations that support this data, the demands have increased exponentially. Only a few years prior, many IT organizations primarily focused on maintaining infrastructure. Servers, networking, and storage components needed to be maintained. Data had to be backed up and protected. Today, however, just maintaining the infrastructure is not good enough—IT leaders must become stewards of the business's data, able to make the right data accessible where and when the business needs it, while keeping budgets under control.

Responding to this surge in demands, the IT and storage vendor community has delivered a cornucopia of new technology options to help IT organizations manage. These innovations include new storage media types, such as solid-state and all-flash storage. This trend has also included new delivery mechanisms for IT infrastructure such as software-defined storage (SDS), converged, hyperconverged, and the cloud. That is the good news. The bad news is that an IT architect must now navigate through this mass of new technologies to find the right solution for his or her particular organization's needs. Additionally, when demands change, the investigation process starts over again. As valuable as these new technologies are, many of them do not ease the burden of transitioning data from one technology to the next. In contrast, many of these solutions often introduce yet another storage silo to manage, albeit a likely more capable one than the solution it replaced.

Architecting a data center best suited to manage the dynamic evolution in IT demands requires an infrastructure that offers not only performance and scale, but also agility. The optimal infrastructure is one that can evolve with the needs of the data center and assign the right capabilities to the right workload at the right time. One method to achieve this necessary agility would be to deploy an intelligent data management layer that can span multiple resources and media types and on- and off-premises infrastructure, while providing performance, scale, or the right blend of both.

One technology that has the potential to meet all these requirements is a scale-out file system. A new entrant to the enterprise IT arena, Elastifile, has developed a scale-out file system designed to provide a single global namespace as the agile data management layer for the modern business. Far more than just another SDS repository for unstructured data, Elastifile is leveraging its technology to serve multiple workloads and consolidate on- and off-premises infrastructure resources across a single global namespace for file and object data. The intent behind this solution is to consolidate and unify data access, allowing businesses to cost-effectively keep pace with changing market conditions.

## The Perpetual Challenge of Data Growth and Accessibility

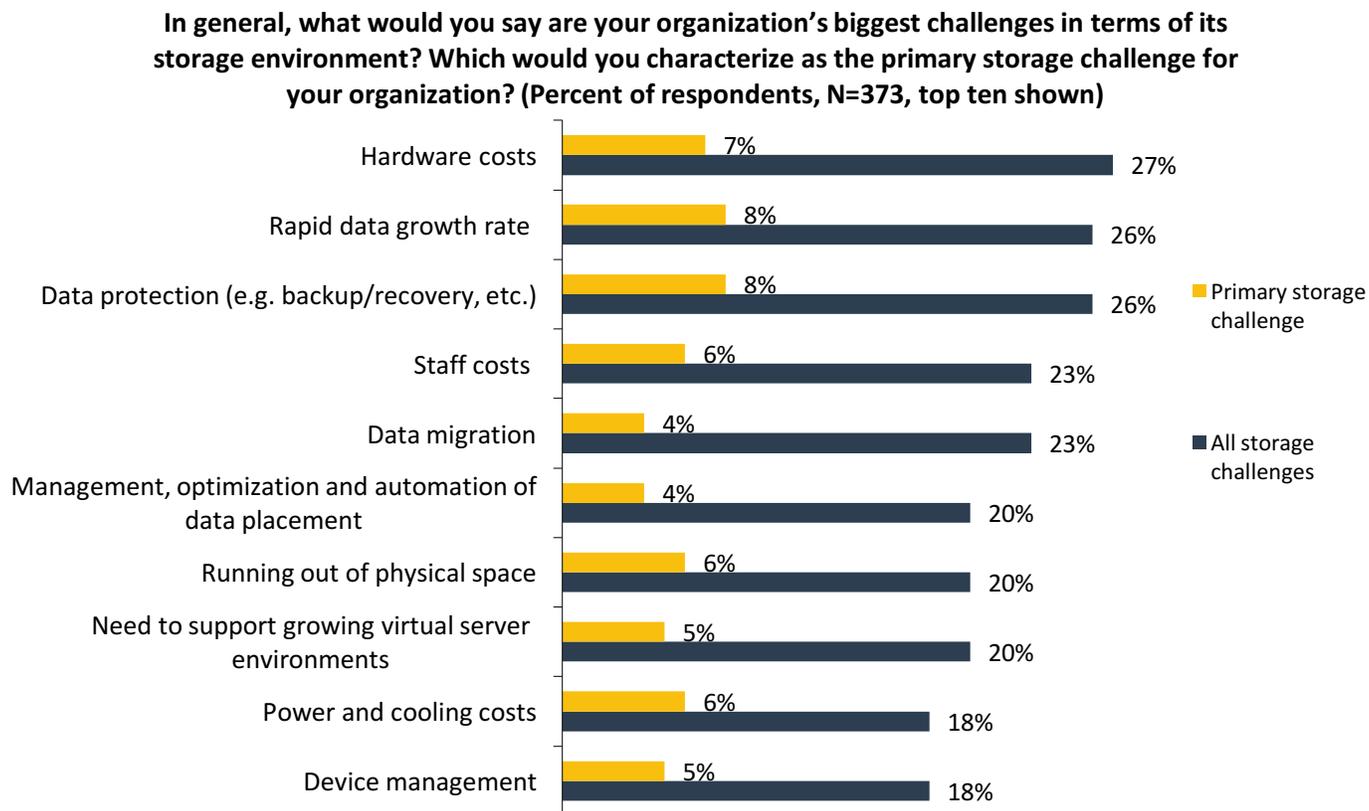
Recently, ESG surveyed 373 IT decision makers responsible for their organizations' data storage infrastructures as part of a research study investigating general storage industry trends. In response to a question asking storage leaders about their top current storage challenges, storage decision makers identified that the rapid growth of data and the associated infrastructure challenges continue to be top pain points for their organizations.<sup>1</sup>

Although it is technically a second most-cited overall storage challenge, the rapid rate of data growth was actually tied for the most popular response when storage decision makers were asked to label their primary storage challenge (see Figure 1). Looking at the rest of the top ten, however, it is not difficult to see a trend. Hardware costs, data protection, staffing, and data migration are all impacted and exacerbated by the rapid rate of data growth and accessibility. Essentially every

<sup>1</sup> Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

response on the list of top ten challenges can be considered to be a manifestation of the impact data growth has on the data center.

**FIGURE 1. The Biggest Storage Challenges**



Source: Enterprise Strategy Group, 2016

Additionally, beyond the list of responses in Figure 1 are what many would describe as new challenges and requirements. Managing the complexities of data growth and increasingly integrated workflows like analytics, while constantly ensuring that content is protected and available remains a persistent concern. These challenges persist despite numerous innovations designed to specifically address the challenges of data growth, such as larger capacity drives or scale-out storage. For many organizations, the cumulative impact of these challenges has either rendered the existing storage infrastructure obsolete or has placed it on a path to become so soon.

When faced with the unsustainability of existing infrastructure, some IT organizations have turned to investigating public cloud services. In an ESG research study investigating IT spending intentions, when IT leaders were asked to identify which measures their organizations would take to reduce or otherwise contain IT expenditures over the next 12 months, they most commonly identified the increased use of cloud computing as an alternative to in-house infrastructure.<sup>2</sup> While this data point indicates strong perceptions of the cloud as a cost-effective alternative to on-premises infrastructure, the cloud is not ideal for every business or for every workload.

For example, in the ESG study of storage leaders, the combination of business intelligence and data analytics workloads was the most commonly identified workload expected to drive capacity growth over the next two years.<sup>3</sup> When IT professionals familiar with their organization’s business intelligence and data analytics environments were asked in a

<sup>2</sup> Source: ESG Research Report, [2016 IT Spending Intentions Survey](#), February 2016.

<sup>3</sup> Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

separate study about their organization's primary deployment strategy, 81% identified on-premises infrastructure, revealing a significant disparity from public cloud (12%) or hybrid cloud (6%).<sup>4</sup> When these IT and business professionals were then asked to identify the disadvantages or barriers to adoption of public cloud resources, the top concerns included storage costs, processing costs, and security. The takeaway is that there is no one-size-fits-all media type for workloads. IT organizations need the agility to select the optimal technology when it best suits the business.

## The Challenges of the Old and the New Storage Silos

For years, enterprise storage architectures have been focused on easing the ability to capture data and then protect it. When data was to leave the system, the expectation was often that it would be read by the application or copied to another location for backup protection or disaster recovery. Rarely do designs consider the idea that an organization may need to migrate large data sets out of the storage system on a regular basis, including increasingly critical on-demand workflows like analytics that need to pull together data from disparate sources and sites. Storage architectures, whether intended or unintended, often isolate the data they store from other devices, creating a silo of storage. Historically, the term storage silo has been used to refer to a storage system that lacked the scale necessary to keep pace with data growth. More modern systems have been architected with the intention of eliminating this particular challenge, adding higher capacity scaling or a scale-out architecture. Yet even these new systems often create storage silos not of scale but of purpose. Often designed to solve a particular storage challenge while forsaking others, the characteristics and challenges of these new storage silos include:

- **Singular benefit or focus:** Whether an all-flash array for high performance, scale-out storage for unstructured data, or a public cloud solution for application development, storage solutions often target a single challenge or workload SLA, focusing on a particular set of benefits. While solving these individual challenges provides value, when application demands change over the life of the data set, the inflexibility of the architecture creates inefficiencies.
- **Data bound to its underlying infrastructure:** Whether in terms of performance, resiliency, or even security, data requirements change over time. The inability to adjust the underlying infrastructure as demands change adds complexity and increases costs. For example, the physical isolation introduced by WAN latencies from public cloud resources to the on-premises data center hinders data movement between locations. Over time, performance or security requirements can demand that data residing on the cloud be stored on-premises, or vice versa. If the architecture was not designed to consider this scenario, the result can be a substantial and unnecessary increase in migration costs.
- **Lack of application intelligence or self-service:** Even if a system is designed to serve multiple data types and move data between storage sites, without the necessary intelligence to understand which data should reside on which site, the benefits of the first two capabilities are severely hindered. Businesses often endure slow and painful processes to identify the needs of data and then migrate the data to the right location at the right time. Ideally, the storage architecture would be able to empower the application and analytics teams to have full visibility into the data and self-service control over data to maximize the business benefit.

Serving the needs of the modern business demands a storage architecture that is not bound to silos. Modern architectures need not only the scale to meet the demands of capacity growth, but also the ability to adapt as business needs change.

---

<sup>4</sup> Source: ESG Research Report, [Enterprise Big Data, Business Intelligence, and Analytics Trends: Redux](#), July 2016.

## Architecting Modern Hybrid Data Center and Agile Workflows with a Scale-out File System

Architecting an IT infrastructure that can effectively evolve with the business requires three fundamental characteristics. It must be able to abstract the workloads (and their data) from the infrastructure. It must deliver the intelligence necessary to help the consumers of that data consistently leverage the right infrastructure technology- whether on-prem or off-prem. And it must enable the ecosystem to be agile for workloads to move between sites when demands change. Delivering these three characteristics necessitates a certain type of storage technology. To deliver the necessary level of abstraction, the ideal solution would likely leverage software-defined storage to simplify deployment and ease the ability to integrate a broad selection of hardware. To deliver the right blend of scale, intelligence, and flexibility, that software would likely be best served by scale-out file system technology. A scale-out file system can deliver on these requirements in a variety of ways, including:

- **Abstraction:** With a software-defined storage architecture, a scale-out file system can abstract data access and control for the applications from the underlying hardware. With storage innovation rampant as of late, new technologies can be integrated behind the scale-out file system layer without impacting the workload, eliminating the downtime associated with provisioning storage and migrating data. This should be extensible for data formats across file, object, and block, and POSIX-compliant application semantics are the natural superset for the more limited object and block semantics. As such, scale-out file systems can unify data management across environments, locations, and applications with a single global namespace. To provide the maximum benefit, the solution must be able to not only span both off- and on-premises infrastructure, but should also offer the ability to integrate multiple cloud environments, as support for only a single cloud will likely become a limiting factor over time.
- **Intelligence:** Scale-out file system technology offers contextual data and application awareness through a deep visibility into metadata (e.g., file types and access patterns). This awareness is critical to supporting the intelligent and real-time decision making required for efficient data placement. In a hybrid cloud environment, this level of intelligence is paramount as WAN latencies hinder data movement across locations. In addition, the intelligence achieved through the understanding of metadata is a critical element to delivering self-service capabilities. File systems can provide policy-based control and copy services tailored to meet the needs of specific service level agreements (SLAs) or customer policies at the most effective and efficient level for the business—file, directory, or workload granularity.
- **Agility:** Unlike other storage architectures, file system technology is designed to be both user- and application-accessible, translating to a wider breadth of workloads that file systems can serve. Thanks to the prevalence of the file system's hierarchical structure, data management and access is efficient and intuitive. As a result, the file system interface is widely supported and understood by both users and applications.

Although scale-out file system technology offers elements designed to deliver these three capabilities, not all scale-out file systems can deliver the agile data management layer the industry so dearly needs. Many scale-out file systems were designed specifically to solve the challenge of storing massive quantities of unstructured data, and were not optimized to deliver the hybrid cloud flexibility critical to a data management layer. One new storage innovator, Elastifile, has developed a scale-out file system architecture to enable all this in a single global namespace, abstracted from the underlying infrastructure stacks and sites. As a hardware-agnostic software-defined approach, Elastifile leverages linear scale-out flash optimizations to deliver elastic performance for massive workload consolidations. Its patented distributed, adaptive metadata model enables both cloud-scale and consistent performance across noisy, heterogeneous, cross-cloud environments. While the firm is still new to the industry and will certainly have to prove itself, its founders bring deep storage, flash, and containers expertise, and its technology is a promising potential answer to delivering such a global data namespace.

## The Bigger Truth

All-flash storage, the public cloud, converged, and hyperconverged are just a few of the new technologies that are disrupting the traditional data center. As the demands to generate, access, and store data increase, businesses, and the IT organizations that serve them, will continue to seek out the best and most valuable new technologies to help them gain a competitive edge. The deciding factor in which organizations reap the greatest benefit will likely be whether they can quickly and efficiently integrate these new technologies with little or no disruption to their existing workloads. The ability to deploy and leverage the right technology at the right time will be determined in part by the ease of integrating the new innovations, and in part by the intelligence of the system to allow the users to self-service and select the right technology for their needs. Achieving this flexibility demands a software architecture that enables application and data access to be abstracted from the underlying hardware. Scale-out file system technology takes this capability one step further and offers a powerful blend of data intelligence and a breadth of workload consolidation services to deliver a data management layer across locations at cloud-scale. All file systems, however, are not created equal. Elastifile is leveraging this technology for the specific purpose of delivering an agile data management layer designed for hybrid cloud infrastructure. As business leaders and their IT organizations look to transform workflows like analytics and break their data out of the storage silos that currently hinder their environments, Elastifile is presenting a compelling option to consider.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.

© 2016 by The Enterprise Strategy Group, Inc. All Rights Reserved.

