commvault®
solving forward®

ADAPTING DATA PROTECTION STRATEGIES
TO CLOUD INFRASTRUCTURE

Eighty percent of organizations have already adopted, or plan to adopt cloud service solutions. Yet often overlooked, is how to protect data moving to a public cloud.

This document is designed for organizations who are **ACTIVELY RESEARCHING DATA PROTECTION STRATEGIES AS THEY MOVE WORKLOADS, APPLICATIONS, AND DATASETS INTO PUBLIC CLOUD INFRASTRUCTURES.**

Cloud Is the New Reality

Cloud computing is moving beyond the hype phase. The technology industry is preparing for a growing shift in consumption toward pay-as-you-go services, as organizations seek to get out of actively managing their own IT departments, with their capital investments and operating headaches, to free up resources to focus on their business core competencies.

According to research from Information Week, 40% of organizations are currently using cloud services, and another 40% are planning or considering doing so[1]. This is an increase from the prior year, albeit at a slower rate than many predicted. Regardless, this tracks to the widespread belief that cloud services will be a key part of most organizations' IT strategies in the not too distant future.

40% of organizations are currently using cloud services, and another 40% are planning or considering doing so.

INFORMATION WEEK2013 STATE OF CLOUD COMPUTING, FEB 2013

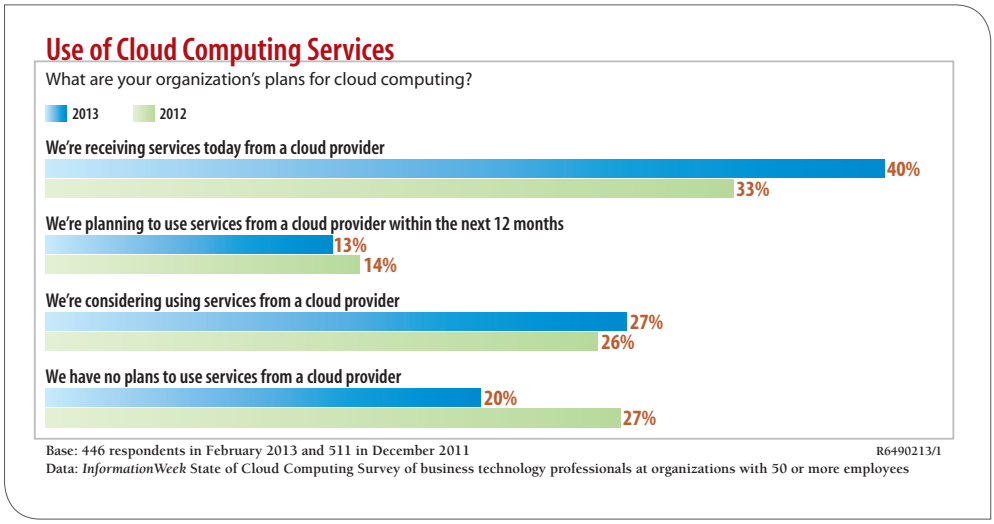

This steady march reflects the process that many organizations are undergoing, as they test, then ramp their approach. Many organizations have been "kicking the tires" on cloud services over the past few years – migrating one remote office, or one departmental workload, or one tier two or three application at a time. Now as many become ready to expand their usage based on successes and learnings from their initial pilots, we will see a tipping point in acceleration of IT spend shift.

Similar to the maturity curve seen with virtualization, cloud computing ramp will be driven by adoption into key workload types. Virtualization first gained a foothold in testing/development environments, where it could be proven out by organizations without impact to core business operations. The technology then became more broadly applied to individual and departmental workloads, and as a cost effective means of delivering disaster recovery, where business needs increasingly sought the flexibility and speed enabled by virtual infrastructures. Now virtualization has matured into use in production environments, as many organizations instituted a virtual-first mentality when bringing on new projects to avoid further investing in legacy approaches.

During this period of evolving infrastructure, organizations adapted their data protection schemes in order to meet backup and recovery needs. The business requirements

[1] Information Week, *2013 State of Cloud Computing*, Feb 2013. Report ID R6490513. 446 respondents from North America with at least 50 employees.

related to protection of data at each stage of this evolution did not change, even if the methods of protection did. Just as a business critical application running in physical or virtual environments on-premises would necessarily have a data protection scheme, the requirements should be no different when running in a cloud environment.

Enterprise Strategy Group finds that 24% of organizations are now using cloud-based data protection services in some way, and another 46% have plans or interest in doing so[2]. So, already almost 1 in 4 organizations are leveraging cloud services as a tool to achieve their backup or disaster recovery needs for some portion of their business.

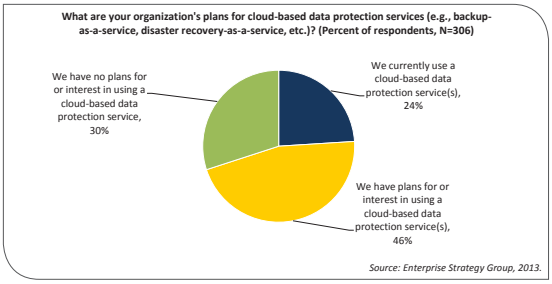

CommVault finds that many customers with data in the public cloud are not necessarily protecting that data. With the storage industry shipping over eight exabytes a quarter[3], and with an estimated 17% of spend going into public IT cloud services by 2017[4], that translates into a lot of data at risk of being left unprotected. In the solution brief, **Cloud Integration with Amazon Web Services**[i], learn how CommVault Simpana software seamlessly integrates with Amazon S3 and Glacier storage services to minimize long term capacity growth and TCO.

The Particular Challenges of Cloud Data Protection

Managing data across physical and virtual infrastructures poses many challenges to IT departments today. Data growth is driving continual increases in storage hardware footprint and costs. That hardware requires staff to manage upgrades, migrations, and day-to-day operations. Often that effort required to maintain the storage infrastructure takes focus away from leveraging the value of the data itself. And that data can be difficult for end users to access when it is sitting in silos across multiple platforms and locations.

Moving data into a public cloud environment brings several benefits as well as some new challenges. Cloud services provide a pay-as-you-go consumption model, freeing up cash and staff time locked into under-utilized infrastructure. Organizations become more agile, with the ability for IT resources to keep up with the constantly changing needs of the business. The data residing in cloud infrastructures becomes more easily leveraged to support business insights, legal, or compliance requirements. And the costs of storing data can be better aligned to the inherent value of that data – with cloud storage providing a cost-effective repository for long term retention of data that needs to be protected but is rarely accessed. However, some of the new challenges with data protection in public clouds include:

- **Understand the business relevance of data in the cloud**. The addition of public cloud infrastructures into an enterprise's IT environment can further add to the challenges of understanding what data is being stored, where, and its value and use to the organization. Data is now residing on multiple storage platforms across multiple data center locations, as well as on multiple cloud services outside those data centers. The need to profile the data as it is created, to ensure it resides on the appropriate tier for recovery time expectations, becomes even more important.

[2] Enterprise Strategy Group, Data Protection-as-a-Service (DPaaS) Trends, September 2013
[3] IDC Worldwide Quarterly Disk Storage Systems Tracker, December 6, 2013
[4] IDC Worldwide and Regional Public IT Cloud Services 2013-2017 Forecast, September 3, 2013 (IDC #242464).

- **Fully evaluate the costs of downtime or data loss**. Attempt to quantify and gain stakeholder alignment on the business impacts of the particular application or data set that will be managed via a public cloud service potentially going down. This enables the right investments in service levels and disaster recovery planning and testing to be commensurate with the value to the business. It seems obvious, but without this you may be paying too much or too little for what you actually need to ensure the appropriate level of availability.
- **Data migration practicalities**. Bandwidth constraints often pose additional challenges in migrating data into, and later recovering from, the appropriate public cloud services. Many organizations find that even when they are ready to start moving large data sets into the public cloud in order to free up capacity on expensive on-premises disk storage, the physical constraints of how long it will take to move the data can often make the endeavor impractical or cost prohibitive.
- **Addressing security concerns**. Public cloud providers have invested heavily in security features to address this key barrier to adoption. Robust features are available to ensure an organization's data is fire walled against intrusion protection, authenticated for user access controls, and ensured of secure upload and download processes. Additionally, encrypting the data itself provides another layer of security, and can be done so both at the source with in-stream encryption, and then extended to the data at-rest in the cloud storage target.
- **Data retrieval of what you need, when you need it**. Many organizations are finding that migrating, storing, and protecting data residing in public cloud services has become easier with continually improving vendor tools. However, the ability to retrieve data is often more complicated. Ensuring the ability to search and retrieve the exact file needed, by end-users self-serving their needs quickly and without IT intervention, can often be overlooked. Without this, public cloud services become just another storage tier, and not necessarily more agile, flexible, or cost-effective in some cases.
- **Portability of data**. As the market continues to innovate, organizations will want to take advantage of increased choices in public cloud infrastructure services. To avoid lock in, many organizations are using multiple providers across their environments. This certainly provides diversity in the features and pricing they have access to, but can also increase operational complexity to manage multiple vendor relationships. Ultimately, the ability to understand what data resides in which platforms, and to retrieve and migrate that data, are key factors in ensuring an organization's data can continue to be leveraged and even moved across a range of cloud infrastructures.

SUGGESTED VIDEO >>
"CommVault's Path to the Cloud"[ii].
CommVault VP for Worldwide Product management, Brian Brockway, discusses CommVault's journey toward creating Cloud-based solutions.
WATCH

Designing Your Cloud Data Protection Approach

There are a number of key considerations to include in designing a strategy to ensure an organization's data in public cloud infrastructure is protected and recoverable. Below are some to evaluate.

- **Align data storage costs with data value**. Just as with the virtualization maturation curve where tier two and three applications were the first to migrate, many organizations are on a similar journey to cloud. Organizations adopting cloud infrastructure as a tier to augment and offload expensive disk storage will want to prioritize which applications and data sets are moved to the cloud first. While public cloud can deliver great agility and performance, moving mission critical applications

requires the underlying services to provide the features and service level guarantees required for that highly valuable data. Conversely, public cloud can also provide cost-effective archival storage for data that requires longer term retention for 90 days, or 7 years, or even 100+ years, but is expected to be rarely accessed. In these cases, the tradeoffs associated with a lower-cost, less feature rich service is appropriate for that data being stored.

- **Availability is not equivalent to protection**. Public cloud service providers will replicate client data across multiple zones, as a standard or upsell service, in order to ensure data access when a zone becomes unavailable. While replication is extremely valuable to ensure availability, it does not guarantee protection of the data. For example, if data were to become corrupted or compromised in one location due to a user human error, then that may be duplicated into the other locations. In that scenario, the only way to recover that data would have been to run a protection copy.
- **Cloud is not just more cheap storage**. With some public cloud services being so easy for any user with a credit card to turn on, it runs the risk of just becoming another bulk offload repository. Using cloud services in this manner exacerbates the problem that many organizations face in ensuring a holistic approach to data management across the enterprise. Ensuring that data is profiled and indexed, before it is moved into the cloud, allows data to be easily retrieved when it is needed – without sifting through mounds of data in order to get to the particular file needed. Furthermore, an organization may find itself paying each month, perhaps into perpetuity, to store data that is actually no longer needed. Over time this can become even more costly than owning the infrastructure. Thus, actively managing the data that goes into the cloud is key to ensuring this becomes an effective tool to realize the ROI improvements that are driving most organizations to the cloud in the first place.
- **Recovery costs are what matter**. With costs continuing to decline, now down to $0.01 per gigabyte per month from some providers, cloud storage looks increasingly attractive to organizations looking to reduce their storage expenditures. However, it's the recovery costs that could potentially drive significant, unexpected costs without the right holistic data management in place upfront. Without defined retention policies, data being sent into the cloud will continue to grow exponentially, just as with on-premises storage. People and process must define those policies, and then technology in the form of data management software tools can ensure execution of those policies to automate the profiling, tiering, protection, and sometimes ultimate deletion of data based on pre-set rules. This process then ensures much more rapid, granular, and time- and cost-effective recovery of data, only what is needed, when it is needed.
- **Management tools are critical**. Integration of cloud platforms into an organization's existing data protection tools help deliver on the intended 'cloud experience'. Ensuring tight control and visibility into data protection operations, reporting and charging back to the business, the ability to recover data rapidly, and supporting self-service access -- all help enable the speed, agility, and accessibility required by users from a cloud service.
- **Snapshots versus traditional backup**. Snapshot technologies are modernizing data protection, by ensuring application recoverability without impacting uptime and performance. This also holds true for workloads running in the public cloud, with some considerations. Most public cloud providers support snapshots, replicas, and simple snapshot management of applications running in their cloud compute services. However, these often lack true lifecycle management and content management

features, making it challenging for organizations to sift through and find the data they need, when they need it. Also, just as with leveraging hardware array based snapshot capabilities, this can often pose challenges to organizations using multiple brands of storage, as well as cloud services. A holistic snapshot management framework provides an enterprise wide view of data as well as a consistent tool set for IT staff to manage.

- **A single platform approach**. The integration of cloud infrastructures could add to the challenges of managing across heterogeneous IT environments if treated as another silo. The need for a single platform that provides a holistic view across stacks of on- and off-premises infrastructure – regardless of the infrastructure chosen -- means that an organization can be freed from needing to know where the data resides, just that it is protected and accessible from just a few clicks away. While the underlying stacks can certainly be quite diverse, a single management platform enables the data itself to be handled consistently – to be easily accessed and leveraged across the enterprise. Only then can organizations begin to leverage technology as a value, rather than cost, to the business.

Getting Started On Your Journey

In summary, the shift towards consumption of IT as a service is widely seen as a matter of "how fast" and not "if" it will ramp. CommVault can help clients leverage this exciting trend successfully, by ensuring that the workloads, applications, and data sets being moved into public cloud services are efficiently protected, managed, and accessible. An approach that includes data protection considerations at every stage of a public cloud initiative -- defining organizational needs, evaluating provider choices, defining a test and rollout plan, and then scaling out -- is critical to success, and CommVault is here to help every step of the way.

To learn more about the full benefits of CommVault Simpana software and its revolutionary approach to cloud data protection, please visit: **commvault.com/cloud**.

Resources

[i] commvault.com/resource-library/1843/commvault-amazon-web-services-solution-brief.pdf
[ii] commvault.com/resource-library/2031/commvaults-path-to-the-cloud.mp4

commvault
solving forward™

www.commvault.com • 888.746.3849 • get-info@commvault.com

COMMVAULT REGIONAL OFFICES: UNITED STATES • EUROPE • MIDDLE EAST & AFRICA • ASIA-PACIFIC • LATIN AMERICA & CARIBBEAN • CANADA • INDIA • OCEANIA