

## ► Recovery in an Ever-Changing IT Landscape

### 5 REQUIREMENTS FOR SMARTER RECOVERY

In the modern era, there are so many possibilities of how enterprises may wish to recover their data. But when trouble does happen, businesses need to return to service as quickly as possible. By adding automation to disaster recovery (DR) processes and cutting down on the touchpoints needed to test or failover virtual infrastructure, a tremendous quantity of manual overhead can be eliminated. Consider these five requirements for smarter data recovery.



The annals of backup are filled with horror stories such as this: a storage manager scrambling through backup tapes to find lost files. He or she either fails to find the right cartridge, or discovers that the data is corrupted or was never backed up properly in the first place. What isn't so well known is that such stories continue to be as commonplace as ever. Backup errors, failures and data corruption issues continue to plague IT.

Efficient backup, then, is certainly important, but only if it facilitates fast recovery – the point is not just to have a copy of data sitting somewhere. Its real purpose is to provide a copy that is accessible rapidly in the event of a data loss incident or disaster.

Additionally, the sophistication of today's IT environments demands the ability to recover workloads, regardless of where they are being protected, and restore them in the cloud. This freedom in terms of recovery location decouples organizations from having to maintain cold/warm duplicate infrastructure that is expensive yet rarely used. The ability to achieve this decoupling takes recovery to a whole new realm – recovering what you need, where you need it, and without being limited to your own four walls.

## ► SMARTER RECOVERY FOR THE AGILE ENTERPRISE

In the modern era, there are so many possibilities of how enterprises may wish to recover their data. For example, they may wish to take virtual workloads protected under one hypervisor and recover them into another. Alternatively, the organization may wish to have its data recovered to a cloud provider so as to avoid the need to operate a secondary data center. And there is even the option of providing operational recovery of primary site applications to a secondary site from an external cloud provider.

To truly enable smart recovery that will support enterprise agility and business continuity, consider these five requirements:

- 1 Automate DR processes.** Recovery should go well beyond the simple retrievable of a backup tape or file and enable the automation of DR business processes. This must encompass application recovery and testing, policy based VM provisioning, and cross-platform operation. The recovery process should also be able to deal with Azure, Amazon Web Services (AWS), VMware or Hyper-V, comfortably transition data from a private cloud to the public cloud, and from a physical to a virtual environment.
- 2 Achieve recovery on-demand.** On-demand live access is another key facet of fast-paced recovery. It takes the indexing of operations on the fly at the point of recovery, as well as the ability to mount volume/ DataStore/RDM to access data in native format. When IT wishes to recover some data, this should not cause disruption. Accordingly, it should be possible to execute recovery operations in the background, create thin copies of databases for secondary use and mirror block changes for automated failover to an alternate site.
- 3 Support instant and native data access.** The concept of recovery has been expanded well beyond the traditional approach of retrieving an old backup tape, and then seeing what it contains. To this, has been added the capacity to instantly access data. Further, protected copies of data should be accessible natively and directly, before even having to restore it. This ability to immediately access data sets makes it feasible to quickly validate that the right data has been found and then allow IT to rapidly direct its restoration from there. The possibilities this type of functionality opens up are almost endless. For example, the storage administrator can take protected copies of a database and activate them in order to use those copies for test/dev/

### Take a Fresh Look at Backup, Recovery and Archiving

Learn how to bring your backup, recovery and archiving processes up to the needs of a virtual and cloud-enabled environment.

READ NOW



QA purposes. Alternatively, those protected copies can be harnessed to manage the lifecycle so that once the data is no longer needed, it can be cleaned up.

**4 Simplify the cloning process.** With these capabilities, the whole subject of database cloning takes on new life. It should be a simple matter to expedite clone operations without having to go back to production to generate the data. Similarly, IT should have the ability to clone any database or snapshot to a separate instance via a one-step cloning process that can spawn multiple copies from the same snapshot. Automatic creation, cleanup and expiration of clones should also be part of the equation. Most importantly, these capabilities should not remain the province of only a few highly trained storage administrators. In order to match the speed of digital business, they should be available on a self-service for DBAs and end-users.

**5 Empower data portability.** Data should never be locked into one set of hardware or one platform. These days, it is essential to be able to take your data anywhere as an essential facet of continuous protection and replication for the migration of private cloud VMs or databases, public cloud databases and VM environments. Any recovery platform should give IT the freedom to migrate workloads across what used to be impassable boundaries – between sites, geographies, and even hypervisors.

In summary, smart recovery requires simplification down to the push-button level. Data loss is inevitable, whether through human error, hardware failure or environmental damage. When trouble does happen, businesses need to return to service as quickly as possible. By adding automation to DR processes and cutting down on the touchpoints needed to test or failover virtual infrastructure, a tremendous quantity of manual overhead can be eliminated.

▶ To learn more about Commvault® data protection, backup and recovery visit: [commvault.com/backup](https://commvault.com/backup).

© 2016 Commvault Systems, Inc. All rights reserved. Commvault, Commvault and logo, the "C hexagon" logo, Commvault Systems, Commvault OnePass, CommServe, CommCell, IntelliSnap, Commvault Edge, and Edge Drive, are trademarks or registered trademarks of Commvault Systems, Inc. All other third party brands, products, service names, trademarks, or registered service marks are the property of and used to identify the products or services of their respective owners. All specifications are subject to change without notice.



COMMVAULT.COM | 888.746.3849 | GET-INFO@COMMVAULT.COM  
© 2016 COMMVAULT SYSTEMS, INC. ALL RIGHTS RESERVED.

