New Best Practices using the latest cloud and data protection technologies that enable Disaster Recovery (DR) in the Cloud

From DCIG president & founder Jerome Wendt

Numerous backup solutions, clouds and cloud storage tiers can lead to escalating cloud storage costs and complexity. These, in turn, hinder an organization's ability to perform DR in the cloud. There are three new best practices that address these new variables. By adopting them, enterprises can better position themselves to execute DR in the cloud faster while incurring lower costs.

THREE NEW BEST PRACTICES FOR DR IN THE CLOUD

1. Deduplication of backup data before storing copies in Cloud storage

2. Simplification of cloud storage management

3. Usage of cloud tiers and cloud lock features of cloud storage

Cloud storage forms the basis for DR in the cloud

Enterprises already store copies of backup data in cloud storage offerings from Amazon Web Services, Microsoft Azure and others. Now they want to use those backup copies to perform disaster recoveries (DR) in the cloud.

However, companies cannot assume that backup copies stored in cloud storage will immediately enable them to perform DR in the cloud. Rather, those backup copies only form the basis for it.

For example, some companies use different backup solutions. As a result, backup data they need to restore/recover to the cloud may not yet exist in the cloud. Other backup solutions may use cloud storage but manage it poorly. This can result in organizations having to spend too much time managing cloud storage. In other cases, organizations store data in cloud storage but place it in the wrong cloud or cloud storage.

Because of these issues, organizations need to consider how they manage and store backup copies in cloud storage. They want to use backup copies in cloud storage to perform DR in the cloud. But they also want to minimize their cloud storage costs, complexity and risks. The following three best practices position organizations to better achieve these goals.

1. Deduplication of backup data before storing copies in Cloud storage
2. Simplification of cloud storage management
3. Usage of cloud tiers and cloud lock features of cloud storage

Best Practice Nr. 1: Deduplication of backup data before storing copies in Cloud storage

Companies might assume that their backup software originally deduplicates all the data they back up and store. Therefore, they might also conclude that only backup copies with deduplicated data are stored in cloud storage. However, this ideal case is not necessarily what happens.

- ***Not all backup data is deduplicated.*** Backup software does not automatically deduplicate all data, especially for applications with high data change rates.
- ***Backup software solutions deduplicate data differently.*** Due to differences in how backup software algorithms deduplicate data, not all perform this task equally well. These differences can lead to potentially poor data deduplication ratios.
- ***Companies can use different back-up solutions.*** Organizations can use different backup solutions to address specific backup requirements of applications. They can assign robust backup solutions to high-priority applications and lower-priority backup solutions to lower-cost applications. These solutions store data in different backup vaults in different storage locations. In addition, lower-cost backup solutions may not deduplicate or compress data.

Failure to deduplicate backups before copying them to cloud storage can result in significant costs. Deduplication can reduce backup data sets by up to 95% if it achieves a data reduction ratio of 20:1. Storing backup copies on cloud storage without prior deduplication can result in costs that are up to 20x higher.

A software-defined, secondary storage management solution that deduplicates data addresses this problem. This solution enables organizations to perform DR in the cloud in the following ways.

- ***Deduplicates all backup data.*** The solution deduplicates all backup data before it is copied to cloud storage. This increases data deduplication ratios, reduces WAN bandwidth requirements and lowers monthly cloud storage costs.
- ***Additional software features improve backup throughput.*** The solution provides additional source-side data deduplication software that works with existing backup software.

It deduplicates backup data before it is sent through the network. This ensures that backup data is stored deduplicated both on-premises and on cloud storage.

- ***Companies store and manage all backup data centrally.*** This solution stores and manages all backup data, simplifying backup data storage management. Because of centralized storage, organizations can more easily schedule replication and verify that backup data copies are stored on cloud storage.

Best Practice Nr. 2: Simplification of cloud storage management

Whether organizations store backup data copies on one cloud storage platform or on multiple platforms, they face numerous challenges. Storing all copies on the default cloud storage tier may be too expensive. Using other lower cloud storage tiers raises concerns about storage management complexity and backup data availability. Organizations can also use multiple cloud providers to offset costs or meet specific recovery requirements.

In this scenario, organizations need to familiarize themselves with each provider's cloud storage interface and tiering options. They will still need to configure each backup solution to connect to each cloud storage offering to place backup copies on it. They may also need to configure each backup solution to be able to move backup copies between clouds and cloud tiers.

A software-defined, secondary storage management solution again addresses these challenges. It starts as an on-site storage solution that virtualizes existing storage. Any backup software can then store data on it by discovering and configuring it as a backup target.

When it hosts the backup data, the solution simplifies both on-premises and cloud storage management. It manages all of its software-defined instances on-premises. It also provides a centralized console for configuring and managing:

- Cloud storage offerings from different providers
- Backup replication in the cloud
- The different cloud storage tiers offered by some providers.
- The placement of backup copies on the appropriate cloud or cloud storage tier.
- Moving backup data between different clouds ie. cloud storage tiers.

Best Practice Nr. 3: Usage of cloud tiers and cloud lock features of cloud storage

Many organizations already understand and use a 3-2-1 backup data management strategy. They first create three copies of data, a production copy and two backups (3). They then store two backups on different types of media, with the third copy stored off-site.

When using a solution that manages cloud storage, organizations should modify this strategy somewhat. Companies still make three initial copies of the data: a production copy and two backup copies. However, the second of the two backup copies is placed on cloud storage.

After the copy is placed on cloud storage, enterprises should pursue two cloud storage functions to manage using the software-defined secondary storage management solution. First is to use cloud storage tiering. The last backup copy is to be kept on a higher-performing cloud storage tier. Then, when new backup copies are placed in the cloud, the older backup copies must be automatically moved to lower-cost storage tiers.

Also, the cloud lock feature should be enabled, if available. A lot of ransomware attacks backup software and tries to encrypt backup data. Many cloud storage providers offer a cloud lock function. Enabling it makes the backups stored in the cloud unalterable, so that ransomware cannot encrypt them.

A stronger foundation for successful DR in the cloud

By applying these three best practices, organizations lay a stronger and more secure foundation for performing DR in the cloud. A software-defined secondary storage solution consolidates backup data and centralizes cloud storage management. As a result, it achieves higher data reduction ratios, which in turn reduces WAN bandwidth and cloud storage requirements and costs.

Cloud storage continues to provide new opportunities for enterprises to perform DR in the cloud more quickly and securely. Organizations can create rules to store data on different cloud storage tiers to meet different RPO and RTO. They can even use the new object locking (lock) feature to better protect their backups from ransomware attacks.

Taken together, these three new practices provide increased confidence to perform DR in the cloud. Organizations have backup where and when they need it to perform DR in the cloud. They even achieve this goal with a simultaneous reduction in their overall cloud costs. ■

Sponsored by

Quest

Quest Software, Inc.
4 Polaris Way
Aliso Viejo, CA 92656
(800) 306-9329
www.quest.com/dataprotection

About DCIG

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analytics. DCIG analysts provide informed third-party analysis of various cloud, data protection and data storage technologies. DCIG independently develops licensed content in the form of TOP 5 report and solution profiles. More information can be found at **www.dcig.com.**