White Paper

How Modern Backup Requirements and HPE StoreOnce Intersect

Sprawling, Evolving Environments Demand an Enterprise-grade, Centralized Solution

By Jason Buffington, Principal Analyst
And Monya Keane, Senior Research Analyst

April 2016

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

Contents

Introduction

Year after year, organizations continue to struggle to ensure that their production systems—which themselves are continuing to evolve and diversify—remain protected according to corporate standards. As a proof point, ESG survey research continues to show (as it has for the past five years) that improving data backup and recovery and managing data growth remain major IT spending priorities for organizations of all sizes (see Figure 1).[1]

FIGURE 1. Top Five IT Priorities for 2016

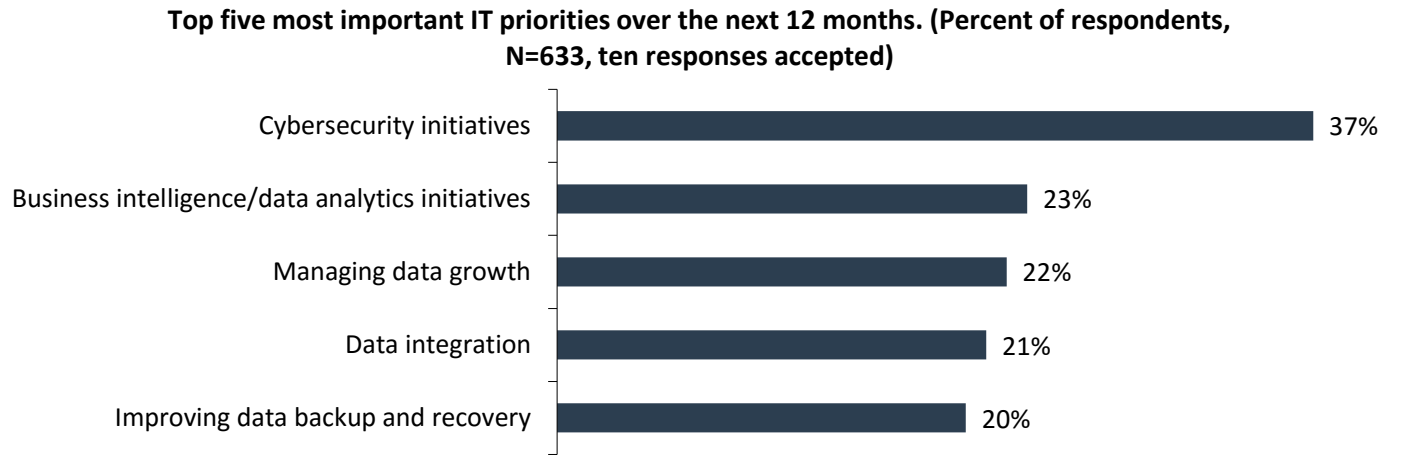

Source: Enterprise Strategy Group, 2016

It is important to note that, unlike popular sentiment or marketing might imply, "backup" is not broken. It is simply that, as production platforms evolve and production storage grows, legacy methods for data protection and recovery quickly become antiquated. Hence the ESG maxim, "When you modernize production, you must *also* modernize protection."

The Decentralization of Data Protection

Along with diversifying the data protection tools that one might leverage in support of those evolving workloads, it is important to recognize that the term "data protection" actually encompasses a variety of IT technologies, each with its own protection, preservation, and productivity-enablement scenarios (see Figure 2).

FIGURE 2. The Spectrum of Data Protection

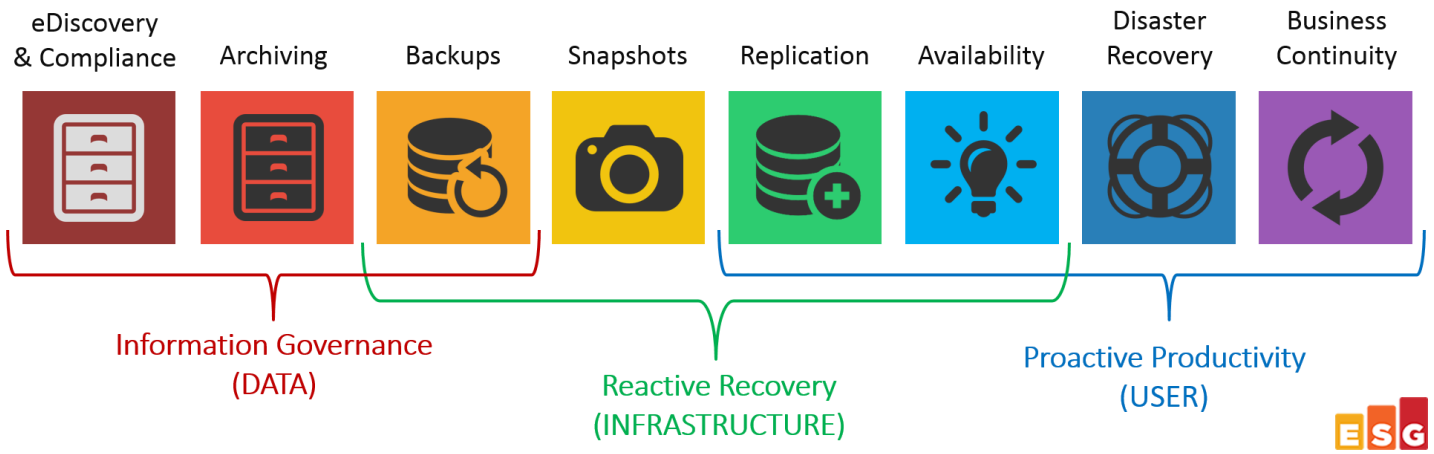

[1] Source: ESG Research Report, *2016 IT Spending Intentions Survey*, February 2016.

Source: Enterprise Strategy Group, 2016

As the image illustrates, backup and data protection are not synonymous terms. Therefore, organizations should be thinking about pursuing a more comprehensive data protection strategy, one that leverages different kinds of protection to enable various capabilities to recover—including protection encompassing backups, snapshots, replication, and archival and availability technologies.

Not Only the 'What,' but Also the 'Who'

It is important to note that along with IT organizations maturing their data protection strategies to include recovery methods other than just backup, the transformation is also being driven in part by IT professionals other than traditional backup administrators—these IT pros include DBAs, vAdmins, and file and storage admins. ESG research frequently finds that these workload and platform owners responsible for individual platforms are requesting to have influence or even demanding authority over the protection of their workloads.

Over the last decade, the industry has seen a shift away from single, general-purpose backup solutions being managed by the backup admin toward a much more diversified approach, whereby workload administrators use best-of-breed approaches to protect their individual platforms. For example:

- DBAs are using database-specific tools such as Oracle Recovery Manager (RMAN).
- vAdmins are using virtualization-specific recovery tools.
- File storage admins are protecting data directly to protection storage.

It is also worth noting that the IT evolution doesn't stop with this decentralization of responsibility across diverse workload owners. ESG also sees a third wave of evolution coming in which IT Operations are actually taking charge instead—shifting back to a unified approach to data protection management, or at least to having oversight over the diverse methods and personnel acting in data protection today (see Figure 3).[2]

FIGURE 3. A Variety of IT Professionals Are Active in Protecting Data

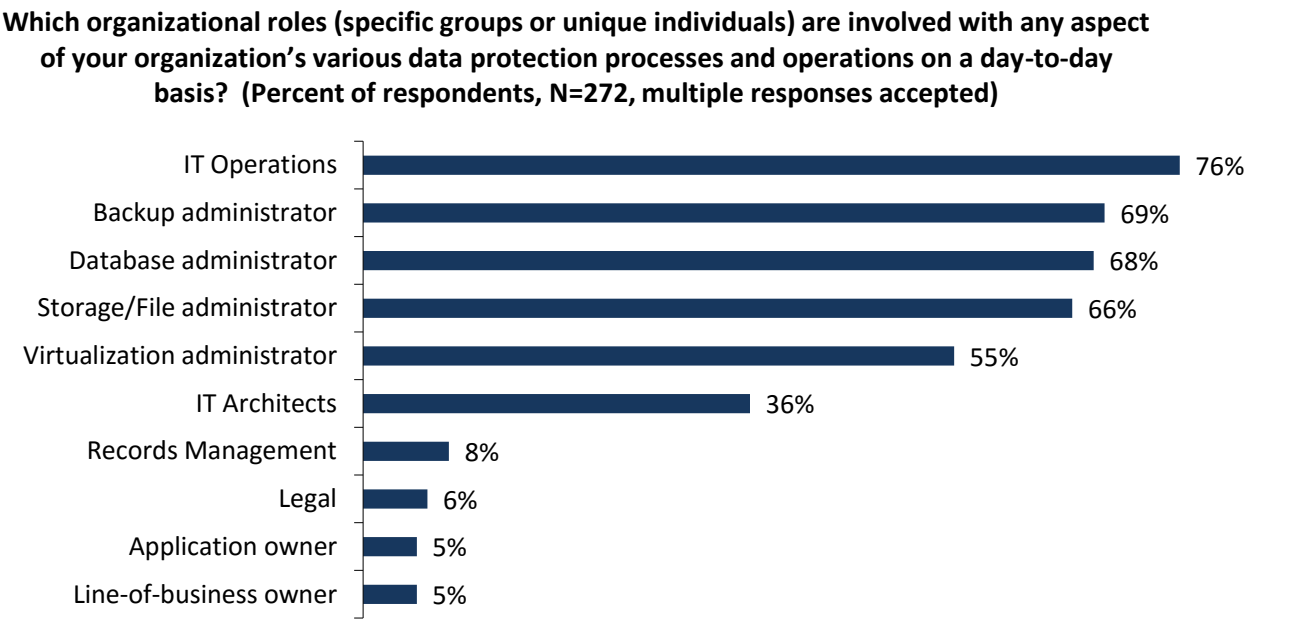

Source: Enterprise Strategy Group, 2016

[2] Source: ESG Research Report, *Data Protection Personas and Methods*, February 2015.

HPE's Approach to Protection Storage

One vendor that understands the myriad workloads and requisite data protection methods to be used is HPE, a tech giant that offers an especially comprehensive portfolio of products and services. Although HPE does have its own unified data protection software in HPE Data Protector, it also recognizes the diversity requirements that heterogeneous enterprises are working under today as they try to accomplish their protection goals. The company has, therefore, invested heavily in its StoreOnce platform for deduplicated storage. StoreOnce is enabling an ecosystem of assorted backup software solutions and tools to share centralized and optimized protection storage (see Figure 4).

FIGURE 4. The HPE StoreOnce Portfolio

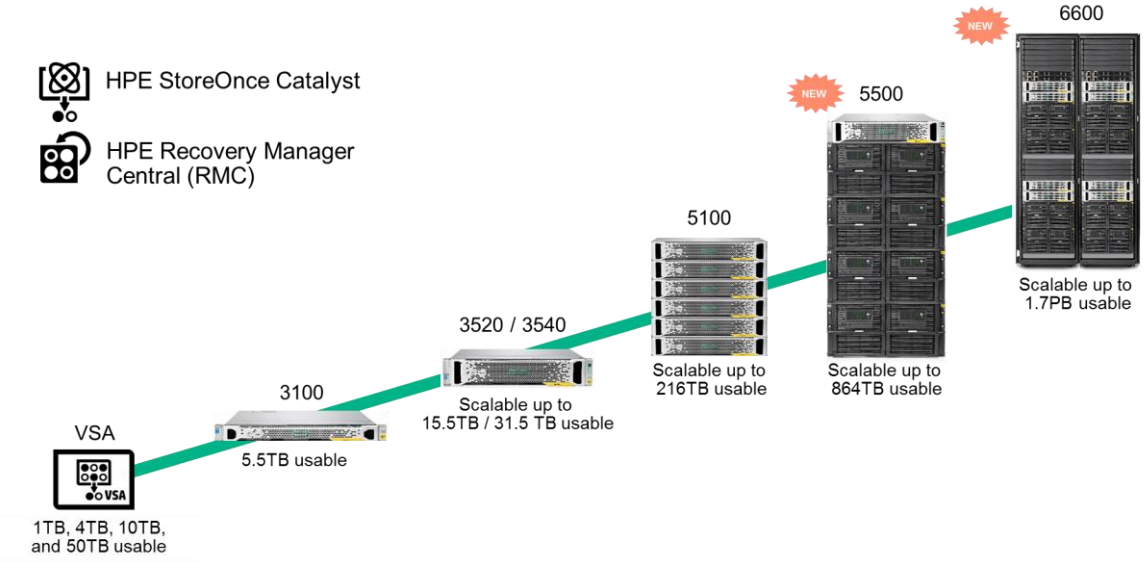

Source: HPE, 2016

StoreOnce encompasses deduplication storage solutions including virtual storage appliances (VSAs), small appliances for midmarket organizations and regional offices, and truly enterprise-class protection storage platforms of up to 1.7PB of usable capacity in HPE's StoreOnce 6600 solution. HPE has established a single architecture across the StoreOnce family, including the virtual appliance and the deduplication logic within the Data Protector software, such that deduplication data moves between devices in its most optimized state.

Notably, the HPE Catalyst product serves as the underlying deduplication technology within and across the StoreOnce Systems. Third-party backup software can leverage the Catalyst APIs to better integrate with HPE StoreOnce Systems.

HPE Recovery Manager Central (RMC)

HPE Recovery Manager Central (RMC) software integrates with HPE 3PAR StoreServ and HPE StoreOnce systems. RMC protects all applications running on an HPE 3PAR StoreServ array, enabling application-centric protection for critical applications and data sets by bypassing traditional backup-server-based processes and transmitting data directly from the 3PAR system to StoreOnce. RMC provides two snapshotting methods:

- Crash-consistent snapshots taken independently of applications (regular disk array snapshots). The data's status would be similar to what occurs in a sudden server power outage, meaning all data in host memory is gone. Most modern applications survive such occurrences by using some type of log/journal check, but they may need time to recover. All applications running on 3PAR arrays can be protected with crash consistency. All array-replication-based disaster recovery (i.e., DR via HPE XP Continuous Access software, or 3PAR Remote Copy) relies on crash-consistent recovery. With crash-consistent snapshots, all data except the data in memory is captured and saved at the same time.

- **Application-consistent snapshots taken after an application has flushed data in memory to disk and quiesced (frozen) the data on disk.** Applications start from a consistent state after failure without experiencing issues. (The app is very briefly "frozen" as the snapshot is taken.) Application consistency, ensured through integrated plug-ins, is available for VMware (RMC-V), all Microsoft apps running in VMware and using VSS, and MS SQL (RMC-S). For other applications, HPE states that customers can use crash-consistent snapshots (sufficient for many applications) or ensure application consistency by writing scripts for the applications and using them alongside the RMC APIs provided in the RMC Software Development Kit (SDK).

Multiple Platforms, One Strategy

Although different IT professionals may use different tools and methods to ensure the protection and reliable recovery of their platforms, there are obvious CapEx and OpEx benefits as well as regulatory compliance implications tied to using centralized storage that all the diverse workloads and platforms can protect their data to. The remainder of this paper explores a few of the most common scenarios and how HPE StoreOnce and/or HPE RMC are enabling them.

Databases

Historically, DBAs often supplemented traditional backup using their own database-specific tools (e.g., Oracle RMAN). In some cases, it was done to satisfy short-term recovery goals—for example, to mitigate a data corruption from recent imports or to ensure faster and more reliable recovery than what traditional backup might provide.

Unfortunately, DBAs doing so in a "rogue" style often used the only storage they had, which was usually expensive, high-performance production storage volumes. Some DBAs might have found secondary storage on other platforms instead, but it was almost always outside the purview of the IT organization's backup administrator (see Figure 5).[3]

FIGURE 5. Tools Used to Protect Databases

Which methods does your organization use to back up production databases? What is the primary way of backing up production databases? How do you expect this to change, if at all, over the next 24 months? (Percent of respondents, N=178)

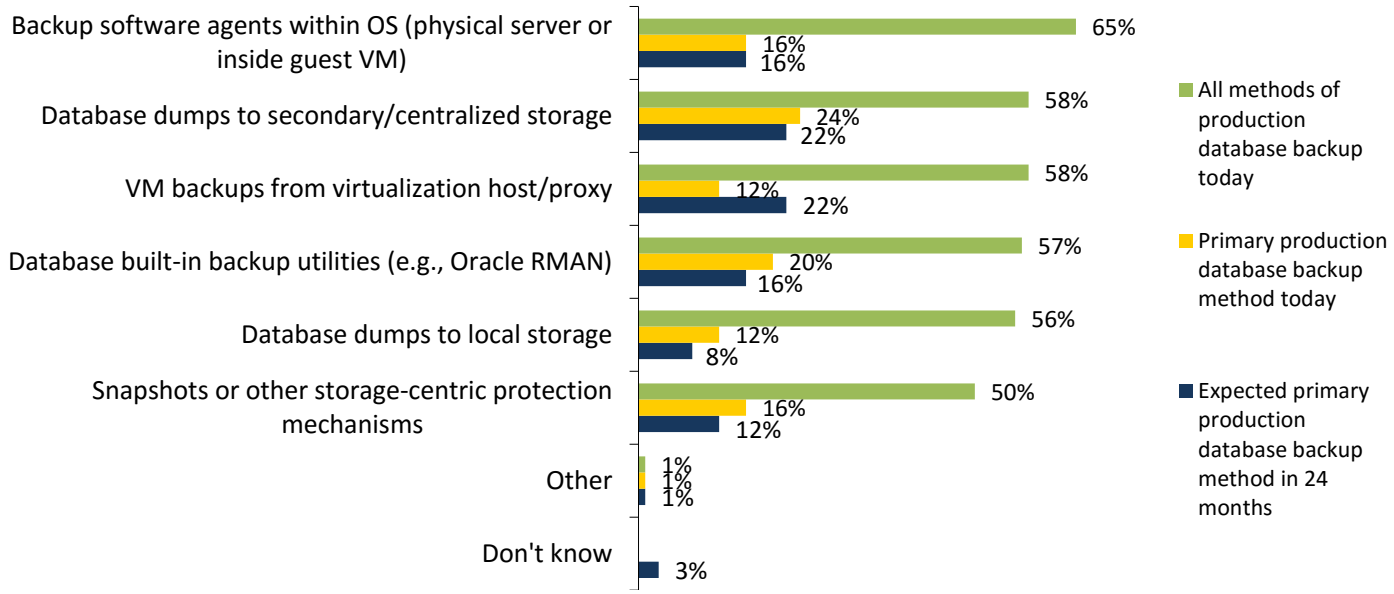

Source: Enterprise Strategy Group, 2016

[3] Source: ESG Research Report, *Data Protection Personas and Methods*, February 2015.

As a superior alternative for most of these scenarios, centralized protection storage such as that offered by HPE StoreOnce provides a more cost-effective solution to DBAs wishing to use database-specific tools. It also enables better long-term retention and IT oversight facilitated by a backup or storage admin in charge of the StoreOnce platform. Additionally, HPE has delivered StoreOnce Catalyst as an API-based accelerator, enabling Oracle RMAN, SAP HANA, Microsoft SQL, and SAP on Oracle to achieve even greater speed and interoperability with the StoreOnce Systems.

Alternatively, RMC for MS SQL (RMC-S), for VMware (RMC-V), and for any MS app running in a VM and using VSS ensures the database is momentarily "frozen" and then snapshotted in a way that the database and its related transaction logs and supporting files will be application consistent and therefore more assured to be recoverable.

Virtual Machines

For the last several years, IT organizations both large and small have struggled to protect and rapidly recover virtual machines adequately. Although API mechanisms have finally given IT the ability to provide adequate protection, the tendency to want to use a VM-specific protection tool in conjunction with traditional unified backup is becoming increasingly commonplace.

vAdmins often begin devising a data protection strategy with hypervisor-based snapshots and VM-centric backup and recovery tools—frequently as a replacement for, but sometimes as a supplement to, traditional backup models (see Figure 6).[4] And just like DBAs, these vAdmins are often using the only storage available to them, which is frequently expensive.

FIGURE 6. Protection Methods Used for Virtual Machines

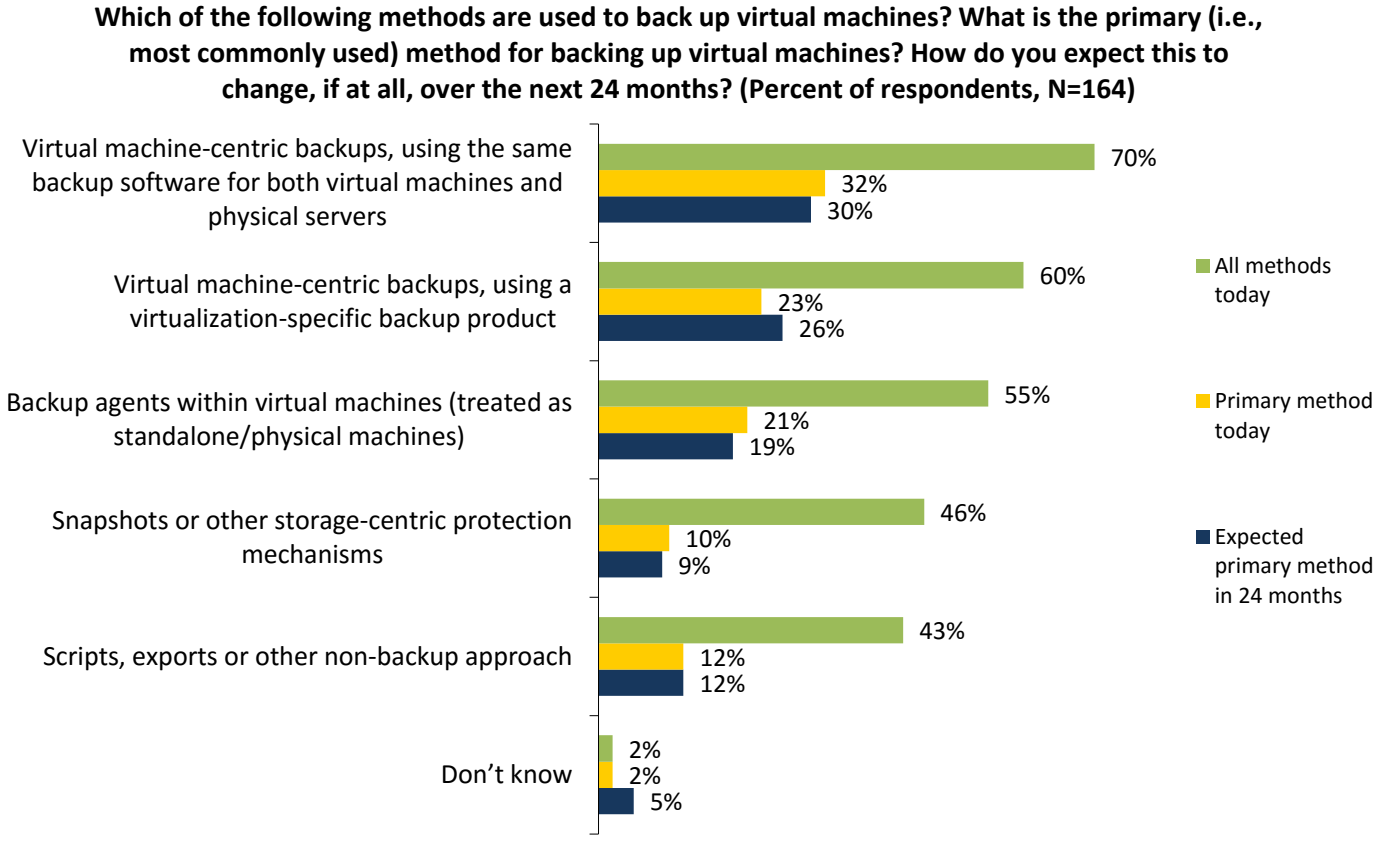

Source: Enterprise Strategy Group, 2016

[4] ibid.

APIs such as VMware vStorage VADP and Microsoft VSS continue to mature in their ability to help data protection tools protect VMs reliably. But there are often compelling reasons why admins continue to want to protect their own platforms. Some HPE customers will utilize HPE Data Protector, and some will leverage a VM-specific backup tool. All of them, however, should seriously consider taking advantage of the efficiency and performance of centralized protection storage such as HPE StoreOnce.

Recognizing how server virtualization continues to be the basis of many IT transformations, HPE provides multiple VM-specific offerings to enhance the protection and recovery scenarios that are powered by HPE StoreOnce:

- **RMC for VMware (RMC-V)** ensures that virtual machines can be momentarily "frozen," then snapshotted, and then replicated from the VMware ESXi datastores running on HPE 3PAR StoreServ directly to HPE StoreOnce systems. In addition to the VMs themselves being quiesced, any Microsoft applications running within the VMs (and therefore utilizing Microsoft VSS writers) will be assured to be application-consistent within the VMs at the time of snapshotting.
- **HPE Element Recovery Technology (ERT)** provides granular recovery within a VM, including individual files, from an RMC Express Protect backup on StoreOnce.
- HPE has also announced integration of **StoreOnce Catalyst** with Veeam (v9), a vendor many regard as the industry leader in VM-specific data protection, to further optimize the integration and performance between virtualization backup and replication technologies with HPE StoreOnce Systems.

Unstructured File Data

Of all the IT workloads whose scale challenges are often underestimated, perhaps nowhere do we see the need for improved data protection more than in unstructured file data, including home directories, shared file directories, and project folders.

With these data volumes growing at approximately 25% year over year[5] but data protection budgets only growing 4-6% year over year,[6] "status quo" approaches to protecting file data are unsustainable. Moreover, with innovations being driven by companies that create both production and protection storage (such as HPE 3PAR StoreServ and HPE StoreOnce), data protection is truly being revolutionized. In this case, HPE offers RMC software as a means by which file data is transmitted directly to HPE StoreOnce without an intermediary backup server of any kind (see Figure 7).

FIGURE 7. HPE Recovery Manager Central: Application-managed, End-to-end Availability and Protection

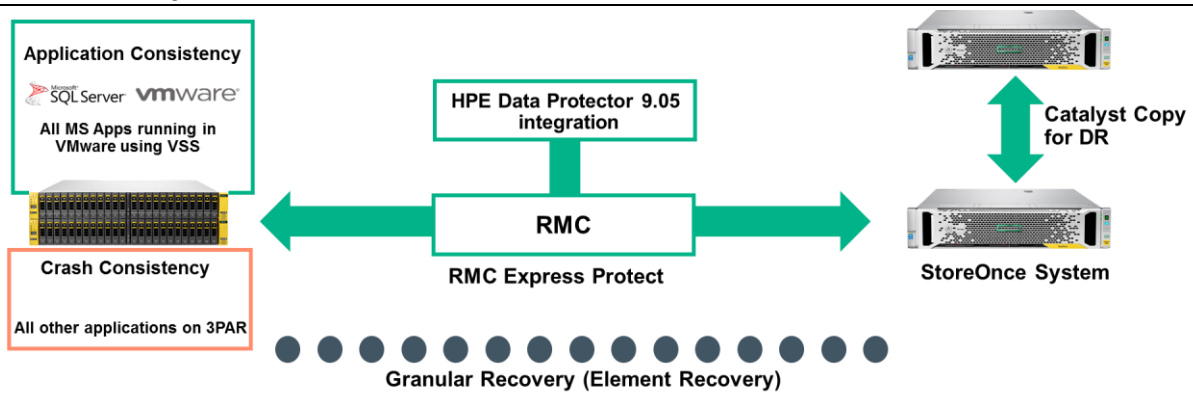

Source: HPE, 2016

[5] Source: ESG Research Report, *2015 Data Storage Market Trends*, October 2015.
[6] Source: ESG Research Report, *2015 Trends in Data Protection Modernization*, September 2015.

The method depicted in Figure 7—data transmitting directly from production to protection storage while simply updating the catalog of a backup or management server—is undoubtedly the future of data protection. As API-based approaches such as HPE's StoreOnce Catalyst and data management conduits such as HPE RMC continue to mature, it is inevitable that more workloads and data protection scenarios will mimic the approach.

Everything Else That Needs Backup

Some workloads have both specific protection and recovery requirements and dedicated personnel who are responsible for ensuring the availability of those services. However, most IT organizations will still find themselves overseeing myriad other platforms without those requirements.

It is important to note that all data should be protected to a corporate standard, regardless of who is responsible for the platform or which tools may be used to accomplish the protection. So, while some dedicated platforms may use workload-specific mechanisms, the rest of the IT environment will assuredly continue to be protected by unified data protection tools, today and in the future.

Even for the workloads listed in this paper, personnel and organizational dynamics, as well as operational and technical requirements, will blur the lines of who protects the data and which tools are used in support of the protection strategy.

HPE's StoreOnce Systems continue to lead the way with integration capabilities via Catalyst, with StoreOnce, as well as with the ability to manage HPE 3PAR StoreServ snapshots using HPE RMC technology.

Data transmitting directly from production to protection storage while simply updating the catalog of a backup or management server is **undoubtedly the future of data protection**. As API-based approaches such as HPE's StoreOnce Catalyst and data management conduits such as HPE RMC continue to mature, it is inevitable that more workloads and data protection scenarios will mimic the approach.

The Bigger Truth

Although data protection continues to be a priority for organizations of all sizes, *how* it is achieved and *who* is invoking those processes continues to evolve as workloads and organizations transform themselves.

What hasn't changed is the need to accomplish data protection in a cost-effective and highly reliable manner. HPE's recognition of this reality can perhaps be best seen in its creation and facilitation of a decentralized data protection toolset within a single data protection strategy—all of it underpinned by HPE StoreOnce.

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