

Solution Showcase

Architecting the Future Data Center with Multi Cloud Data Management

Date: February 2017 **Author:** Scott Sinclair, Senior Analyst

Abstract: The advent of public cloud services has introduced new questions in IT architecture design. Hybrid cloud ecosystems have become the norm. Workload deployments must select not only between on- and off-premises models, but also which public cloud service to leverage. The future data center will not simply be hybrid; it will be multi cloud. Future data centers will need not only the ability to pool multiple cloud infrastructures together, but also the flexibility and the agility to leverage the best aspects of each cloud, even as the cloud market landscape evolves. The future data center will demand multi cloud data management technology.

Overview

In a recent video, Steve Duplessie, ESG founder, delivers a warning to IT decision makers about the dangers of remaining devoted to the status quo in the wake of rapid technology evolution. It should come as no surprise that the IT industry has undergone a revolution fueled by the rise of public cloud services. As a result, many of the design tenets that IT infrastructure architects once held as fundamental truths may no longer be valid. The first of those truths that should be called into question is that business leaders want to care about IT infrastructure.

Letting the detailed minutia of IT infrastructure dictate how businesses interact with data has become a hindrance, and some of the largest culprits may be storage administrators. As Duplessie mentions in his video, “We still talk protocols. Millennials don’t talk protocols. Business people don’t talk protocols. But storage infrastructure people still do.” The bottom line is that necessitating an understanding of the specifics of storage infrastructure is unnatural and slows down business initiatives. The high growth rates of technologies that help to abstract away the specifics of the infrastructure, such as public cloud, hyperconverged, and even software-defined storage (SDS), help support this point.

These technologies alone, however, do not take abstraction and automation far enough. According to Duplessie, “Ninety-nine percent of the world does now, and will for the foreseeable future, live in a hybrid cloud world. You will have stuff on-premise[s]. You will use public cloud infrastructure—and you may even think it’s private. You will live as a hybrid or as a dinosaur. Those are the only two real options.” The bottom line is that the current and future data center will be a hybrid, a blend of on- and off-premises capabilities. As public cloud adoption increases, more IT decision makers will take a multi-vendor approach to public cloud services, similar to what is considered a best practice for on-premises infrastructure. To thrive in this emergent multiple cloud landscape, IT organizations will need multi cloud data management technology to ease and automate the management of these disparate and ever-evolving cloud elements.

The Need for Multi Cloud Data Management

Though already well established by some measures, the public cloud landscape is still evolving. As with any high growth industry, cloud service providers continue to innovate and deliver new capabilities as a means to differentiate offerings and capture market share. For IT leaders, the resulting cloud innovations will help businesses further maximize the value of their applications and their data. While all this sounds positive, the result will be increased complexity. As businesses endeavor to identify, integrate, and leverage these new cloud capabilities, deployment or migration delays and management complexity likely will generate unnecessary costs.

In a future where multiple clouds become the default, IT decision makers will need a mechanism not only to simplify the management of multiple public cloud services, but also to expedite the integration of these disparate cloud services as well as to automate data movement across the resulting multi cloud ecosystem. As Duplessie mentions in his video, “In the new world order, the administrator is the data controller at the highest level—you set the policies. You use tools to enforce and ensure those policies are carried out—but YOU are in control. Or you should be. If not, you are letting Amazon make decisions you might not want made.” For multi cloud data management solutions to be the technology enabling administrators to excel in this role as the ultimate data controller, they will need to deliver several critical capabilities, including; a single architecture that consolidates multiple disparate cloud buckets, the ability to automatically select a specific cloud service and then quickly leverage the desired capabilities, and the ability to manage the co-location of compute and data as a means to reduce that latency.

A potential counterargument against the need for multi cloud data management would be to refute the need for hybrid or on-premises infrastructure at all. Why not simply use a single public cloud service provider for all IT infrastructure needs? While some IT organizations may pursue a single cloud strategy, just as some IT organizations standardize on a single on-premises vendor today, multiple factors suggest multi cloud ecosystems will become the norm. There are performance and security concerns related to moving data across the WAN, encouraging IT organizations to retain certain workloads on-premises. For off-premises solutions, concerns over becoming locked into a particular cloud vendor persist and fuel interest in a multiple cloud supplier model. Additionally, from a supply side, cloud services providers will likely continue to differentiate their offerings and pricing models as the cloud market matures. In an effort to achieve the maximum benefit from this diversified cloud landscape, IT decision makers will seek to leverage the differentiated services that meet their specific workload needs. With the right data management, cloud infrastructure administrators can take advantage of that differentiation without requiring their developers to change their applications. SDS-based technologies appear well poised to deliver this needed multi cloud data management functionality.

Achieving Infrastructure Flexibility with Software-defined Storage

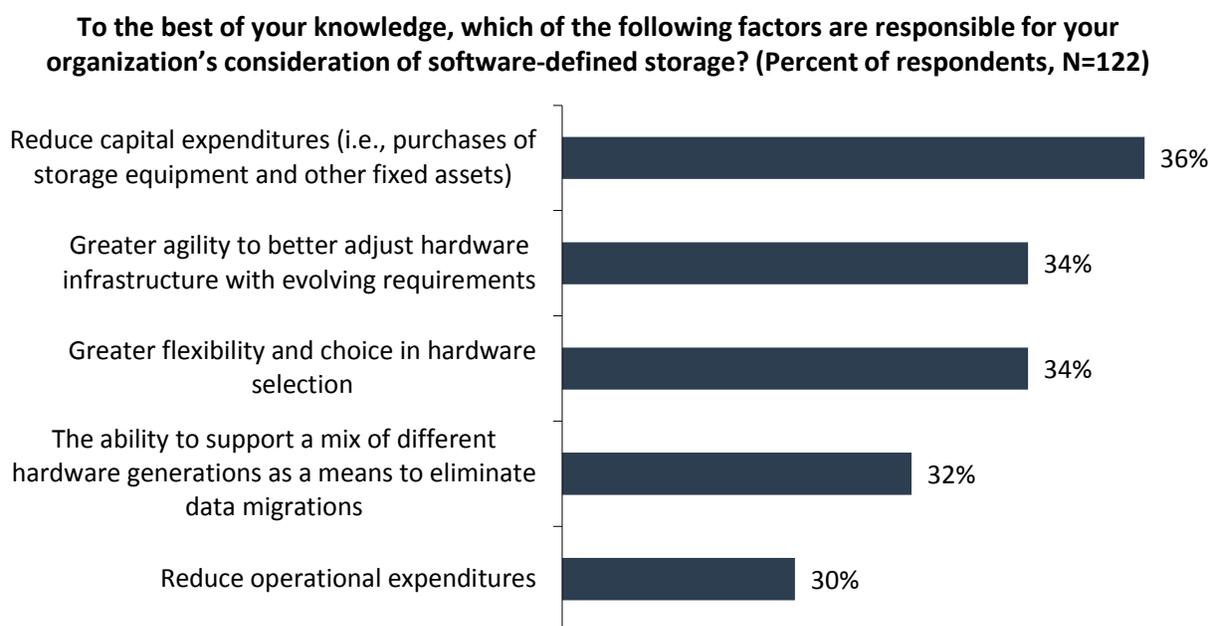
The ability to deliver storage intelligence as software has introduced new variations in storage deployments, opening opportunities while simultaneously creating some confusion when discussing the specifics of SDS. A more effective way to think of SDS is not necessarily as a new storage technology category, but rather as an alternative delivery model for enterprise storage functionality. With SDS, storage intelligence is delivered as software abstracted from the hardware, providing, among other benefits, the ability to consolidate multiple infrastructure elements into a single capacity pool and then integrate new elements as demands or technologies change.

Recently, ESG conducted a detailed research study focused on SDS market trends. When storage decision makers familiar with SDS technology were asked to identify which SDS implementation strategy they were using or expected to use, the most common response was deploying SDS technology as a virtualization layer (40%). For comparison, the alternative SDS implementation models were as SDS for hyperconverged (27%) and SDS as an alternative to a traditional storage array (26%).¹ Respondents who identified a preference for deploying SDS technology as a virtualization layer were then asked to

¹ Source: ESG Research Report, [Software-defined Storage \(SDS\) Market Trends](#), February 2017. All ESG research references and charts in this solution showcase have been taken from this research report.

identify the factors that led their organizations to consider SDS technology (see Figure 1). The four most-cited responses highlight a desire to reduce the cost of infrastructure and improve hardware deployment flexibility and agility. These include reduction of capital expenditures (36%), greater flexibility and choice in hardware selection (34%), greater agility to better adjust hardware infrastructure with evolving requirements (34%), and the ability to support a mix of different hardware generations as a means to eliminate data migrations (32%).

Figure 1. Top 5 Factors Responsible for Consideration of SDS for IT Organizations that Identified the Ability to Deploy SDS as a Virtualization Layer as Their Expected SDS Implementation Strategy



Source: Enterprise Strategy Group, 2017

These concepts of flexibility can be extended to other infrastructure elements, beyond hardware and storage systems, to include public cloud services as well. SDS technology does not, however, automatically translate into multi cloud data management. The solution must provide a number of capabilities beyond simply the ability to deliver storage intelligence as software.

Multi Cloud Data Management: A Logical Conclusion

For SDS-based architectures to deliver multi cloud data management, they must extend beyond traditional SDS definitions and deliver several key capabilities including:

- **A single control layer:** This is the most obvious element. Multi cloud data management must consolidate multiple infrastructure elements, on- and off-premises. It should also extend beyond being simply a single management pane, and deliver the ability to deploy a workload and then have the infrastructure evolve underneath while data accessibly remains unaffected. This abstraction is necessary to ensure that new infrastructure or cloud services can be added to increase resiliency, access new capabilities, or protect against being locked in to a single vendor.
- **Automatic multi cloud management:** Once the infrastructure is consolidated, administrators will need the ability to manage policies that adjust data placement and access when application or business demands change. This will require automated policy-based data placement and movement across storage tiers, data centers, and clouds. Without the necessary intelligence and automation, multi cloud infrastructures quickly will become unwieldy.
- **Universal data accessibility:** Beyond simply the ability to place data anywhere, a multi cloud data center must have the option to spin up new workloads on-premises, or within the public cloud of their choosing. This flexibility must also be

timely. New workloads cannot take months to shift from one location to another. The multi cloud infrastructure must be able to keep up with the pace of business. Improving accessibility to public cloud services, however, must not neglect on-premises infrastructure. For certain workloads, on-premises deployment may be preferred for a variety of business reasons including security, performance, and cost. In addition, some legacy, possibly low priority, workloads may be retained on-premises despite potentially benefiting from the cloud due to other, more pressing business priorities, further increasing the need to integrate on-premises infrastructure.

- **Data intelligence for future application development:** An emerging generation of cloud-enabled workloads has increased the importance of metadata, the organizing principle around how data is stored and accessed. Increased sophistication in how metadata is accessed, searched, and manipulated provides the opportunity to further maximize value from data. Already being leveraged for analytics, metadata holds the potential to eventually surpass the file-tree hierarchy and become a dominant method to locate and categorize digital information.

While multiple technologies could potentially serve as a multi cloud data management layer, unstructured storage architectures, such as object storage, however, may be best suited. Object storage technology possesses an innate ability to achieve massive scale and leverage metadata-based intelligence. These solutions, however, would target a vastly different segment than the traditional object storage systems designed for archiving data. A variety of SDS technology vendors are developing technology in or around the area of multi cloud data management. Some of the vendors poised to potentially deliver solutions in this area include Cohesity, Datera, Elastifile, Formation Data Systems, Hedvig, and SwiftStack, just to name a few. Currently, these SDS providers differ in terms of the workloads they target, as well as their focus, from predominantly on-premises models to others that target a more multi cloud architecture. All of these SDS providers, however, possess the potential to deliver the multi cloud data management the industry needs.

The Bigger Truth

Public cloud services are reshaping the IT landscape. With data quickly becoming a key determinant to business success, the future data center will need access to the best tools available to compete. Future levels of innovation and differentiation in the public cloud arena will therefore translate into the necessity of the multi cloud data center. Ultimately, businesses will need to achieve a balance between retaining key capabilities on-premises and simultaneously leveraging the best-of-breed public cloud services as a means to lower costs, distribute risk, reduce lock-in concerns, and leverage the latest in cloud capabilities. Managing each cloud separately as a disparate entity will quickly become unwieldy, limiting the benefits of the cloud. Achieving, and then maintaining, this equilibrium point will require multi cloud data management technology to help businesses keep pace with cloud innovation, and SDS-based architectures likely hold the key to this much-needed class of solutions. It is still early. With a number of technology providers already offering multi cloud data management solutions or at least having the right foundational architectures in place, these solutions have already begun to experience traction. Given the expected demand, this space could easily become the IT segment to watch over the next few years.

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.



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

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

