ESG

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

SOLUTION SHOWCASE

Understanding How Customers Can Achieve an All-IBM Cloud-powered Data Protection Solution

Date: July 2017 **Authors:** Jason Buffington, Principal Analyst; and Monya Keane, Senior Research Analyst

Abstract: When talking about cloud services to achieve the kind of agility and recoverability organizations are demanding moving forward, data protection should always be about "*and*" instead of "*or.*" In other words, modern *hybrid* cloud protection must be the foundation. Few vendors in the world can deliver the breadth of products and services that IBM can. Its "One IBM" solution of IBM Spectrum Protect leveraging IBM Cloud Object Storage is a strong example.

Introduction

Today's best hybrid cloud backup architectures combine the speed of an on-premises solution with innovative remote-protection cost and agility models in ways that most pure-disk or pure-onsite options cannot. Aware of this reality, savvy IT managers combine and deploy multiple IT topologies to back up their organizations' primary production data effectively and establish an ideal protection plan.

According to findings from the *2017 ESG IT Spending Intentions Survey*:[1]

- Cost reduction is the second most common business initiative expected to drive IT spending in 2017, behind cybersecurity.
- When asked how they would be reducing costs, more than one in four organizations said they would increase their use of cloud services.
- Storing data as a repository for backup/archive is once again among the top use cases for current users of cloud infrastructure services.

Looking at the hybrid cloud data protection architectures of many organizations today, snapshots and backups within disk-based systems serve as the first line of recovery. Tape supports long-term retention. And the cloud provides disaster recovery. Each medium has its strengths:

- **Disk** is the highly agile, first-tier choice for item-level restores and deduplication and compression optimization. Per ESG's recent data protection research, 74% of organizations leverage disk as their first tier of recovery.[2]
- **Tape** isn't as ideal as a first recovery tier, but it outshines disk for long-term retention and archiving—serving as a reliable, economic way to transport data and store it for years or decades. ESG finds that nearly half (48%) of surveyed organizations use onsite or offsite tape for archiving or long-term retention as part of their backup process.[3]

[1] Source: ESG Research Report, *2017 IT Spending Intentions Survey*, March 2017.

[2] Source: ESG Research Report, *2017 Trends in Data Protection Modernization*, to be published.

- The cloud offers long-distance data survivability, provides on-demand agility for BC/DR and analytics, and reduces onsite or DR-site hardware expenditures. Seventeen percent of ESG survey respondents anticipate using cloud services as part of their backup strategy (typically alongside disk and/or tape on-prem). And adoption rates appear to be rising.[4]

  Technically, of course, "the cloud" is not an actual type of media. It is an alternative consumption model for disk/tape capacity, and it is often coupled with expert management services and delivery at scale to achieve higher reliability and economic savings.

Catches Exist

For large enterprises, offsite cloud data protection is neither indefensibly better suited to all scenarios, nor is it equitable from all providers. For example, some cloud backup solutions started as consumer services that were "beefed up" to attract enterprise clients, but not all made the transition effectively.

Also, using a cloud service doesn't equate to getting out of the backup business. Many IT managers think they'll start using a cloud solution and everything will be magically better—forgetting they'll still need to maintain/manage backup agents and onsite disk to meet stringent requirements for uptime (see Figure 1).[5]

FIGURE 1. Downtime Tolerance for Applications Protected by a Cloud

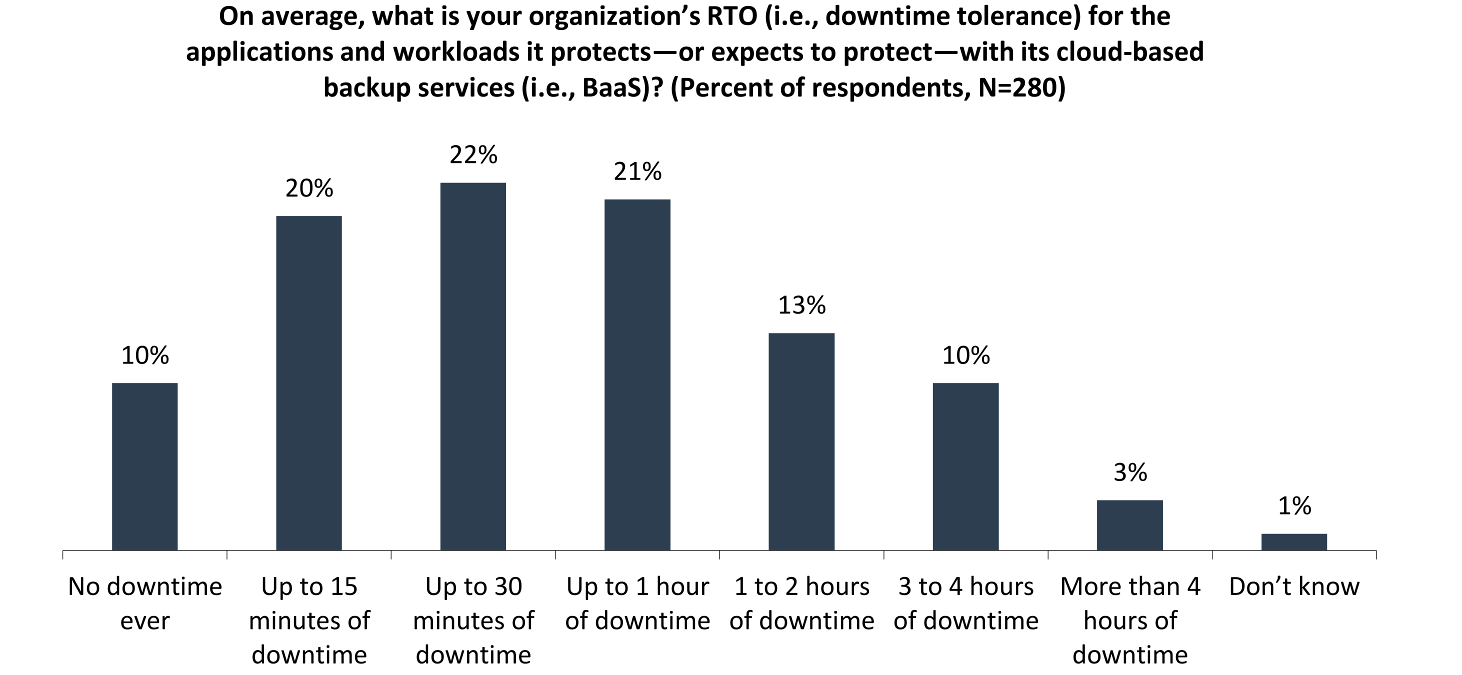

Source: Enterprise Strategy Group

Additionally, until recently, some cloud providers didn't store data for as long as their clients required, so those clients still needed economical cold media (i.e., tape) for truly long-term archiving. Fortunately, that situation has improved, with many of the bigger cloud service providers now offering economical pricing to store data on disk or tape for as long as their clients require to achieve truly long-term archiving and compliance.

[3] ibid.
[4] ibid.
[5] Source: ESG Research Report, *Data Protection Cloud Strategies*, December 2016.

As Figure 2 shows, well over half of the organizations surveyed by ESG tend to regard the cloud as a service to store data for three to five years. Many organizations actually have retention windows of six years or more for their "cold" data, but they can leverage the DR-related agility of the cloud for their "warm" data over shorter retention durations.[6]

FIGURE 2. Duration BaaS Users Plan to Store Data in the Cloud, by Length of Time Backup Data Is Retained

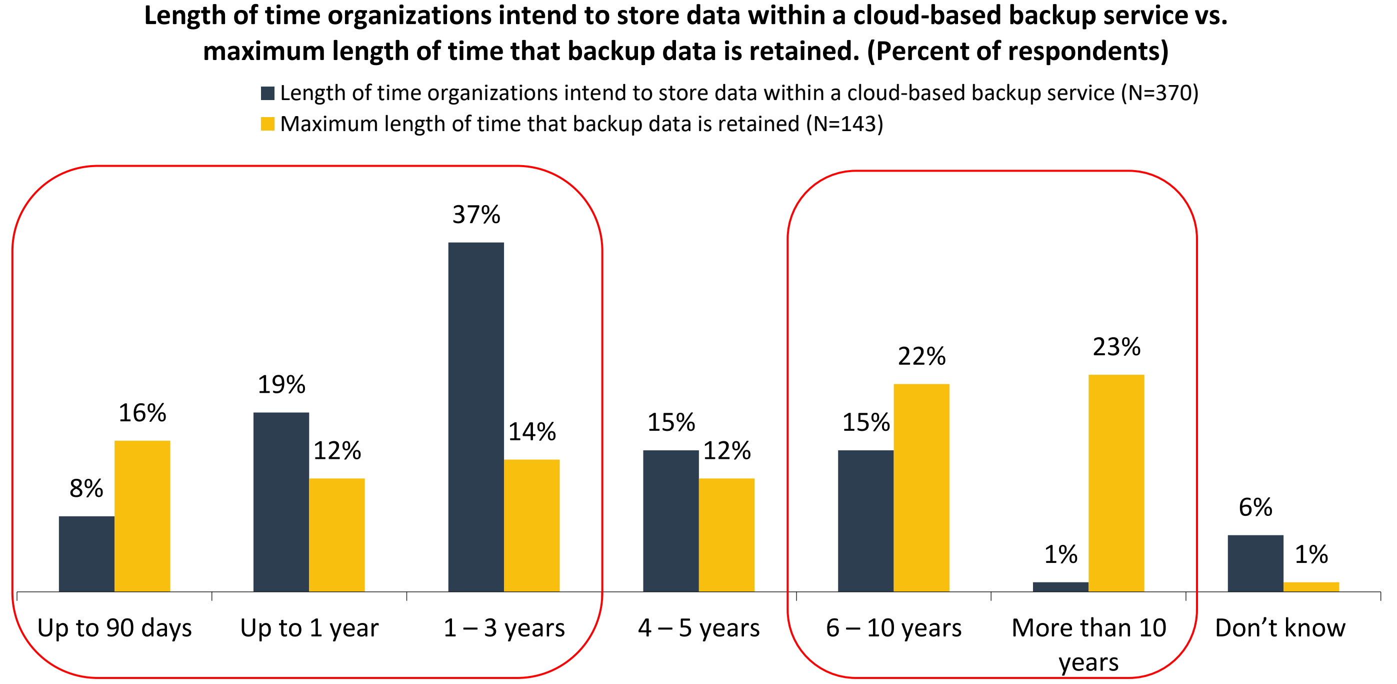

Source: Enterprise Strategy Group

It's Time to Do Something Better

Protection challenges can't be overcome based solely on which protection methods or media the IT team chooses—at least, not if the methods are insufficient or are perceived as being insufficient. Interestingly, 43% of respondents surveyed by ESG report that they would change their current backup solution if they were given the chance to rearchitect it from scratch.[7]

In some cases, their dissatisfaction is warranted. The top problem IT pros cite in protecting virtualized IT environments is *data recoverability*.[8] Occasionally (in 16% of cases), backup software caused a VM recovery effort to fail. More often, application problems or infrastructure bottlenecks were the cause.[9] And 11% of the time, recovery problems arose because the source was never designated for backup.[10] Those findings are specific to VM protection, but they illustrate how challenging "accomplishing backup better" can be.

Recommendations

A hybrid data protection architecture composed of on-premises and cloud capabilities is beneficial. ESG offers four recommendations to optimize such an environment.

[6] ibid.
[7] Source: ESG Research Report, *2017 Trends in Data Protection Modernization*, to be published.
[8] Source: ESG Research Report, *Trends for Protecting Virtualized Environments*, August 2015.
[9] Source: ESG Research Report, *Reliable Virtualization Protection Continues to Elude Many Organizations*, to be published.
[10] ibid.

- **Plan for disk-to-disk-to-cloud (D2D2C).** The uptime levels that IT is expected to achieve today can rarely be met by a "cloud-only" protection solution. So organizations *must* combine local recovery capabilities (i.e., snapshots and backups) with the agility and survivability options of one or more cloud services, potentially including backup-as-a-service (BaaS) for endpoints and remote offices, cloud backup storage for servers throughout the environment, and disaster recovery-as-a-service (DRaaS) for IT resiliency and BC/DR.

- **Look for cloud solution providers that serve not only as data repositories, but also as sources of expertise.** Economics and agility can make any cloud solution appear compelling. But the expertise to accomplish new recovery scenarios or unlock new value from data is also important. Some service providers offer "turnkey" remote management. Others provide consulting help with BC/DR or IT optimization. Regardless, local partners/resellers trained in relevant cloud technologies and services have a unique perspective, contextual know-how, and empathy with the local IT environment. They are experienced in accelerating this type of backup modernization/digital transformation.

- **Recognize that disparate repositories (such as local disk and remote clouds) can serve as "air gaps" to mitigate malware/ransomware attacks.** Although IT should address cybersecurity first through prevention instead of recovery, using multiple media connected only by a data mover for backup can create an additional layer of separation—one that may be key to stopping an infection from propagating. When organizations combine an air gap with analytics that identify radical changes to what had previously been dormant data sets, they may even be able to identify some malware events from within the data protection solution.

- **Use warm data for more than cold storage.** Cloud storage is often first considered because of its economic advantages. However, a lot of its real value comes from its ability to utilize secondary, natively accessible cloud-hosted data (i.e., warm data) in ways that are more viable than with tape-stored data (i.e., cold data). Examples include performing analytics or generating reports using the information, testing system patches or conducting DevOps testing without changing the protection copy, and performing sandbox orchestration as part of BC/DR planning. These scenarios are possible when one has 1) a dormant but near-current copy of data available, 2) on-demand compute for running processes only when needed, and 3) orchestration and/or expertise. The right cloud provider can deliver these services when the cloud is combined with the right data protection solution.

A Cloud-enabling Solution to Consider: IBM Spectrum Protect

IBM Spectrum Protect software supports enterprise-scale virtualized and hybrid IT environments, offering:

- **VM-optimized, application-aware backup services** for a range of operating systems and production workloads. IBM Spectrum Protect supports a unified strategy for physical and virtual server protection, incorporating hardware-assisted, application-aware snapshots and traditional backups.

- **Cloud optimization** including support for multiple public-cloud environments. Spectrum Protect enhances an organization's architectural flexibility, with public/private/hybrid options via the IBM Cloud or others.

- **Built-in efficiency** to support requirements for deduplication, compression, and incremental-forever backup. IBM Spectrum Protect builds high efficiency into software, reducing or eliminating the need for additional hardware.

- **Integrated, offsite, policy-based replication** to reduce the storage needed to support multi-site topologies, including hybrid clouds and DR facilities or services.

- **Self-service portals** to maintain service levels and reduce complexity.

- **Encryption and multi-tenancy in the cloud**, including encryption of data in flight between subscriber and cloud service provider, and encryption of data at rest within the cloud service provider's repository. Multi-tenancy lets a cloud service provider completely isolate each subscriber's data.

The Power of a "One IBM" Solution—Including IBM Spectrum Protect and IBM Cloud Object Storage

IBM Cloud Object Storage (COS) is built on the foundation of Cleversafe (acquired by IBM in 2015) and accessed through what has become an industry standard in cloud storage: S3 APIs. Although it would be easy to assume that all S3 storage providers are equal, IBM COS technology provides some very notable distinctions that make the IBM Cloud solution compelling:

- **Security—**The IBM Cloud platform uses erasure coding on the data, thereby slicing it, spreading those slices across three or more data centers, and then recombining them in real time for recovery/retrieval. All the slices are encrypted, with an erasure-coded key residing inside the data. No external key management is needed.

- **Availability—**Even if a region-wide outage occurs, data will be available "with strong consistency" via cloud-based data centers in other regions. IBM customers can choose resiliency within a single region or "cross-region," whereby mission-critical data is ensured to be always available around the globe.

- **Economics**—Unlike legacy technologies that rely on replicated copies for their durability, the IBM Cloud uses data slicing and dispersal, thereby delivering ensured data access without the 2X to 3X storage burdens internally plus the added operational costs to sustain each copy across sites. Because of IBM COS's data slicing and dispersal, there's only one copy of the data—making capacity usage, according to IBM, about 1.3x to 1.8x as IBM extends each slice across multiple IBM Cloud locations. Those cost savings are then reflected in the pricing for IBM's highly resilient cloud service.

By combining IBM's Spectrum Protect solutions with IBM's COS services, organizations large and small can achieve the same breadth and depth of recoverability that legacy data center architectures provided—while leveraging the power of cloud-based infrastructure.

For all the reasons that "backup alone is not enough" for data centers with varied recoverability mandates (a situation that encouraged integration with snapshots and replicas), organizations should not be forced to degrade their recovery options as they move forward with cloud modernization. A hybrid approach to data protection (specifically in regard to backup, snapshotting, and replication) is a best practice of a hybrid on-premises/cloud-based IT infrastructure.

All Solutions (Even Cloud-based) Are Better with Expert Services

This combination of capabilities is even more meaningful and achievable when it is delivered by certified partners, including IBM Business Partners and IBM Global Services offerings such as IBM Resiliency Services. By leveraging the technical expertise and on-premises insights of local partners who are well-versed in helping clients migrate production and protection environments to the cloud, organizations have even less of an excuse to put off reimagining their IT future in a cloud-first world. In keeping with the "One IBM" mantra, IBM Resiliency Services can operationalize the combined power of the data protection software/hardware and IBM Cloud—truly providing "one hand to shake" in a way that ensures customer success.

The Bigger Truth

For many organizations, D2D2C backup is likely in their future—secondary protection disk for agile recovery, and tertiary cloud services providing offsite survivability and boosting BC/DR preparedness. One huge indication of the value of a hybrid architecture centers on how well it enables IT to use warm data for more than cold storage. The IBM Cloud offers that "something more" quality in the form of agility, economics, *and* benefits tied to enabling IT to perform analytics, conduct tests and checks of all different sorts, take good advantage of outside expertise, and so on.

That outside expertise, in particular, can be invaluable—particularly considering the complexities of modern heterogeneous IT infrastructures and the assorted skills needed to back up data, recover it, and architect durable BC/DR frameworks and policies. Organizations gain a huge benefit when they can leverage cloud services, impressive software, and expert know-how together.

The word "hybrid" encompasses media (disk, tape, and cloud), topologies (onsite and offsite), and expertise (in-house and external). *But the common denominator is the software.* It leverages the media, operates across the topologies, and is the tool of the skilled experts.

Modern data protection strategies require a hybrid approach to media, where most recoveries should and do come from disk-based mechanisms (backups and snapshots) within the operating environment. For many organizations, tape remains a medium of choice for cold-data archiving and regulatory compliance; as such, IBM is a core innovator in mainstream tape (LTO) as well as its own enterprise tape format—and with its COS offering, it provides disk-based archiving for long-term retention as well.

Although most organizations carefully consider their tape and disk options for long-term storage, few understand the differences between the various cloud providers. Those organizations that conduct such due diligence will appreciate what the IBM Cloud and its IBM Cloud Object Storage (based on Cleversafe technology) provide to data protection solutions, including IBM Spectrum Protect.

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