

White Paper

Enterprise Cloud-Based Data Protection with HPE

Considerations and Solutions for Successful Cloud-based Data Protection

By Christophe Bertrand, ESG Senior Analyst

August 2018

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



Contents

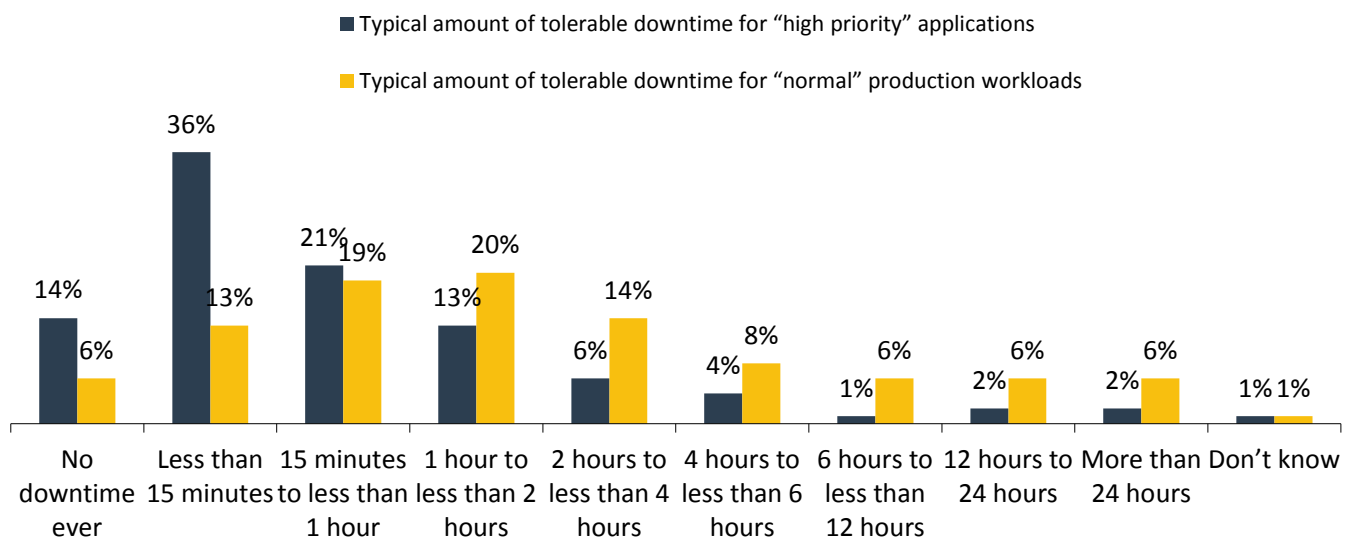
Data Protection Service Levels Drive IT Strategy	3
Leveraging Cloud for Resiliency and Long Term Retention	4
Protecting Data in the Cloud with HPE	7
Intel Inside.....	7
Storage-integrated Data Protection With Recovery Manager Central.....	7
Built for Cloud Data Protection with HPE StoreOnce Next Generation and Cloud Bank Storage	7
HPE Cloud Volumes.....	8
The Bigger Truth.....	8

Data Protection Service Levels Drive IT Strategy

Enterprises rely heavily on their IT infrastructure to support their organizational and business endeavors. Any interruption to the infrastructure or mission-critical applications exposes them to serious financial and reputational consequences. ESG research confirms that the tolerance for downtime is very limited across the board for mission-critical workloads—71% of organizations can only tolerate less than one hour of downtime for high priority applications before invoking a failover.¹ Breaking down this number reveals that 14% of respondents can tolerate no downtime whatsoever for “high priority applications.” More interestingly, even “normal” production workloads are becoming the target of higher availability demands, with 6% of respondents not able to tolerate any downtime, and close to 40% able to tolerate only less than one hour.²

Figure 1. Amount of Downtime Tolerance

What is the amount of downtime your organization can tolerate from servers running “high priority” workloads before making the decision to “failover/recover” to a BC/DR secondary site or service provider? What is the amount of downtime your organization can tolerate from servers running “normal” workloads before making the decision to “failover/recover” to a BC/DR secondary site or service provider? (Percent of respondents, N=320)



Source: Enterprise Strategy Group

These are significant findings because they demonstrate that application and data availability service levels are simply business-critical. This places the role of IT and its ability to deliver on availability service levels at the heart of corporate strategy, simply because no or little downtime can be tolerated in the digital world we now live in. Failing to deliver results in this arena can result in heightened business risk, such as non-compliance with data availability regulations, as well as financial loss and/or heightened operational costs. This is true across the supply chain, whether interacting with an end-user transactional application, like a non-responsive e-commerce site, or a manufacturing environment that has become idle. Unfortunately, the list of possible negative outcomes is long. Further research provides great insight into these impacts, with loss of customer confidence and direct loss of revenue most-reported as concerns by ESG research respondents, making them the top two responses (see Figure 2). Additional impacts can have long-lasting effects such as damage to brand integrity—a marketer’s nightmare—and loss of employee confidence. While loss of customer or brand

¹ Source: ESG Master Survey Results: [Real-world SLAs and Availability Requirements](#), May 2018. All ESG research references and charts in this white paper have been taken from this master survey results set, unless otherwise noted.

confidence following an availability issue are reasonable and expected outcomes, the domino effect also clearly reaches inside the organization. This only places additional pressure on IT teams to put in place an efficient and effective data protection and availability infrastructure to negate or minimize these impacts.

Figure 2. Impacts Resulting From Application Downtime or Lost Data



Source: Enterprise Strategy Group

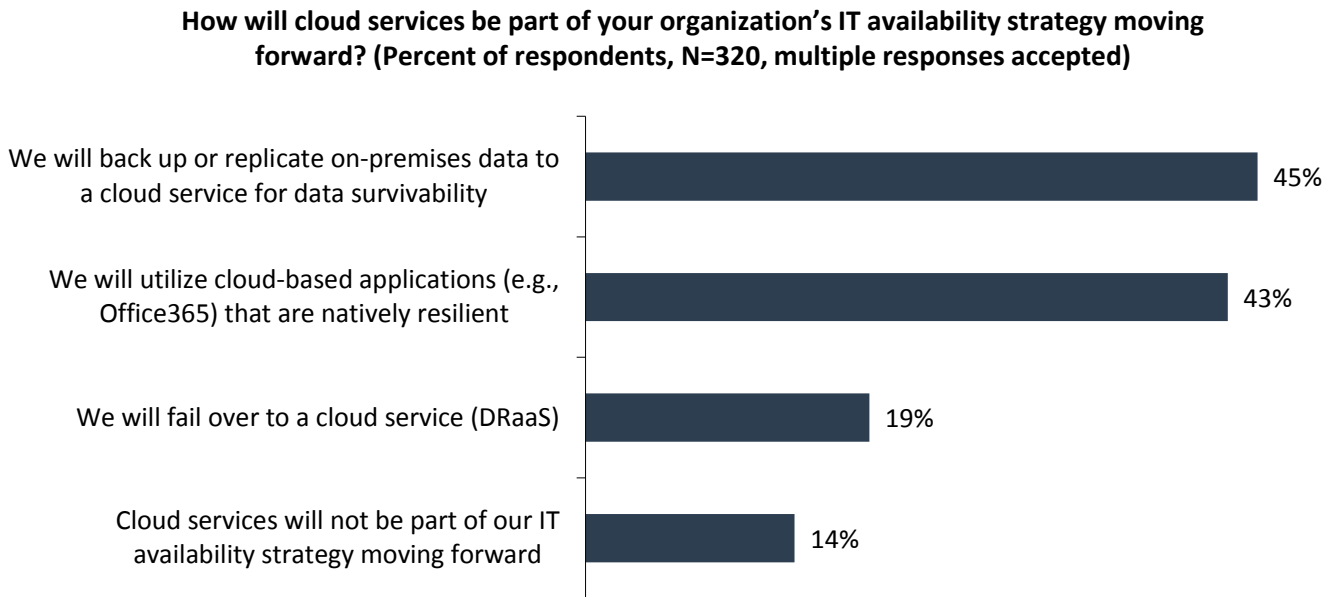
When it comes to the determination of service level agreements (SLAs), 40% of organizations reported that IT is the primary driver, while 26% reported that executive leadership sets SLAs and 23% reported that they are set by the business units. This healthy and close split should act as a motivator for both “sides” of the organization to actively collaborate in setting proper business and technical expectations. It should also be noted that, from a reporting perspective, uptime and outages are only visible to IT in 15% of cases, meaning that, in most cases (82%), they are visible to either executive leadership (33%), business units (16%), or both (33%), making uptime a true business key performance indicator (KPI).

Leveraging Cloud for Resiliency and Long Term Retention

Cloud technology and services have clearly established themselves as key components of a successful availability strategy, adding to the first line of defense provided by on-premises topologies. IT leaders use multiple cloud-based technologies, such as cloud-based applications with native resiliency between sites, cloud-based storage for data survivability, backup-as-a-service (BaaS) with cloud-based data restoration capabilities, and replication to a cloud host (DRaaS), for their IT resiliency needs. ESG research additionally shows how cloud services are getting woven into availability strategies, with 45% of organizations indicating that backup or replication to a cloud service will be part of their IT availability strategy moving forward and 19% indicating the same for disaster recovery-as-a-service (DRaaS).

These are significant trends that highlight the fact that IT decision makers do leverage these cloud options to deliver on the data protection SLAs that support their businesses and augment on-premises hardware and software.

Figure 3. How Cloud Services Will Be Part of IT Availability Strategy



Source: Enterprise Strategy Group

In addition, when asked about their perceptions about the resiliency of cloud-based applications and servers compared to on-premises applications and production servers, the majority of organizations (66% for server, 68% for applications) agree that cloud-based options are equally or somewhat more available/resilient than self-managed applications. While these are perceptions and not strictly measured operational survey results, the reality is that cloud passes the test for most IT leaders when it comes to resiliency levels, supporting business SLAs.

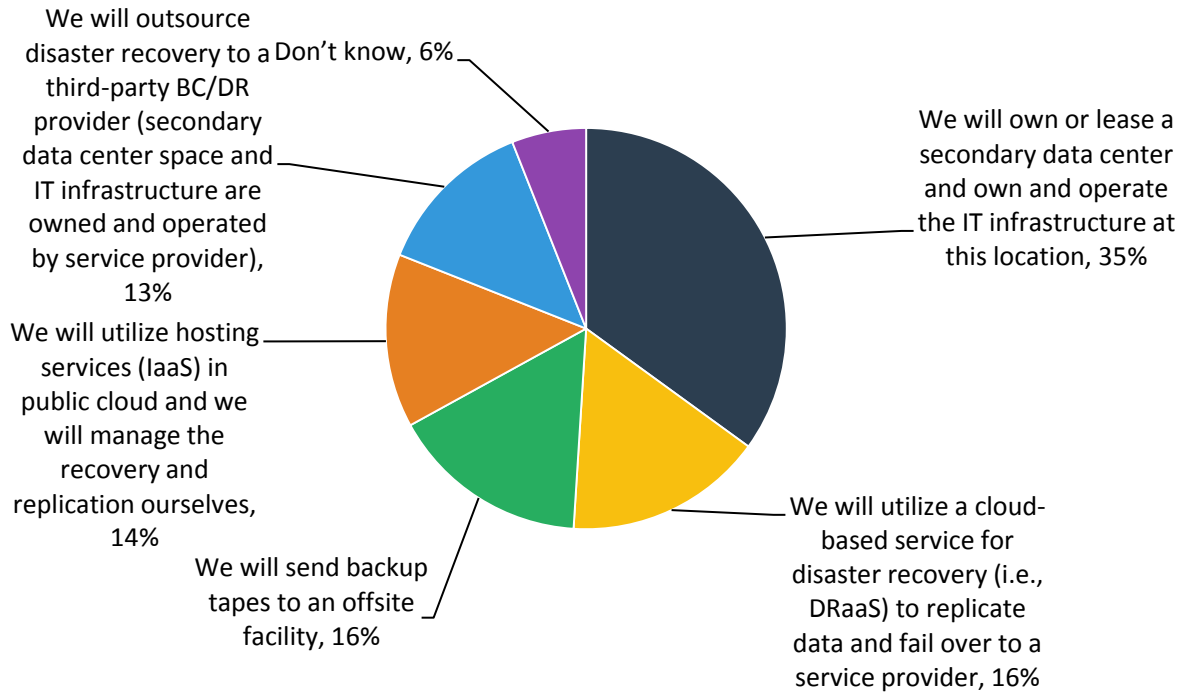
The market is evolving, and projecting ahead, while we see an evolving role for cloud-based technologies for availability and disaster recovery (DR), many respondents will still heavily rely on their own IT infrastructure on-premises or outsourced to a third party where the infrastructure is owned and operated by a service provider. Thirty-five percent of organizations believe that they will own or lease a secondary data center and operate the IT infrastructure within it as their primary disaster recovery process in the next 24 months (see Figure 4). This clearly means that on-premises IT is not going away any time soon. However, the use of public cloud and cloud-based disaster recovery services will be favored by 30% of organizations in the next 24 months.

Many end-users also favor cloud “destinations” for additional use cases beyond disaster recovery or resilience. Cloud infrastructures clearly allow for a tiered approach to backup retention and offer a convenient solution for long term data retention or archiving processes, with 40% of organizations reporting the use of the cloud as a repository for backup and/or archive data.²

² Source: ESG Master Survey Results, [2018 IT Spending Intentions Survey](#), December 2017.

Figure 4. Primary DR Process 24 Months From Now

Which do you believe will be your organization’s primary (as determined by the method providing resiliency to the greatest percentage of recoverable servers) disaster recovery process in 24 months? (Percent of respondents, N=320)



Source: Enterprise Strategy Group

With the data and trends in mind, it has become clear that IT leaders, if they haven’t done so already, need to adjust their infrastructure with a clear line of sight to cloud integration for backup, resiliency, archive, or disaster recovery. While traditional on-premises-based data protection is here to stay, fully integrating with cloud-based infrastructure is the name of the game. Hybrid architectures are becoming the new norm across enterprises: Organizations are increasingly using or looking into hybrid “on-premises plus cloud-hosting” architectures as part of their BC/DR strategy, with 31% already doing it, and another 48% planning or interested in deploying a hybrid cloud approach.³

Operationally, best practices must be deployed to ensure SLAs and successful recovery processes regardless of the circumstances. For example, the ability to use multiple copies of data, in two or more locations, with “air-gapping” capabilities (offline copy) is a best practice that can be augmented with cloud integration.

Modern and innovative data protection techniques that allow the delivery of stringent data protection SLAs, providing a balance between on-premises and cloud, are needed. At the same time, the balancing act for many IT investments is to also provide a robust foundation architecturally, using a combination of proven vendors and infrastructure building blocks for storage and its components. One of these rare vendors is HPE.

³ Source: ESG Survey, 2017 Trends in Data Protection Modernization, December 2017.

Protecting Data in the Cloud with HPE

On the basis of constant innovation built on many years of expertise in developing enterprise-class data protection solutions, the HPE portfolio provides a complete set of built for cloud solutions that provide customers the best of both the on-premises and cloud worlds. The portfolio provides options and flexibility to end-users, allowing them to use individual solutions or combine them as needed.

Intel Inside

The HPE solutions are designed for resilience and performance from the ground up and leverage a strong and long partnership with Intel. Intel technology inside HPE servers and storage, including Intel Xeon processors and Intel SSDs, offers the performance, orchestration, and security features organizations need to better manage shared resources, modernize their data centers, and support their IT service levels. HPE and Intel together provide these organizations with a range of innovations designed for optimized performance in a hybrid infrastructure environment.

Storage-integrated Data Protection With Recovery Manager Central

HPE offers the **3PAR StoreServ** family of mid-range and enterprise flash-optimized data storage systems for tier-1 data center environments and boasts a guarantee of 99.9999% data availability. The solution features a well-established high-performance and automated provisioning feature set and is designed for multi-tenancy. The Intel-powered hardware provides for built-in deduplication and compression and sub-1ms latency.

HPE Nimble Storage also leverages Intel flash storage and is designed for NVMe and Storage Class Memory (SCM). The Nimble family, well known for its signature predictive analytics, offers, like the 3PAR family, a guarantee of 99.9999% availability. It includes native cloud integration with AWS, Azure, and cloud-native APIs. It is positioned as a very efficient family of arrays, with the ability to store more data per terabyte of flash storage than other all-flash arrays. It can also be combined with HPE Cloud Volumes for full cloud integration.

HPE Recovery Manager Central (RMC) software integrates HPE 3PAR StoreServ All Flash Arrays with HPE StoreOnce Systems to provide a flash-integrated snapshot, replication, and backup service that augments traditional backup approaches. RMC allows for rapid recovery with non-disruptive, *application consistent* local and remote snapshots.

RMC does not require an application server or a backup server in the data path, which significantly reduces performance impact on the application. RMC leverages SnapDiff technology in HPE 3PAR StoreServ, which means that only changed blocks are sent to the StoreOnce backup system. Every backup is quick, as it is essentially only an incremental that then gets combined into a synthetic full backup, which accelerates application recovery. This also gives end-users more granularity in achieving their backup and recovery point objectives by having many points in time available for recovery.

Storage-integrated data protection will be extended to Nimble Storage with the upcoming release of RMC 6.0.⁴ The new version of RMC will also include a new SLA policy engine, enabling the creation and automation of backup, replication, and archival SLA policies with just a few clicks.

Built for Cloud Data Protection with HPE StoreOnce Next Generation and Cloud Bank Storage

The recently launched next generation **HPE StoreOnce** platform encompasses a broad range of deduplicated disk-based solutions, from virtual appliances and small form factors suited for SMBs, to enormous enterprise-scale deduplication platforms providing petabytes of deduplicated capacity. It can be combined with **HPE Cloud Bank Storage** software, which

⁴ Available December 2018.

natively moves backup data to the public, private, or hybrid cloud, enabling long-term retention and reliable disaster recovery.

Cloud Bank Storage sends, stores, and retrieves only unique data for lower TCO and is highly scalable, enabling the protection and retention of more than 100 PB of backup in the end-user's cloud of choice, such as AWS S3, Microsoft Azure Blob, and others. An ecosystem of ISVs and business apps also support this solution.

A number of security features are built into HPE Cloud Bank Storage to ensure safe data movement across tiers, including in-flight/at-rest encryption. It also protects against site loss by enabling the recovery of self-describing backup data from the cloud to any local or cloud-based HPE StoreOnce system.

HPE Cloud Volumes

HPE Cloud Volumes helps enterprises move their business applications to the cloud while maintaining an enterprise-grade storage experience. This solution provides the ability to use HPE Cloud Volumes as a replication target for any HPE Nimble Storage systems in the infrastructure, which makes migration for an existing data center and any other data movement operation to the cloud very easy. Another advantage is the ability to rapidly on-ramp to the cloud, move to another provider since this is a multi-cloud service, or just repatriate data back in-house without incurring costly data egress service charges. The solution is designed to deliver proven storage uptime, security, and availability. From a data protection standpoint, it should be noted that it includes instant snapshots for faster backups and recovery and instant zero-copy clones, which can be leveraged for additional uses cases, such as development/testing and analytics, for example. In addition, end-users can benefit from advanced capabilities with Infosight predictive analytics.

The Bigger Truth

Hybrid data protection and resilience infrastructures are quickly becoming the norm across enterprises, making any investment in technology a cloud integration decision. Selecting the right vendor on the cloud integration journey when it comes to data protection should be motivated by key criteria such as the depth and breadth of its portfolio and features set, its integration points to public and private clouds, its ability to consistently deliver efficient and fast backups or recovery processes...and a proven architecture to do so. The prize is not related to technology, but to business availability, measured in terms of resiliency SLAs.

This is where HPE shines by not only providing a robust portfolio of solutions that can be deployed individually or in combination but also by constantly innovating and further bridging the SLA gap with new releases such as next generation StoreOnce and RMC 6.0.

To deliver against key data protection challenges, the fundamental building blocks of HPE's solution leverage a long partnership with Intel around which advanced storage capabilities are built and combined with integrated cloud, software, and services solutions.

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.



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

© 2018 by The Enterprise Strategy Group, Inc. All Rights Reserved.

