White Paper

Exploring How HP Delivers Data Protection for Modern Infrastructures

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

September 2015

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

Contents

Introduction

Business transformation—altering how "business gets done" in reaction to marketplace shifts—drives IT modernization. In turn, IT modernization forces modernization in data protection approaches and solutions.

This fact is supported by research. According to the ESG *2015 IT Spending Intentions Survey*, improving backup and recovery was this year's second most-cited priority (along with managing data growth) among IT managers at companies of all sizes (see Figure 1).[1]

In fact, data backup has appeared among the top-three most-cited priorities for IT spending at enterprises and midmarket organizations alike for the past five years in a row.

Figure 1. Top Ten IT Priorities for 2015

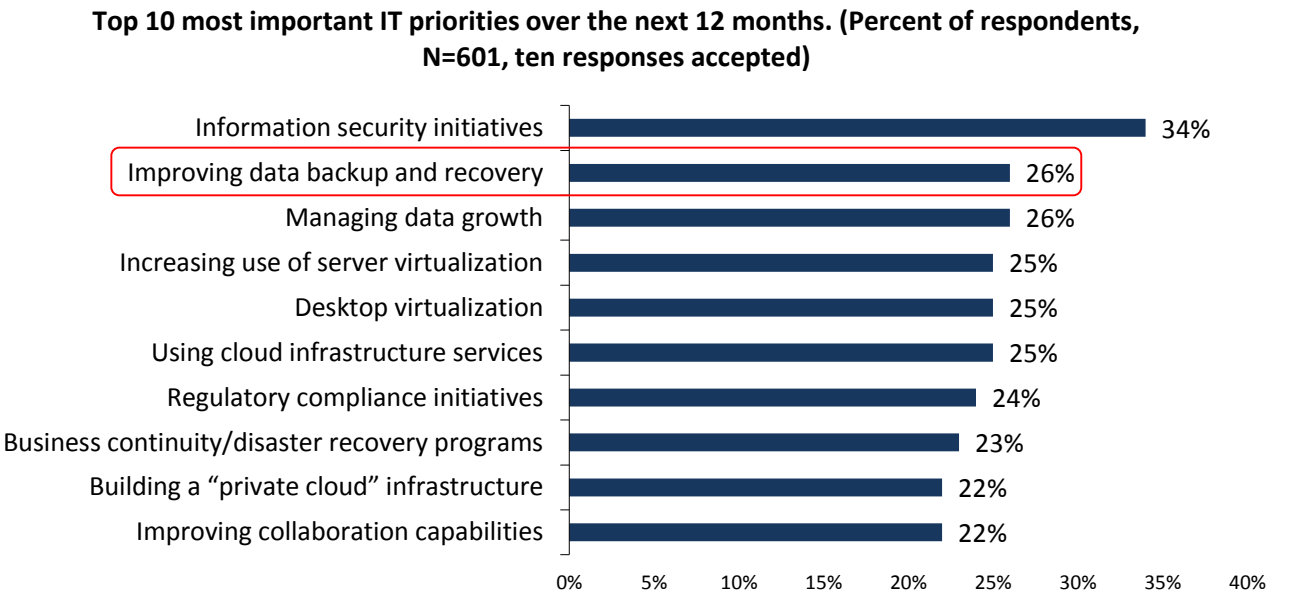

Source: Enterprise Strategy Group, 2015.

Of course, these days, when people refer to data protection, they aren't speaking of backup and recovery exclusively; they're speaking of a comprehensive approach to protecting, preserving, and managing the production data being generated by a modern and diverse IT infrastructure. And that IT infrastructure is itself distinguished by the use of scalable storage, virtualization, private/public cloud, pervasive heterogeneity, and of course, constant collaboration among end-users who demand continuous access to their data. In fact, every effort to modernize production infrastructures from new platforms (e.g., SaaS or hybrid-IaaS) to new optimizations (e.g., flash storage) points to an ever-increasing reliance on IT systems and services, with an ever-decreasing tolerance to lost data or downtime.

A Modern IT Infrastructure Demands Modern Protection

One of today's most universally applicable (yet disruptive) IT transformations relates to the adoption of server virtualization. Seemingly everywhere, organizations have been implementing virtual machines to improve consolidation at the hardware level and make their IT efforts more agile in general.

A Highly Virtualized Environment Will Undoubtedly Reveal Antiquated Data Protection Methods

Nowhere do we see the "IT agility-boosting" phenomenon unfolding more clearly than inside organizations that embrace virtualization mechanisms as a way to truly transform their IT infrastructures. It is one thing to look at

[1] Source: ESG Research Report, *2015 IT Spending Intentions Survey*, February 2015.

server virtualization as simply a means of provisioning new servers more effectively; it's quite another matter to understand how the convergence of compute, network, and storage resources within preconfigured platforms such as converged or hyperconverged systems can truly transform an IT architecture. As an example, Figure 2 depicts the broad range of approaches to virtualization available as an organization matures, from 20% virtualized (using any myriad of host platforms) through racks of standalone but consistent hosts, to blade systems to a truly converged or hyperconverged platform.

Figure 2. Virtualization Maturation—Part of the Journey to Hyperconvergence

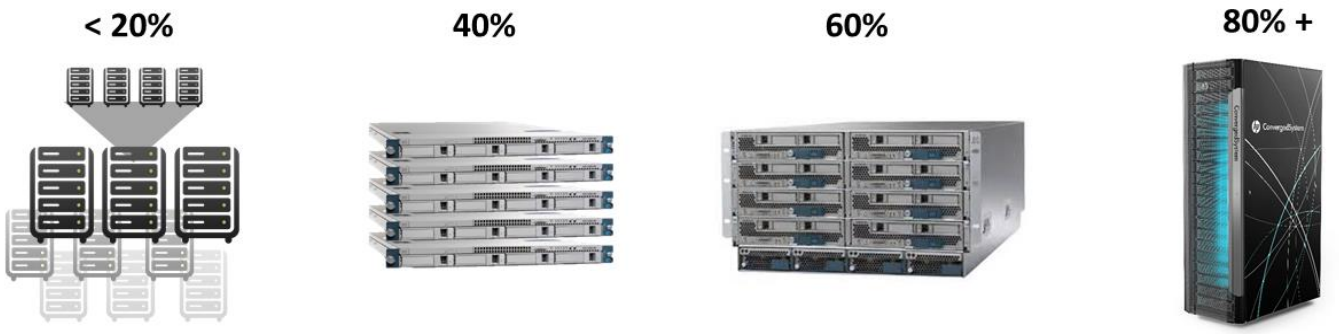

Percentage numbers reflect examples of the infrastructure used as an organization increases its virtualization density.

Source: Enterprise Strategy Group, 2015.

Systems that are converged (i.e., racks) or hyperconverged (i.e., wholly contained appliances) are the epitome of how IT transformation can happen through virtualization, providing easier infrastructure building blocks and enhanced management that comes through a single platform approach.

Unfortunately, as VM density increases, data protection-related insufficiencies can result in less reliable backups/restores and even hinder the *production* environment. Old approaches to data protection (such as simply installing traditional backup agents in virtual servers, similar to physical servers) cannot deliver the reliability, recoverability, or agility that organizations with highly virtualized environments require. For example, trying to protect a new infrastructure in an old-fashioned way can bring about:

- **Penalties—**Legacy approaches to backup can cause CPU and I/O penalties when the backup processes consume so much of a server's resources. Extra resource consumption is usually acceptable in a dedicated physical server. But the same penalties can severely impede a highly virtualized server, with the negative impacts extending not only to the VM being protected, but also to its neighboring VMs and the underlying host.
- **Magnified outage-related effects—**A single non-virtualized physical server in the data center can be down for an hour, and that's probably okay. A different server could be down for four hours, and that still might be okay. But when an organization puts 20 VMs into one physical host, *and then that host goes down*, there's nothing okay about it. Even if all the VMs are individually low priority (supporting non-critical functions or test/dev copies), at least some end-users still rely on them. And cumulatively speaking, all those "unimportant" VMs being down at once can have an effect that is arguably as dire as the outage of a truly mission-critical server.

And let's not forget the cost factor. Although embracing virtualization ultimately pays great dividends, it comes with a price tag. IT modernization requires underlying production investments, and it brings management and protection-related costs.

Downtime Is Unacceptable

Because of the high density of virtual machines in today's servers, data protection has to be less obtrusive than it has ever been. A modern approach to VM-based backup and recovery is required: Modern data protection mechanisms have to be able to recover everything from individual files and datasets to entire virtual servers and entire physical hosts—all in a highly agile way.

Backup is the cornerstone of a data protection strategy, but that strategy also should include other measures and methods if it's going to meet the resiliency demands of the business units. In particular, ensuring data availability is important. A 2015 ESG survey revealed that:[2]

- About half (51%) of tier-1 applications have a downtime tolerance level of **less than one hour**.
- Roughly seven out of eight (86%) tier-1 apps are marked by a downtime tolerance of **four hours or less**.

The numbers above seem reasonable considering that they are for "tier-1" applications, but "normal" applications are also intolerant of interruption.

- One out of four (28%) normal applications have a downtime tolerance of **less than one hour**.
- Five out of eight (62%) normal apps have a downtime tolerance of **four hours or less**.

The numbers paint a clear picture of how dependent we have become on having reliable access to our digital data.

Considerations for Protecting a 'Modern' Environment

With end-users' intolerance of downtime in mind, it is easy to understand why nearly half (46%) of organizations using a modern hypervisor—including VMware or Hyper-V—leverage snapshots or other storage-centric protection mechanisms as part of their broader data protection strategy.[3] In general, a multifaceted approach to data protection (all part of the effort to ensure agile recovery and resiliency) has been driving organizations to adopt a "data protection spectrum"-oriented approach (see Figure 3).

Figure 3. The Spectrum of Data Protection

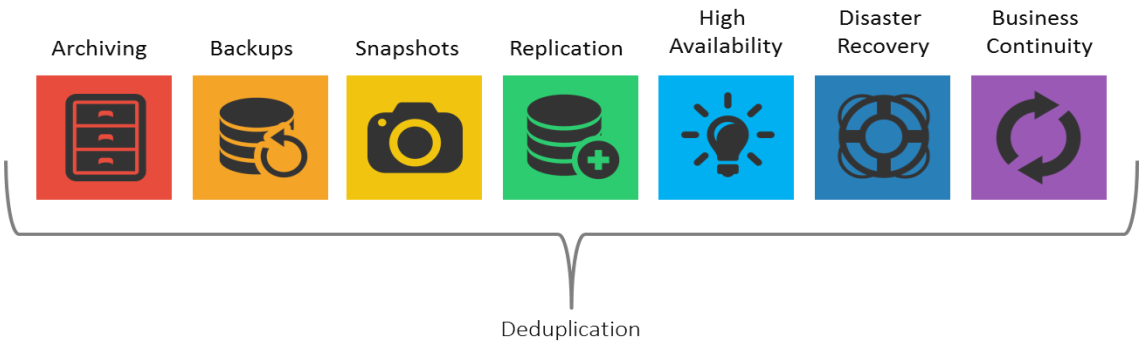

Source: Enterprise Strategy Group, 2015.

When considering which "color" of the spectrum to use, organizations of all sizes should:

- Plan on **archiving** stagnant data off high-performance production storage (particularly in cases where primary storage is becoming flash-based). Be sure you are preserving data that needs long-term retention and simply removing the rest, per policy, as its business value wanes.

[2] Source: ESG Research Report, *The Evolving Business Continuity and Disaster Recovery Landscape*, to be published.
[3] Source: ESG Research Report, *Data Protection Personas and Methods*, February 2015.

- Remember that although **backup** will continue to be the core of almost every data protection strategy, **snapshots** are almost always more effective for rapid recovery, and **replication** will ensure that the data is agile (or survivable) when you need it.
- Combine the data protection methods listed above with **availability** or resiliency technologies to provide the foundation for IT durability, which (when used with proper planning and processes) results in more effective **Business Continuity/Disaster Recovery** (BC/DR) preparedness.

The goal is to align the organization's data and usability requirements with the right IT mechanisms to deliver them—without failing on all accounts due to lack of integration, too much operational complexity, or low economic viability.

To achieve that goal and address their business demands, many organizations look for a single vendor that offers a complete portfolio of solutions or services across the data protection spectrum.

A Closer Look at How HP Is Addressing Modern Data Protection Demands

One vendor to consider as a supplier of this comprehensive approach is Hewlett-Packard—notable for its breadth of data protection technologies (see Figure 4) that can address varied requirements for protection, preservation, and resiliency. The technologies encompass:

- Snapshots and replication within primary storage systems.
- A range of modern data protection software approaches and products for backup, archiving, and replication.
- Deduplication as the underpinning of HP's secondary storage.
- Instant recovery of mission-critical applications, restore/power-on, and migrate for rapid recovery.
- Flexible tertiary retention mechanisms.

Figure 4. How ESG's Spectrum of Data Protection Can Be Mapped to HP Offerings

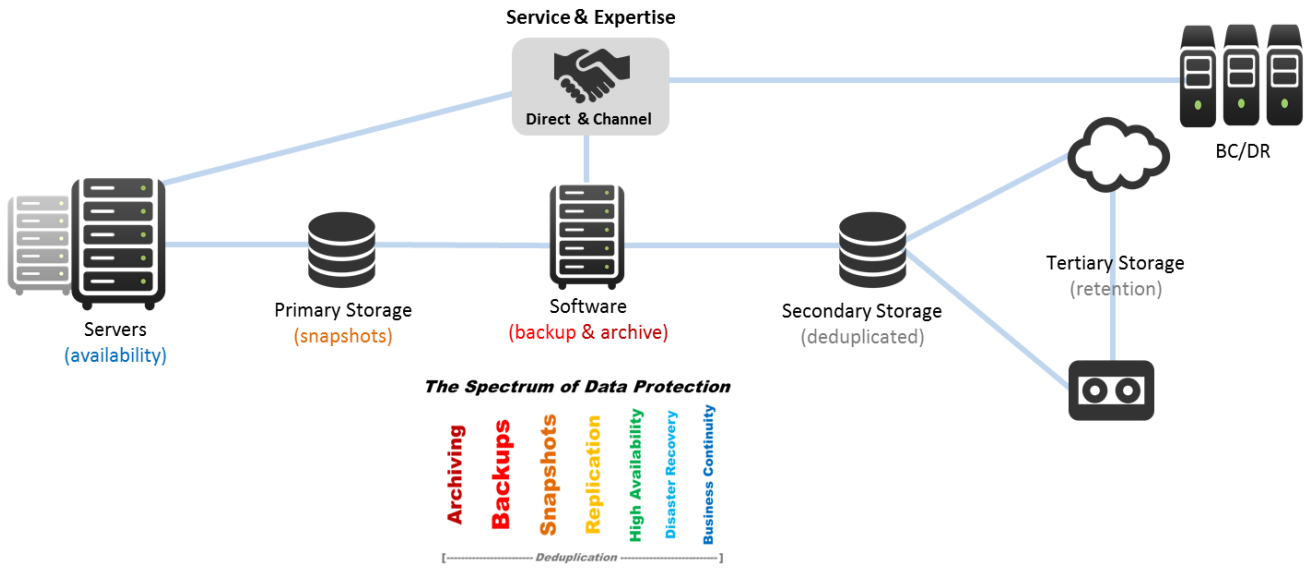

Source: Enterprise Strategy Group, 2015.

It is worth noting that HP has offerings at each point along the data protection solution architecture shown in Figure 4, as two other ESG papers (see sidebar) discuss.

For more ESG insights on HP's offerings across the data protection spectrum, read *'Multiple Data Protection Solutions' Does Not Have to Mean 'Multiple Vendors.'*

To learn more about why ESG finds HP StoreOnce deduplication technology to be compelling, read *Four Data Protection Imperatives for 2015 and HP's Answers to Each.*

Deduplication as the Underpinning of HP's Secondary Storage

HP StoreOnce is a family of data protection technologies (see Figure 5) composed of deduplication storage solutions including virtual storage appliances (VSAs), small appliances for midsized organizations and regional offices of large companies, and truly enterprise-class platforms offering up to 1.7PB of usable capacity.

Noteworthy is the single architecture across the StoreOnce family, which is present even in the VSA and in the deduplication logic within the HP Data Protector software. The single architecture enables deduplicated data to move between devices in the most optimized state possible.

Figure 5. The HP StoreOnce Product Line

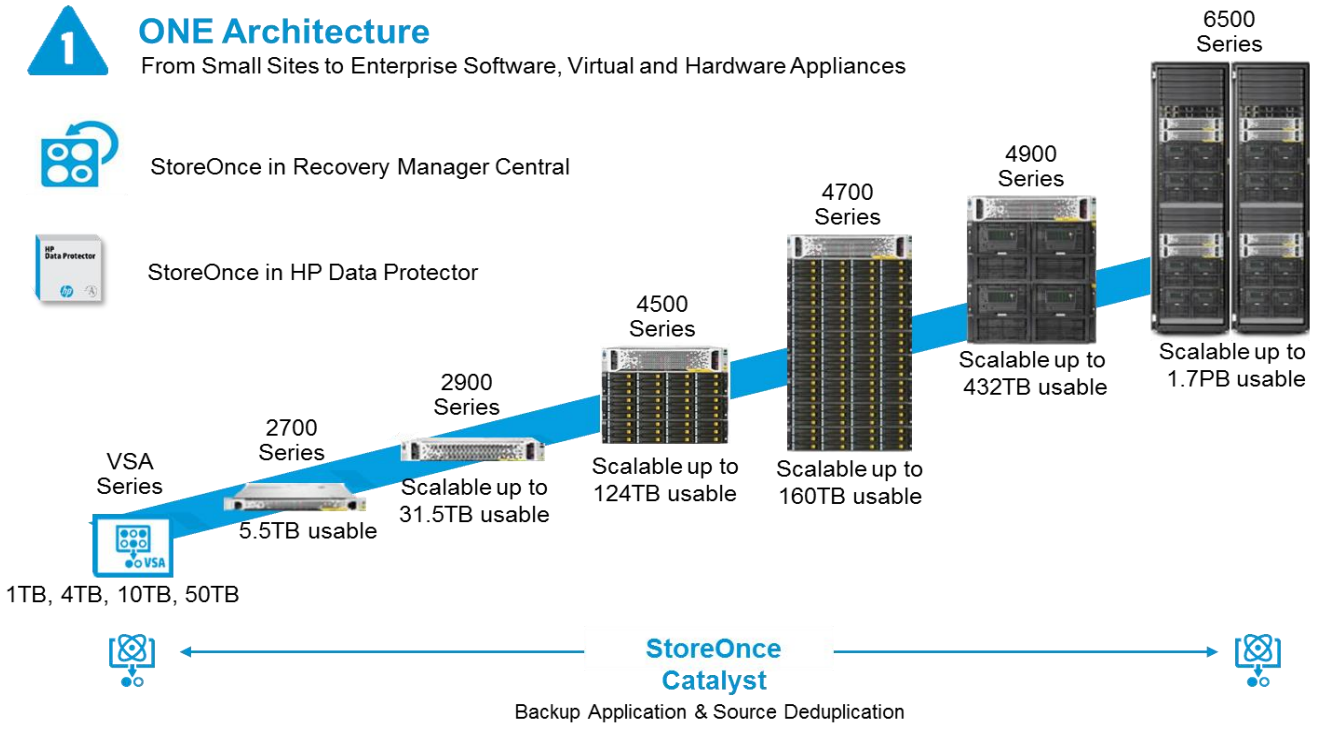

Source: HP, 2015.

Plan on a Range of Modern Data Protection, Resiliency, and Preservation Methods

It is important to remember that to achieve comprehensive protection and recoverability, organizations should plan to adhere to the spectrum approach to protecting data shown in Figure 3. But even a single vendor promoting a cohesive approach will likely need multiple products to achieve that IT agility organizations are looking for. *Unified strategy* and *interoperable technologies* should be key considerations.

Data Protection solution components from HP

HP offers a complete portfolio of data protection software supporting IT resiliency initiatives through various data movement, protection, and recoverability scenarios:

- **HP Data Protector**—Data Protector is backup and recovery software that delivers comprehensive data protection based on real-time intelligence about the backup and recovery processes in effect. Specifically to meet aforementioned strict availability requirements and service levels, it provides policy-based protection, real-time analytics, and advanced recovery options including rapid recovery and the ability to restore/power-on virtual machines and live migrate into production. In addition, Data Protector can serve as a single point of management and control to manage data protection locally and remotely. And with recent enhancements to replication and automation workflows, Data Protector provides additional

recovery capabilities through automated replication synchronization of backup catalog metadata and corresponding backup data (via StoreOnce Catalyst).

- **HP 3Par StoreServ Virtual Copy—**This provides rapidly recoverable snapshots within production storage.
- **HP StoreOnce Recovery Manager Central (RMC)**—RMC complements traditional backup by providing integrated convergence between primary and backup storage for a variety of purposes, including transporting incremental snapshots directly from 3PAR primary storage to HP StoreOnce protection storage and facilitating both snapshots and better backups through API integration with traditional backup software (e.g., HP Data Protector and other third-party backup offerings). ESG expects that RMC will likely continue to be an area of differentiable agility within HP's data protection solution portfolio as HP continues to invest in new scenarios such as synthetic fulls within HP StoreOnce (from the snapshots) and granular recovery of VMs from snapshot-based versions.
- **HP StoreOnce Catalyst**—StoreOnce Catalyst is an underpinning technology that makes HP deduplication work even better. It optimizes deduplication and improves data transmission on production servers and within backup servers—even before data is stored in the deduplication appliance. This single technology is usable in multiple locations on a network. It eliminates the need for data rehydration as data moves around among source servers, backup devices, and target appliances.
- **HP Backup Navigator**—HP Backup Navigator provides IT staff with an intuitive, interactive dashboard and analytical reports based on more than 75 key performance indicators related to backup and recovery operations. Using this dashboard, IT staff can immediately identify inefficiencies within backup operations, detect an unbalanced use of backup resources, and uncover failures before they affect a recovery process. By getting end-to-end insight into the physical and logical data protection infrastructure, IT staff can make faster, smarter decisions concerning backup and recovery. Quickly isolating data protection problems and uncovering their root causes gives IT staff precious time at a point when business needs center on vital information recovery.

Data Retention solution components from HP

Disk-based protection and recovery is important because IT organizations have to meet stringent SLAs. But it is equally important to recognize that, in many environments, disk-based data protection won't satisfy long-term retention or long-distance data survivability requirements. With that in mind, HP offers not only software technologies, but also modern tape and cloud solutions to provide tertiary-level preservation and retention:

- **HP Storage Optimizer—**This offers software-driven file archiving and storage usage reporting/analytics to help IT administrators preserve what data needs to be retained, while ensuring that more critical data remains on high-performance primary storage.
- **HP StoreEver**—Any organization that needs to retain data for several years should consider tape as an economical, reliable form of long-term preservation. The HP StoreEver line of tape drives/cartridges, autoloaders, and tape libraries are powered by modern LTO formatting. Modern LTO provides dependable retention, speed, and media durability. Another benefit: organizations can transport tapes to achieve economical disaster recovery or simply to provide long-term, secure retention offsite.
- **HP clouds (public, private, and hybrid)**—In recognizing organizations' various protection and retention requirements, HP has built a formidable portfolio of cloud services that can meet a range of needs:
    - **HP Helion Cloud:** HP created this service to help its enterprise customers leverage infrastructure-as-a-service (IaaS) for a hybrid IT experience. Helion Cloud is HP's newest cloud offering, and it boasts a combination of IaaS and platform-as-a-service (PaaS) in an effort to combine the benefits of a third-party public cloud and HP's service-based expertise.
    - **HP Connected:** HP's endpoint data protection product/service includes both HP Connected Backup and HP Connected MX (converged backup/recovery, synchronization, and sharing). Connected MX builds on the market strength of Connected Backup and combines managed backup and recovery

services for endpoint devices with file synchronization and sharing. It delivers enterprise assurance services that provide corporate intelligence, control, security, and analytics for information regardless of where it resides, which is, according to HP's marketing, "a smarter way to protect mobile information." Connected MX leverages HP Helion as a core part of its architecture and service delivery.

The Bigger Truth

Business units are demanding new capabilities from IT, and those demands are driving the modernization of production systems. The effort includes a rapid, industry-wide embrace of server virtualization—particularly to achieve the agility benefits that can be gained through the use of converged and hyperconverged appliances and infrastructures.

But, regardless of how a production modernization effort unfolds, data protection has to evolve with it in order to ensure that those impressive new production platforms are reliably protectable.

Value can be gained from choosing production systems that come from a single vendor (with converged and hyperconverged systems being the pinnacle). Similarly, there may also be significant value to be gained in using data protection offerings from a single vendor, particularly one that not only offers a breadth of software, deduplicated hardware, and tertiary tape/cloud options, but also can deliver those data protection solutions in integrated lockstep with the converged production systems that many organizations are coveting and deploying, such as HP.

HP Document No. 4AA6-1937ENW

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