

# White Paper

---

## **Five Key Considerations to Ensure Data Recovery**

### *A Guide for Small and Midsized Organizations*

*By Jason Buffington, Senior Analyst; and Monya Keane, Research Analyst*

June 2015

---

This ESG White Paper was commissioned by Hewlett Packard Enterprise and is distributed under license from ESG.

## Contents

The Modernization of Production Is Driving the Modernization of Protection .....	3
Five Considerations for Midsized Organizations .....	4
What Do I Need to Protect? .....	4
What Kinds of Recoverability Should I Plan For? .....	4
How Can I Reduce My Costs? .....	6
How Long Do I Need to Keep My Data? .....	6
Where Does the Cloud Fit into Modern Data Protection Activities? .....	7
Understanding How HPE Can Address Modern Data Protection Requirements .....	7
Modern Data Protection Software from HPE .....	7
Modernizing Backup with HPE .....	8
Modern Data Retention from HPE.....	9
Modern Backup via the HPE Cloud .....	10
The Bigger Truth .....	10

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.


## The Modernization of Production Is Driving the Modernization of Protection

Businesses modernize IT production because it enables them to be more productive. But often, as they go about modernizing their production environment, they are forced to recognize that they are using inadequate approaches to protect the data in it.

This year’s “production-centric” IT modernization priorities that were most frequently cited by ESG research survey respondents (see Figure 1)<sup>1</sup> are all clearly noble, logical ways to improve the productivity of a business. But for midsized organizations especially, the collective impact of those priorities on the IT infrastructure can be significant. Thus, many SMBs are finding themselves prioritizing efforts to modernize their data protection solutions in addition to updating their production platforms and workloads.

Figure 1. Top Ten IT Priorities for 2015 for Midsized Organizations

**Top 10 most important IT priorities for midmarket organizations (100 to 999 employees) over the next 12 months. (Percent of respondents, N=233, ten responses accepted)**



Source: Enterprise Strategy Group, 2015.

The bottom line is that the modernization of different aspects of production (i.e., increasing storage capacity, virtualizing, improving collaboration, etc.) also drives the modernization of protection.

Perhaps that is why, year after year, improving data backup and recovery has been mentioned by ESG research survey respondents even more frequently than some of their production-related priorities are mentioned.

Today’s production modernization priorities are noble and logical. But their collective impact on an IT infrastructure forces an organization to prioritize modernizing its data protection solutions as well.

<sup>1</sup> Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

At face value, one would think that any production-centric IT modernization priority would be recognized as important by more respondents than just backup.

But that's not the case. Because IT groups have been modernizing so many aspects of their production infrastructures to meet business demands, improving backup has become, collectively, an even higher priority for them than it would have been otherwise.

And that situation is problematic in itself:

- Most production-modernization projects expose an unfortunate level of inefficiency in the legacy data protection solutions that are commonly in use at midsized organizations.
- Midsized organizations are just as dependent on data as larger organizations and have always needed data protection as much, too. But traditionally, they have had to settle for mediocre approaches to backup or have had to take on the costly, complicated job of trying to squeeze a large-scale data protection solution into their midsized environment.

Thankfully, today, solutions are available that provide modern, comprehensive data protection without the cost and complexity of a traditional large-scale data protection product.

## Five Considerations for Midsized Organizations

When IT decision makers at midsized organizations are assessing the modern data protection options available to them, five key questions can help them uncover the technologies that might be most appropriate. The questions center on:

- What does the organization need to protect?
- What kinds of recoverability should it plan for?
- How can it reduce its costs?
- How long does it need to keep its data?
- Which cloud method or approach is best for the organization?

### What Do I Need to Protect?

The days of a small or midsized business (SMB) having a homogeneous, "simple" IT environment are long gone. A modern midsized IT organization uses a wide variety of infrastructure components and services—perhaps not as many as big companies do, but still a more diverse collection than what many people may presume.

Furthermore, that product diversity actually increases over time: departments and business units are continually expanding the number of IT tools they use to achieve their operational goals.

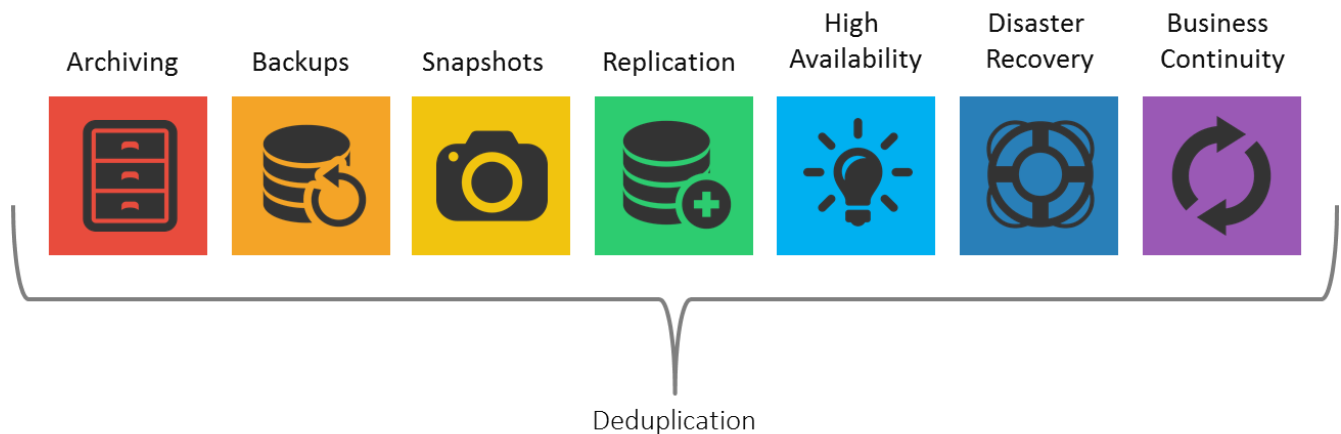
An ever-growing variety of production platforms will need to be backed up—not only the various server operating systems (including multiple releases or versions of each), but also the data that lives outside of the traditional servers, i.e., data on endpoint devices, remote offices, or potentially even in the cloud.

### What Kinds of Recoverability Should I Plan For?

In this case, the word "recoverability" could arguably be typeset as "*recover-ability*." That's because the term describes more than a straightforward data-restoration effort. "Recover-ability" describes an organization's true degree of aptitude for, and fitness in, assuring or restoring access to data properly and efficiently in general.

The kinds of "recover-ability" an SMB should plan for most likely will involve combining backups with snapshots, replication, and other protection activities, as shown in Figure 2.

Figure 2. The Data Protection Spectrum



Source: Enterprise Strategy Group, 2015.

For example,

- To preserve data for content-specific purposes, implement **archiving**.
- To recover data selectively or en masse according to a range of previous timeframes, use traditional **backups** in a deduplicated storage pool or contemporary tape or cloud solution.
- To recover to near-current points in time, use the **snapshot** technology (ies) in the primary storage.
- To ensure that data lives in more than one location, **replicate** it.
- If certain data must *always* be accessible, employ **high-availability** mechanisms on production servers and resilient storage.
- If data needs to be recoverable from a remote office or other physically separate location, leverage the replicated copies for **business continuity/disaster recovery** (BC/DR). DR relates to the survivability of data, and BC takes that idea to its ideal end by ensuring that the business never suffers an outage—regardless of the scope of the crisis.

### **Take End-users' Downtime Intolerance Seriously**

Regardless of an organization's size, its workers tend to have extremely high expectations about the availability of the applications and systems they depend on. To them, even a few hours of downtime are usually intolerable. In fact, ESG research shows that 62% of respondents have a downtime tolerance of three hours or less related to normal applications and 83% reported the same downtime tolerance for high-priority applications.<sup>2</sup>

With such a high dependence on IT and such a low tolerance for downtime and data inaccessibility, traditional backup alone is insufficient. It's unrealistic to assume that a typical IT-generalist admin at an SMB can become aware of, diagnose, restore, and then restart a production application in three hours or less using only traditional backup/restore methods.

### **Establishing SLAs Should Be a Partnership Effort Between Business Leaders and IT Implementers**

Decision makers overseeing modern IT environments should think about combining backups (which are still absolutely necessary) with snapshots and replication to provide a broader, stronger degree of recoverability. Often, which method of data protection to use (i.e., which color of the spectrum shown in Figure 2 to use) depends on

<sup>2</sup> Source: ESG Research Report, [Data Protection-as-a-service \(DPaaS\) Trends](#), September 2013.

understanding the recovery point objective (RPO) and recovery time objective (RTO). Ask, “How long can we be down?” and “How much data can we afford to lose?”

Although “backup” is addressed in this paper, a broader range of snapshotting and replication may be necessary to achieve the business’s resiliency goals—and managing those complementary approaches to data protection through a single framework or through components acquired through a single vendor can have both operational and agility-related benefits.

## How Can I Reduce My Costs?

As Figure 2 illustrates, IT teams should be performing a range of data protection activities to give the business the resiliency it requires. Of course, they also need to reduce the environment’s overall storage and data protection cost and complexity—including capital expenses related to the data protection infrastructure. One economical way to protect everything that needs protecting is to deduplicate data. In this manner, only unique data and versions of data are stored over time. Deduplication eliminates duplicate copies of repeating data, yielding significant savings by reducing storage capacity consumption and thereby changing backup’s economics.

### *Take Advantage of Today’s Advances in Deduplication*

With so many major IT challenges arising due to storage growth, deduplication has never been more of a necessity. It provides organizations with substantial efficiency advantages when deployed as a part of the data protection infrastructure. Although deduplication does not stop the growth of protected data entirely, it does provide an effective strategy for reining in the size of backups.

Deduplication technologies continue to evolve as vendors develop progressively better ways to transport and store data efficiently—the algorithms that analyze the stored data are constantly improving. That’s fortunate because even a fraction of a percent of improvement in a deduplication ratio can lead to gigabytes of capacity saved.

In addition to algorithm-related improvements, the times and places in which deduplication occurs also are evolving. Deduplication is no longer centered on simple enhancements to purpose-built backup appliances; it now works across a network—within the data path and across multiple devices.

## How Long Do I Need to Keep My Data?

The short (albeit unsatisfying) answer to that question is “it depends.” When deciding how much data should be stored for how long, remember that two forces are at work:

- Regulatory compliance often requires the retention of data for periods longer than one might presume, and company size doesn’t matter. Compliance is not a “theoretical” activity that only big organizations have to deal with (or are *able* to deal with). Compliance involves following data preservation or destruction mandates based on such things as industry or geo-political boundaries rather than company size. For instance, in the U.S., any publicly held firm—whether it has 40,000 employees or 400—must adhere to the same Sarbanes-Oxley rules for retaining and destroying data. Similarly, all U.S. medical practices and offices, regardless of size, have to adhere to HIPAA patient privacy and records retention rules.

As a term, “regulatory compliance” may seem daunting or reflective of a complex effort. In actuality, compliance is really just about retaining the right data for the right reasons. And most regulations do provide actionable, specific guidance.

- Even if a company’s industry or country does not require data retention or destruction to meet regulations, good business practice dictates keeping data for a significant period of time. In considering the earlier question of managing cost, long-term data retention requirements will likely be better served by retaining the data on economical (but modern) tape, or protecting it in the cloud.

## Where Does the Cloud Fit into Modern Data Protection Activities?

It is impossible today to have an IT modernization conversation that does not include the cloud in some form. And for many people, that fact leads to frustration and confusion instead of better IT-modernization insight.

Essentially, *cloud-based data protection solutions are best suited to enhance one's on-premises data protection and data management strategies*. And they can take multiple forms:

- **Endpoint backups:** Today's modern endpoints—whether they are corporate issued or employee owned devices—are natively connected to the Internet and increasingly less often tethered to intranets and similar corporate IT networks. Thus, backing up those endpoints to a cloud-based service is an obvious solution that many mid-sized organizations should consider as part of a broader data protection strategy.
- **Collaboration enablement:** Knowledge workers continue to leverage cloud-based file sharing and synchronization to increase their personal productivity and enhance teamwork/collaboration. File sync-and-share solutions are invariably cloud based, and in that context, it is important to remember that the data is still corporate data and should be protected by every company—even small and mid-sized organizations.
- **Disaster recovery preparedness:** The cloud opens up exciting opportunities for business continuity and disaster recovery (BC/DR) at mid-sized organizations. As Figure 1 shows, improving backup and recovery is a top-ten priority, but so is BC/DR. Thankfully, an economical secondary site (i.e., the cloud) is available.

## Understanding How HPE Can Address Modern Data Protection Requirements

It is important that mid-sized organizations consider that they may benefit greatly from a *portfolio set* of technologies including:

- Modern data protection software.
- Efficient, purpose-built backup appliances with integrated deduplication.
- Options for long-term data retention, including tape and/or cloud.

One vendor that SMBs should take a closer look at in the context of data protection capabilities and requirements is [Hewlett Packard Enterprise](#). HPE has an entire portfolio of offerings encapsulated under what it refers to as BURA (backup, recovery, and archive) designed to meet the needs of SMBs.


### Modern Data Protection Software from HPE

Two HPE products address modern data protection: HPE Data Protector software is for larger enterprises, while **HPE StoreEasy plus LiveVault (from j2 Global)** was designed for SMBs.

- *As a standalone product*, **HPE StoreEasy** is actually not software but a “designed-for-SMBs” file-sharing NAS appliance acting as a cornerstone for production and protection storage alike. But when used in conjunction with HPE's **TurboRestore Appliance (TRA)** software, the StoreEasy appliance becomes a disk-based repository for data protection solutions.
- **LiveVault software (from j2 Global)** provides protection directly to the HPE cloud data centers for physical environments, virtual environments, and endpoint devices. LiveVault is able to use the HPE StoreEasy appliance with TRA as its local target for fast recovery before replicating data to an HPE cloud. The combination of StoreEasy with TRA *plus* LiveVault gives SMBs reliable protection and fast recovery from the StoreEasy appliance as well as extended data retention and survivability from the LiveVault cloud service.

StoreEasy and LiveVault are both products that have been around for a while. HPE proposes that SMB customers think of StoreEasy plus LiveVault as a single turnkey solution (see Figure 3) to keep their businesses running at a reasonable cost.

Figure 3. Hybrid Cloud Backup with HPE StoreEasy Plus LiveVault



Source: HPE, 2015.

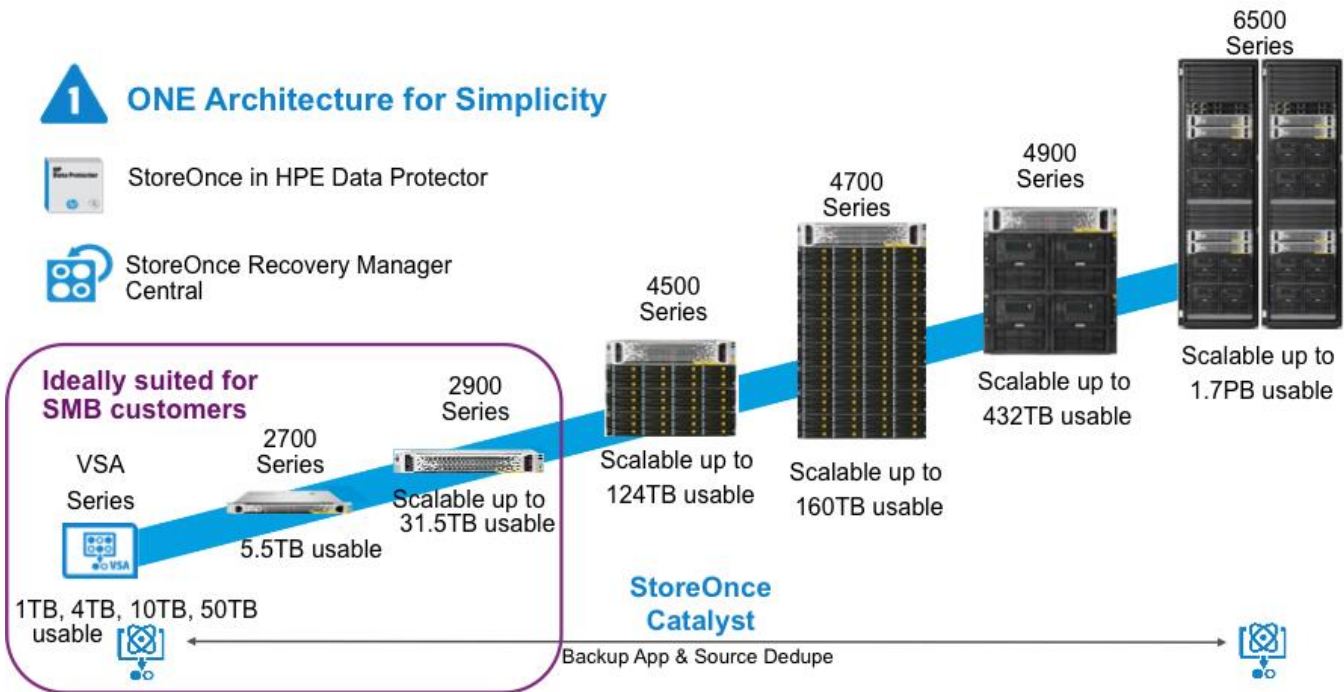
### Modernizing Backup with HPE

HPE offers two varieties of storage hardware that complement a modern data protection infrastructure. The products include the StoreEasy NAS appliance (described above) and the **HPE StoreOnce Backup system**, which is specifically designed to optimally store backups and archival data.

As stated, it is imperative to retain a variety of copies of data at multiple points in time in order to satisfy the business’s requirements for agility. But doing so can be cost prohibitive without some level of highly efficient deduplication to remove redundancies and reduce CapEx and OpEx.

To solve that problem, HPE offers the StoreOnce family of backup storage (see Figure 4). These solutions include virtual storage appliances (VSAs), small appliances for midsized organizations and regional offices, and enterprise-scale protection storage. The StoreOnce systems share a single architecture, allowing deduplicated data to move between devices in the most optimized state possible. HPE guarantees its customers a 95% reduction in backup capacity through its [Get Protected Guarantee Program](#).<sup>3</sup>

Figure 4. Single Backup Platform: HPE StoreOnce Portfolio



Source: HPE, 2015.

<sup>3</sup> HPE legal disclaimer: “As compared to a fully hydrated backup. Subject to customer qualification and compliance with the Get Protected Guarantee Terms and Conditions, which will be provided to you by your HPE Sales or Channel Partner representative.”



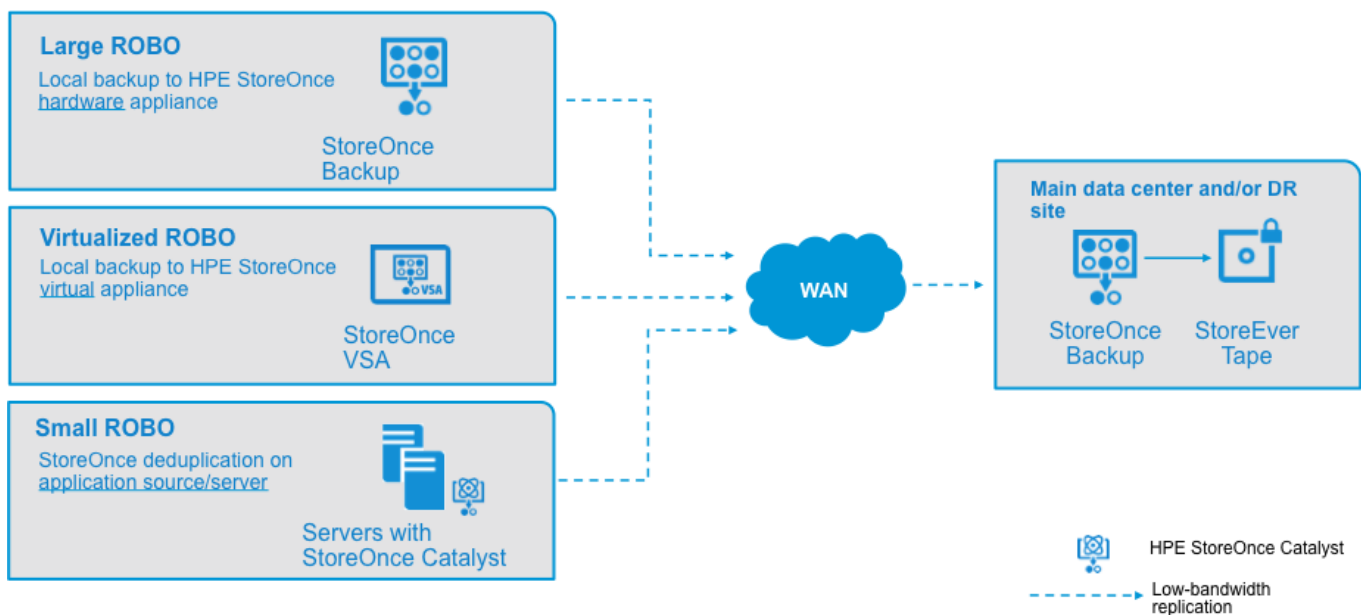
## StoreOnce Catalyst

A StoreOnce component of particular note is HPE StoreOnce Catalyst—a backup- and recovery-optimized interface with performance and manageability benefits over network-attached storage (NAS) and virtual tape library (VTL) backup. It is a key component in HPE’s federated deduplication architecture. It leverages a common deduplication algorithm across the enterprise and allows deduplication anywhere, rather than just at the points in the network where specific vendor technologies allow it.

StoreOnce Catalyst provides a single technology that can be used in multiple locations on the network without requiring rehydration when data is transferred between source server, backup device, and target appliance. It not only further optimizes deduplication, but also improves data transmission by extending its advanced optimization technologies for use on either production servers or within the backup servers—even before the data is stored in the deduplication appliance itself.

Figure 5 shows where Catalyst operates in each scenario as data moves from a remote office or branch office (ROBO) back to the data center.

Figure 5. Backup Deduplication Solutions for SMB Customers



Source: Enterprise Strategy Group, 2015.

HPE also provides Federated Catalyst for deduplication across an entire environment. This software enables an SMB to manage several StoreOnce appliances as one virtual store of data. Small sites, mid-sized organizations, and broadly distributed environments can optimize data protection and retention through this Catalyst software.

## Modern Data Retention from HPE

Disk-based data protection solutions provide agility and optimization that meets the demands of modern organizations. But any SMB that needs to retain data for three, five, seven, or ten years should consider tape as an economical and reliable form of long-term preservation. With that in mind, HPE customers should be pleased to discover the **HPE StoreEver** line of tape drives, autoloaders, and tape libraries—each powered by modern LTO-formatted drives and cartridges.

Many outdated concerns about the longevity or speed of tape have persisted for decades. However, those concerns simply do not apply to modern LTO, which ensures long retention, high speed, and durable media for organizations of all sizes. Another benefit of tape is the ability to move the tapes offsite for economic disaster recovery or data retention purposes.

## Modern Backup via the HPE Cloud

Cloud-based solutions can take many forms, even just within the realm of data protection. HPE has multiple cloud-centric solutions that enable a modern organization to leverage the cloud as part of a holistic, hybrid data protection strategy:

- For protecting endpoint devices, HPE offers **HPE Connected**.
- For protecting servers and whole sites as part of disaster preparedness, **LiveVault (from j2 Global)** is a perfect complement, which can be coupled with **HPE StoreEasy** to provide a vaulted secondary location, if a primary site experiences a catastrophic failure.
- For leveraging infrastructure-as-a-service (IaaS) for a hybrid IT experience, HPE offers **HPE Helion Cloud**.

## The Bigger Truth

Modern SMBs need the same kinds of recoverability and reliability as their larger counterparts do ... but without the complexity or cost that enterprise solutions are often marked by. Furthermore, SMBs' previously unmet needs have been exacerbated even more of late, as they embrace modern production-related enhancements that end up revealing that their legacy backup solutions are inadequate.

But this isn't about backup alone. Data protection for a modern SMB requires a combination of backups, plus snapshots, plus replication—it's all in an effort to quantify the RPO/RTO requirements that the business units' leaders have helped establish and then meet those SLAs through the right data protection and resiliency technologies in response to their end-users' extreme dependence on IT services.

Midsized organizations don't necessarily need to try and cobble together various data protection point products from a range of different IT vendors. The products on their own may be complicated to deploy, and trying to integrate them as a collection would be prohibitively difficult.<sup>4</sup>

Instead, midsized organizations that wish to focus on their business instead of IT should consider using a portfolio *that is designed for organizations of all sizes but sized appropriately for midsized organizations*, and is built to protect and recover data across a variety of workloads, media, and recovery-agility scenarios. One company to consider as a provider of this entire data protection solution is HPE.

---

<sup>4</sup> HPE claims compatibility with all leading backup ISVs and points to its compatibility matrix: <http://www8.hp.com/us/en/products/data-storage/data-storage-solutions.html?compURI=1226240#tab=TAB4>

HPE Document No. 4AA5-8754ENW



Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | [www.esg-global.com](http://www.esg-global.com)