ENHANCE YOUR DATA PROTECTION STRATEGY WITH

CLOUD-BASED BACKUP, REPLICATION AND DATA ARCHIVAL

TABLE OF CONTENTS

INTRODUCTION

Information technology (IT) is always changing.

Innovative technologies — and applications that leverage them — seem to debut daily. Buzzwords that dominate industry news and blogs come and go. Today's trends are tomorrow's "back in the good old days."

However, one topic has stayed top-of-mind for many IT professionals over the last decade: the cloud. It's changed the way business is done and had an equally significant effect on the IT world.

The cloud remains a leading IT topic largely because it offers so many benefits: scalability, cost savings, security, availability, performance and more. As the technologies behind it evolve, new uses continually emerge across just about every industry and market sector.

Among those uses is data protection. This document provides a brief history of data protection and outlines some of its main components —backup, data archival and replication. It also discusses the benefits of using cloud-based backup, data archival and replication to enhance your data protection strategy.

67%

of enterprise IT infrastructure and software will be cloud-based by the end of 2020[1]

[1] LogicMonitor's Cloud Vision 2020: The Future of the Cloud Study

THE EVOLUTION OF DATA PROTECTION

Humans have always had a thing for data.

Some of the earliest known examples of using data go back to around 18,000 BCE. Paleolithic tribespeople would mark notches on bones to keep track of inventories and trading activities.

Jump forward to around 344 BCE. That's the approximate date when the Great Library of Alexandria was founded. Housing an estimated half million scrolls that represented the accumulated knowledge of scientists, physicians and philosophers up to that point, it was considered the largest collection of data in the ancient world.

A tally stick was an ancient device used to record and document quantities and messages. One of the most notable found is Ishango Bone. Found in the Belgian Congo, it's believed to be 20,000 years old.

In 1943, the first Colossus computer was used to decipher secret codes on paper tape — a predecessor to magnetic tape.

But bones break. Libraries get destroyed by fires, wars, and other disasters. The data they contain can easily be lost forever.

No one is sure how or when the Great Library of Alexandria was ultimately destroyed. Its contents were lost — but perhaps not entirely. It's believed that the library's guardians moved some of the library's collections to other buildings in the city, demonstrating what may have been the world's first data protection strategy.

In addition to geographically dispersing data, there's also plenty of evidence of data backup efforts throughout history. Among the most notable: scribes routinely made hand-written copies of writings on papyrus and other materials and hid them.

Data has evolved since ancient times, with bones, clay tablets, scrolls and even books replaced by 0s and 1s. Data protection has evolved as well. Developed in early 19th century France, punch card technology became the predominant way to collect and store data.

Then in the 1950s, magnetic tape was introduced. Each roll of tape could store as much data as it once took 10,000 punch cards to preserve — and it cost less to process.

Next came the hard disk drive. Even today, many businesses save their backups on them, storing them on site or off site in dedicated data centers.

The use of cloud technologies is the latest in the evolution of data protection. But before diving into the benefits of cloud-based data protection, it's important to understand the most recent iterations of the traditional data protection strategies that's its positioned to replace.

THE BACKUP SIDE OF DATA PROTECTION

Many people immediately think of backups when they hear the words data protection.

A backup is a copy of information that is used to safeguard against unexpected data loss and corruption. If you lose your original data, just use the backup to make it available again.

Backup copies of data are made at intervals. It could be hours or days between the copies, depending on business requirements. The copies are then saved to an on-site or off-site physical hard drive, tape, disk or to a virtual tape library (VTL). For quicker recoveries, performance disks or flash drives are used.

For optimal data protection and often for compliance purposes, data copies shouldn't be stored anywhere near the primary data storage location. This mitigates the risk of losing both the primary and backup storage in the event of a catastrophic event that wipes out where your data is stored.

In fact, best practices for backups include the 3-2-1 rule. Keep three copies of your data stored on two types of media with one copy being off site.

3	**# of copies of data**
2	**# of types of media for data storage**
1	**# of copies to keep offsite**

The data remains stored until something goes happens to the primary data storage. At that point, the backup data can be accessed and used to restore lost or damaged data.

Other than the cost for storage media, backup is relatively inexpensive. But it may take time for your IT staff to retrieve and recover the data. That's why it's usually reserved for data you can do without for 24 hours or more. Application performance can also be affected each time a backup is done.

Another issue: the physical media used for storing data can fail. There are ways around this, such as using synchronous mirroring and RAID protection. However, there are time-related costs associated with recovering from a media failure and returning to a protected state.

THE ROLE OF DATA ARCHIVES

Data archival can also play a role in data protection.

While a data archive is similar to a data backup, it serves a different function.

A backup is a copy of your data is made to protect against loss of that data. The original data isn't deleted when a copy is made, but older backups maybe to make room for newer backups.

A data archive is a copy of data made for long-term storage and reference. The original data may or may not be deleted from the source system after the archive copy is made and stored. However, it's common for the archive to be the only copy of the data.

Data archives protect information that isn't needed for day-to-day business but must be retained for a long period — often for compliance purposes.

Across a wide variety of industries, compliance regulations require certain kinds of data to be stored for a specific time, often for years. With the right kind of storage, backup media can be stored indefinitely. Should auditors require that data, the backup media can be retrieved and recovered as necessary.

Because they aren't frequently accessed, data archives are stored on a lower-cost tier of storage. Consequently, moving data to this lower cost storage reduces primary storage consumption and related costs.

The problem is that lower-cost forms of storage are much slower. Archived data is indexed and has search capabilities, so files can be located and retrieved. But if you need your data quickly, data archival is not the way to go. And, as is the case for backup, failure of the storage media is always a possibility.

THE REPLICATION ASPECT OF DATA PROTECTION

Yet another component of traditional data protection is replication.

Replication is like backup in that it copies and moves data to a target or secondary site (for disaster recovery). The main difference is that replication copies data in real- or near-real time, so you have a more up-to-date copy.

Changes to the source system or virtual machine (VM) are captured and copied to one or more target systems. This can be done between two onsite appliances or between appliances in different locations. The target system can then be rapidly brought on-line in the event of the failure of the main server.

Onsite replication is a high-availability scenario with replicas stored on the same site. Offsite replication protects critical data from problems at the production site by keeping data copies at a different location site.

Replication can be synchronous or asynchronous.

+ Synchronous replication creates copies of data in real time. It's best for environments that require reduced RTOs. It requires capable computation capacity because it creates latency and slows the primary systems. It also tends to be expensive, but it's very reliable in the event of a disaster.

+ Asynchronous replication creates copies of data per a defined schedule. It's designed to work over distances and uses less bandwidth than synchronous replication.

Replication can also be implemented using several different methods. Among them:

+ Array-based replication uses built-in software in compatible storage arrays to automatically replicate data between them. However, it requires homogenous storage environments as the source and target arrays have to be similar. The advantage of array-based replication is that it's robust and requires little coordination when deployed.

+ Host-based replication uses application servers paired with software to move copies of data from one site to another. This kind of replication is mostly file-based and asynchronous. It's also storage agnostic.

+ Hypervisor-based replication is designed to copy virtual machines (VMs) from one host server or cluster to another. This facilitates disaster recovery by easing failover to the copy of the primary system. Hypervisor-based replication can run on servers that don't natively support replication. But because it uses CPU resources, server performance can be affected during the replication process.

+ Network-based replication requires an additional switch or appliance between storage arrays and servers. It can support any host platform and work with any array. It's typically used with heterogeneous storage environments.

While backup is designed to provide data protection and archiving, replication is designed to reduce recovery time objectives and recovery point objectives (RTOs and RPOs). RTOs define the time it takes to recover from an outage while RPOs define the acceptable amount of data loss a business can tolerate. Because of its focus on RTOs and RPOs, replication is often used for disaster recovery (DR).

As with other DR strategies that employ a second site, using offsite replication can be costly. You have to pay for the real estate, equipment, power, cooling and staffing of that second site. Additional time, resources and expenses are also required if the site has to expand and more equipment is need particularly as storage capacity requirements grow.

Another potential downside of replication is that it copies every change, even if the change is because of an error or virus. To access data before a change, the replication process must be combined with snapshots, continuous data protection or another type of technology to create recovery points to roll back to if needed.

In addition, replication costs more than backup. As such, it's often reserved for mission-critical applications that must be up and running for business operations to continue during any business interruption.

THE CASE FOR CLOUD SERVICES

Cloud-enabled backup, replication and data archival solve many of the issues of their counterparts in traditional data protection strategies.

For example, all three are more cost effective because there's no need to purchase or maintain infrastructure. They also are able to take advantage of many of the cloud's benefits, including scalability, flexibility and built-in security.

The sections that follow cover some of the specific advantages that the cloud-powered versions of these data protection components have over their non-cloud counterparts.

TRADITIONAL VS CLOUD BACKUP

Cloud backup works like traditional off-site backup, moving a copy of data to another location.

The difference is that the location is the cloud, and the data copy is sent to it over a proprietary or public network.

The data remains stored until it's needed for recovery. At that point it can be accessed and restored from any internet-connected device.

Backups can be made to a company's own on- or off-site private cloud or to a cloud service provider's (CSP) public or private cloud. Going with a CSP offers greater cost savings because there's no infrastructure to purchase or maintain.

The following are some the benefits of cloud backups:

+ Traditional backup solutions are limited to the capacity of the storage media used to store data backups. Cloud backups are scalable. It’s easy to grow capacity as your data increases.

+ There’s no need to worry about planning capacity as with traditional backup. Most providers allow you to add storage capacity on demand, so there’s no upgrade or rollout wait for increased storage.

+ Cloud backups are extremely reliable. Due to being stored in a cloud environment, redundant drives compensate for possible hardware corruption and facilitates improved data integrity.

+ Data is accessible any time. Multiple network and internet connections to the cloud backup storage facilitate reliable access from most devices. There is no need for waiting for physical transport of backup tapes or drives from an offsite facility.

+ Cloud backups work well for scheduled backup. With a reliable cloud service provider and suitable online backup software, backups can be performed at regular intervals and be transferred automatically over a WAN or internet connection.

+ Most CSPs provide cloud backups using a pay-as-you-go model in which you’re only charged for the resources used. There’s no capital expenditures for hardware to store the data or software to automate and manage the backup process. There’s also no need to worry about upgrades, migrations or technology obsolescence. The CSP is responsible for the infrastructure.

- CSPs are likely to use the most advanced technologies and security mechanisms, which help to make their backup services faster and more secure.
- Compared to traditional backup solutions, cloud backup solutions are easy to deploy. There is generally minimal configuration, so you can be up and running in minutes on a fully redundant, enterprise-class backup platform.
- Cloud backup solutions require less management and maintenance overhead Most have easy-to-use interfaces.

TRADITIONAL VS CLOUD ARCHIVE

Like any other data archive, a cloud archive holds infrequently accessed data.

However, storing that data in the cloud can be more cost-effective. There's no need to buy or upgrade onsite hardware systems or software to manage and store non-essential data.

The disadvantage is that it can be more expensive to move data out of the cloud when it's needed. That's why some companies use a hybrid setup. Data that requires faster access is stored on premises with only rarely accessed data moved to the cloud.

Others use services that automatically move backup data into low-cost object storage based on certain rules or policies. That enables them to get unlimited capacity for long-term data retention at the lower cost of object storage. The integration with object storage also results in a space-saving, smaller footprint on your local storage that doesn't hinder recovery operations.

However, recovering data stored on tape can be difficult because it requires sorting through a large collection of tapes to find the desired data. When cloud archives make use of object storage, the data is stored with extensible metadata that is typically highly searchable.

Cloud archives also offer greater durability than other kinds of data archives. Over time, tape physically degrades. The cloud doesn't.

TRADITIONAL VS CLOUD REPLICATION

Cloud replication is similar to other types of replication.

All require a storage target. The biggest difference is that when the target is in the cloud, there's no need to invest in capital resources.

With most cloud replication solutions, you simply identify the source volume and the destination volume. You then choose a schedule and a replication policy, which defines how the storage system replicates data from a source volume to a destination volume.

You can failover to the replication environment and back to source as needed. Depending on the service, you'll experience fast failover and failback with almost zero data loss and minimal disruption. Plus, many CSPs that offer cloud replication only charge you for what you use in your target storage after deduplication.

Cloud replication also drives faster RTOs and RPOs. In addition, it allows you to scale your offsite storage quickly if your capacity requirements expand.

THE REST OF THE STORY

There’s much more that can be said about the advantages and disadvantages of both traditional and cloud-based protection.

What it comes down to is that there’s a wide variety of options available. The ones that work best are those that best fit your organization’s use cases, budget and other specific needs.

If you choose to go with cloud-based data protection, it’s important to thoroughly vet the service provider. At first glance, many cloud-based offerings look similar. The difference is largely quality of infrastructure and level of service versus just having a cheap target for your backup and replication files.

You should also look for service providers that can demonstrate their ability to meet your security and compliance requirements. They should have clear SLAs for the services you want.

In addition, find out how the service providers you’re considering handle capacity planning, change management, lifecycle management activities, data destruction and disaster recovery.

THE US SIGNAL ADVANTAGE

With US Signal, you get a partner who will walk with you on your journey towards data resiliency.

You get technology that will work when you need it. And you get support your long-term needs, even if your requirements change.

To complete your data protection plan, consider US Signal's Cloud Availability suite for Veeam. Those using Veeam Availability Suite 9.5 can now seamlessly back up and replicate to an offsite US Signal data center — creating ultimate availability for any disaster recovery (DR) plan, with no additional licensing fee. As a Gold Veeam Cloud & Service Provider (VCSP) partner, US Signal delivers fully-integrated, cloud-based DR and replication together with fast, secure cloud backup — enabled by Veeam Cloud Connect through US Signal's vCloud Director.

To learn what US Signal can do for your organization, call 866.2.SIGNAL, email us at info@ussignal.com, or visit our website at www.ussignal.com.

US SIGNAL®

veeAM