

Solution Showcase

Scale-out Data Protection

How It Promises to Revolutionize Recovery Capabilities for Enterprises and Managed Service Providers


Date: October 2017 **Authors:** Jason Buffington, Principal Analyst; and Monya Keane, Senior Research Analyst

Abstract: Enterprises and managed service providers (MSPs) have some similar challenges in the realm of data protection. Both have to deal with expansive environments—with enterprises protecting distributed offices/geographies, and MSPs protecting many independent subscribers. Both also have been seeing production data storage and secondary protection storage demands rising. Enterprises and MSPs alike should therefore be looking for data protection architectures that have scale-out capabilities built into their core foundation.

Introduction

Research by ESG (see Figure 1) reveals that the data protection capabilities most frequently mandated by IT executives reflect a pattern. All the top mandates relate to executives simply “wanting better” when it comes to protection-related agility, reliability, and efficiency.¹

Figure 1. Top Five Data Protection Mandates from IT Leadership



Source: Enterprise Strategy Group

¹ Source: ESG Research Survey, *2017 Trends in Data Protection Modernization*, December 2016.

This ESG Solution Showcase was commissioned by Commvault and is distributed under license from ESG.


© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

The issue of cost control is always going to be on the mind of any IT leader. But in this case, mandates for cost control were not their universal, paramount priority. Many of the surveyed executives indicated that cost reduction, while important, was less urgent than attaining better protection-related functionality, nimbleness, and dependability.

ESG also asked hands-on IT implementers a similar question—inquiring about their greatest data protection *challenges* (see Figure 2).² As the senior executives did, many surveyed IT pros mentioned cost(s). But they also reported that inadequate reliability and agility are among the main challenges they and their organizations face.

Figure 2. Top Five Challenges with Current Data Protection Processes and Technologies

Which of the following would you categorize as challenges with your organization’s current data protection processes and technologies? Which would you characterize as the top challenges for your organization? (Percent of respondents, N=387)



Source: Enterprise Strategy Group

Clearly, both the “heads” (executives) and the “hearts” (hands-on implementers) desire something better.

Current environments are typically built upon scale-up architectures, and as the survey results imply, they are no longer meeting today’s IT needs. A big challenge with these scale-up architectures is their limited-scalability compute and storage, a limitation that impacts growth, service levels, and availability.

Of course, achieving something better doesn’t simply mean “replace what we have with something else.” Instead, a new approach, involving both a shift in mindset and a shift in architecture, is warranted. Truly achieving something better requires actively reimagining the foundation upon which data protection will be delivered.

Specifically, the new approach should center on *taking advantage of a scale-out architecture*. It will address many of the data protection challenges described above.

The Differences Between Scale-up and Scale-out Architectures

Scale-out architectures offer a number advantages over scale-up architectures. These advantages include:

- Dynamically expandable storage.
- Flexible services to end-users.
- Greater resiliency and availability.
- More predictable costs, and even reduced storage costs through standardized x86 hardware.

² *ibid.*

- More predictable performance.
- Greater freedom of choice in hardware.
- Being able to leverage the benefits of cloud services in an on-premises environment.

The Growing Importance of Scale-out in Data Protection/Data Management Architectures

Just as backups, snapshots, and replication are components of a comprehensive data protection function, data protection is itself part of an even broader data management strategy. Data management touches primary production data, secondary protection data, and tertiary copies retained for many business reasons.

Both managed service providers and individual organizations might benefit by looking into how a *scale-out data protection and management architecture* could help them. Of course, the specific consumption models, mechanisms, and methods will vary:

- **Large enterprises** deal with everything in greater volume. Yet many of them attempt to protect their sprawling data repositories using component-level solutions that create inherently isolated silos. For a large enterprise, the first step toward deploying scale-out data protection could be as simple as utilizing whatever capacity might be suitable/available across their myriad platforms.
- **Midsized organizations** also want to maximize the utilization of their available storage, although their repository componentry isn't likely to be as diverse as a large enterprise's. Some smaller and midsized organizations focus on looking for enterprise-caliber protection solutions that are right-sized for them. Others are rearchitecting their data protection strategies more dramatically—implementing a “cloud-first” vision that utilizes backup (BaaS) and disaster recovery (DRaaS).
- **Managed service providers** find scale-out data protection architectures appealing because they have goals in common with both enterprises and SMBs: They experience the same issues related to scale and IT heterogeneity that large enterprises do, and they need to support multiple independent SMB subscribers (similar, in a way, to how a large enterprise has to support multiple departments and branches).

If they haven't already started, it is almost inevitable that enterprises and MSPs will pursue scale-out data protection architectures to abstract and consolidate their diverse storage hardware into one manageable pool of optimized capacity. This trend is yet another validation of the value of and need for software-defined storage (SDS).

SDS is storage managed and automated by intelligent software instead of being managed in the “old-fashioned” way by the storage hardware itself. SDS, combined with a scale-out data protection/data management architecture, not only supports balanced data growth on an as-needed basis, but also supports multi-protocol access. The result is that applications can concurrently access data and capacity regardless of whether that data/capacity is onsite, in a secondary facility, or in the cloud.


Essentially, a good scale-out storage architecture uses its distribution and resiliency strengths to provide seamless extensibility and durability across multiple underlying repositories. The result is better cost control, higher performance, and all the other IT-oriented advantages described earlier, especially—importantly—*more agility* to serve a wider range of departments, remote sites, and subscriber clients.

Commvault HyperScale

[Commvault](#) has been innovating in the data protection market for more than 20 years and continually introducing evolutionary enhancements aligned with customers' evolving needs. In 2017, it announced [Commvault HyperScale Software](#) (see Figure 3) as the underlying enablement layer of the Commvault Data Platform.

The actual functionality and usage experience may appear quite familiar, but Commvault HyperScale Software now provides an extended file system within the repository, and it offers clustering/failover capabilities between nodes, thereby providing scale-out data protection.

Figure 3. The Commvault HyperScale Architecture



Source: Commvault

Supported by Commvault’s existing multi-tenancy features and distributed architecture, the new scale-out capabilities explicitly address the scale and flexibility challenges that enterprises and MSPs are looking for as they rearchitect themselves for the future.

The Bigger Truth

It appears that agility and efficiency are coveted by basically everyone involved in protecting and managing data—especially those people struggling to simultaneously keep up with sprawl and meet ever-heightening expectations. One answer to these storage-related challenges centers on introducing a software-defined layer that abstracts and normalizes underlying storage repositories while still enabling already-deployed best-of-breed componentry to do what it does best.

In the realm of data protection, the newest area to start seeing software-defined evolution is protection storage. Such storage requires highly efficient optimization intended specifically for data management, data protection, and data preservation at increasingly higher scales.

One company that has continually endeavored to innovate ahead of current market demand for more than 20 years is Commvault. As such, Commvault’s HyperScale Software delivers a scale-out architecture and should be considered as it evolves and expands the solution offerings through both MSP and on-premises software and appliance offerings.

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. 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.