

EBOOK

Disaster Recovery with Cloud Volumes ONTAP

A reliable, cost-effective, and flexible solution for both on-premises and cloud-based ONTAP storage environments.



Table of Contents

Executive Summary	2
Introduction	3
Disaster Recovery Challenges	4
Building Cloud-Based DR Environments	6
Cloud-Based DR with Cloud Volumes ONTAP	7
Benefits of using Cloud Volumes ONTAP	8
Conclusion	10

Executive Summary

Whether you are running NetApp ONTAP on-premises or already in the cloud, you need a good DR solution in place if you want to protect your data. AWS, Google Cloud, and Azure offer some native solutions and on-prem ONTAP users have their backup data centers, but both of these users could benefit from cloud-based DR using Cloud Volumes ONTAP.


Disaster recover allows you to ensure you can failover your business operations to a secondary location and later recover and failback to your primary copy reliability: Cloud Volumes ONTAP makes it possible to do that faster and more efficiently while paying less for updates and data storage.

Introduction

All enterprise platforms, whether on-premises or in the cloud, require a DR (Disaster Recovery) environment to ensure that business-critical services and applications can continue operating in the event of a major infrastructure outage. For example, a natural disaster such as a hurricane may cause outages at a primary site or even on a regional level. Security threats such as massive, targeted malware attacks can also take sites completely out of action. In these circumstances, it is essential to bring services back online as quickly as possible, which requires having standby systems in a secondary location that can be failed over to.

Deploying disaster recovery environments is challenging for most organizations due to the requirement for infrastructure and site independence. There are huge costs involved in establishing and maintaining such sites physically—basically the same costs as the entire production environment but for a site that will sit idle for most of the time. This is where the cloud helps to lower the barrier to entry by providing scalable Infrastructure-as-a-Service solutions that can be used to build DR environments for both on-premises or cloud-based systems. After building out all DR services, the challenge then becomes to synchronize data from the production environment, and to keep it in synchronized going forward.

In this guidebook, we examine in detail the challenges involved in setting up a DR environment, discuss the available services in AWS, Google Cloud, and Azure that can be used to build DR solutions, and look at how [NetApp Cloud Volumes ONTAP](#) provides cost-effective enterprise-grade support for data replication and disaster recovery both for existing NetApp storage system users and for cloud-based deployments with the major public cloud providers.



Disaster Recovery Challenges

Effective disaster recovery requires a full complement of the applications and services used in production environments. Everything you need to run your primary workload has to be reproduced from the bottom up. The necessary infrastructure must be planned for and implemented in advance, and deploying physical infrastructure for the setup of a DR site increases operational costs significantly.

Making sure this all works is a paramount concern. The longer it takes a site or application to operate normally, the more business losses will be incurred. Getting the site up and running is so important in DR, it has its own metric. The time taken to bring services back online after a DR event is known as the RTO (Recovery Time Objective), and the aim should be to reduce this interval as much as possible.

Setting up a replica DR environment also requires having an independent, up-to-date copy of all enterprise data, including database systems, file services, iSCSI storage, etc. As data in the production environment will be constantly updated, these data changes must be transferred to the DR site on a regular basis. The frequency with which data in the DR site is updated will determine the period that data loss will be incurred for after a disaster. The amount of data that can acceptably be lost is known as the RPO (Recovery Point Objective). For some enterprises data is so important that in a DR event no data loss can be tolerated, which means their RPO will equal zero.

RPO


Recovery Point Objective

How much data in a time period a company can lose during a disaster event.

RTO

Recovery Time Objective

How long it takes to get back to normal operation after the disaster event.




A fully-functioning DR site is useful for both unplanned and planned outages. For example, a DR site may be made live to perform updates to the primary production environment. This requires a failover to the DR site and then a failback after the updates have been performed. Failover and failback would also be used to make the DR site live after a disaster and to eventually restore services back to the primary site in the future, and so being able to perform these operations easily is a requirement for an effective DR solution.

Enterprise workloads and services are constantly evolving, and new software releases must be applied to both primary and DR environments. The primary site is used actively and so it will be possible to quickly verify that it is operating correctly. The DR site, however, may be left unchecked for a long period of time, and then may fail to come online when it's actually needed. To prevent that from happening, it is essential to perform regular testing of DR services and ensure that they are functioning as expected.

Finally, after completing the deployment of a DR site, it may seem as though it simply remains idle for much of the time. In the interests of improving cost effectiveness, it would be ideal to make use of DR systems for peripheral requirements, such as read-only reporting or for setting up software development test environments.




DR Challenges at a Glance




Scheduled syncs

Data must be synchronized efficiently and regularly to the secondary location to ensure it is up to date.




Failover & failback

Provide the capability for data storage to be failed over to the DR site and then failed back to the primary site as required.



Regular testing

Ensure that DR systems work as expected in the case a DR failover is required.



Controlling costs

DR compute and storage resources should be allocated in such a way to remain cost effective since the system is normally not in active use.

Building Cloud-Based DR Environments

Disaster recovery environments must implement redundancy at the compute, network, and storage layers. In this section, we will look at how this can be done using the three major public cloud providers native services.



Cloud DR Environment Building Blocks

Compute

The CPU processing power that runs applications.

Amazon Elastic Compute Cloud (Amazon EC2), Google Cloud Compute Engine instances, and Azure Virtual Machines provide flexible cloud compute resources that can be used to build and scale the most demanding enterprise workloads. A range of different instance types makes it easy to find the right fit in terms of CPU processing power and memory capacity. For containerized applications, all the major cloud providers offer native Kubernetes services. On AWS that service is Amazon Elastic Kubernetes Service (Amazon EKS), Azure offers Azure Kubernetes Service (Azure AKS), and Google Cloud offers Google Kubernetes Engine.

Some DR deployments make use of an architecture known as pilot light, where only the most critical applications and services are active at the DR site. When a failover is performed, the rest of the

infrastructure can be instantiated on demand, which can dramatically reduce the costs associated with the DR environment for regular day-to-day operation. AWS CloudFormation, Google Cloud Deployment Manager, and Azure Resource Manager make it possible to recreate compute and other cloud resources from a predefined template.

Network

How traffic is managed to the primary and secondary DR site.

In event that a failover is required, client hosts and applications must be able to automatically find the active site that is hosting the services they need to access. This is usually performed through DNS, which allows a network name to be repointed to a different resource without requiring any client-side changes. Amazon Route 53, Azure DNS, and Google Domains and Cloud DNS can be used to manually failover services to a DR site when this is required. Amazon Traffic Flow, Google Traffic Director, and Azure Traffic Manager take this a step further, allowing for automatic failover when the primary site has been deemed to be unhealthy.

Storage

The repository for all the data, optimized for usage and costs.

AWS, Google Cloud, and Azure provide a variety of data storage solutions, such as managed file services, block-level iSCSI devices, and low-cost, highly durable object storage. Some of these services provide redundancy within an availability zone, such as block-based Amazon Elastic Block Store (Amazon EBS), Google Persistent Disks, and Azure disks. Amazon EFS, Google Filestore, and Azure Files offer cross-zone availability. Availability can also span across entire regions with object-based Amazon S3, Google Cloud Storage, and Azure Blob.


Each solution in use at the primary site would need to be catered for separately and could require end-user administrators to set up additional processes and workflows. For example, Amazon EC2 compute instances using Amazon EBS would need the data being stored at the primary site available at the DR site as well. Amazon EBS snapshots could be used to create a solution, however the actual failover and failback processes would need to be manually developed and tested. That can be a difficult, risky, and costly process to carry out and maintain. In the next section we will look at how NetApp's Cloud Volumes ONTAP helps solve those problems out of the box.

Cloud-Based DR with Cloud Volumes ONTAP

Cloud Volumes ONTAP is NetApp's solution for enterprise data management in Google Cloud, AWS, and Azure. Building on native cloud compute and storage resources, Cloud Volumes ONTAP offers a wide variety of data storage services from a single platform, including NFS, SMB/CIFS with Active Directory integration, and iSCSI. Any and all of this data can be efficiently replicated from on-premises NetApp ONTAP systems, or from Google Cloud, AWS, or Azure-based deployments of Cloud Volumes ONTAP, using NetApp's replication technology, [SnapMirror®](#).

Cloud Volumes ONTAP provides SnapMirror as a solution for block-level data replication that keeps the destination up to date through incremental updates. SnapMirror uses NetApp Snapshot™ technology as the basis of its replication process, which first creates full copy of the data and subsequently replicates only the delta data, saving on transfer and storage costs. Users can specify a synchronization schedule ranging from few minutes to hours, at which frequency data changes from the source will be transferred over to the copy. Creating a new SnapMirror relationship is very easy: simply drag and drop the source system onto the destination in [NetApp Cloud Manager](#). This is a process that the Cloud Manager wizard will walk you through from start to finish. These are all central parts of the NetApp data fabric vision.

Cloud Manager is the web-based UI used for deploying Cloud Volumes ONTAP and managing hybrid and multi-cloud ONTAP storage environments. Existing on-premises and cloud-based ONTAP deployments can be discovered and added to the main dashboard, making it possible to set up replication relationships in any direction. Cloud Manager gives users the ability to failover data storage to the SnapMirror destination, as well as facilitating efficient failback to the source through a reverse re-synchronization operation.



Cloud Volumes ONTAP Features for DR

Storage efficiencies

Reducing overall storage space cuts DR costs

Storage tiering

Stores DR data on inexpensive object storage, automatically shifting it to performant block-based disks only when needed, reducing costs

SnapMirror

Data replication technology keeps DR sites up to date

FlexClone data clones

For fast and space-efficient DR testing

Cloud Manager

Easy management of all primary and DR systems

Benefits of using Cloud Volumes ONTAP

Reliable Data Protection

NetApp Snapshot™ technology is used to create instant, space efficient, backups of storage volumes of any size so they can be instantly restored.

Cloud Volumes ONTAP provides reliable data protection in the cloud within an Availability Zone, or across Availability Zones using the dual-node Cloud Volumes ONTAP High Availability configuration which guarantees RPO=0 and RTO<60 seconds workload SLAs.

Cross-region data protection is provided using SnapMirror replication. These comprehensive data protection capabilities are easily accessible from the Cloud Manager UI, which reduces the complexity of protecting cloud, hybrid cloud, and multicloud storage environments.

Cost Efficiency: Save Space, Save Costs


When using Cloud Volumes ONTAP, storage space requirements can be significantly reduced, in some cases by as much as 70%, through the use of built-in ONTAP technologies such as thin provisioning, data compression and deduplication. These storage efficiency features are applied transparently at the block-level, and so require no changes to client applications. In fact, SnapMirror replicates data in its compressed and deduplicated form, improving the speed at which transfers complete and reducing network bandwidth usage and costs.

Data tiering is another compelling storage efficiency feature provided by Cloud Volumes ONTAP that automatically and seamlessly shifts data between performance and capacity tiers as required. The capacity tier uses object storage on Amazon S3, Google Cloud Storage, or Azure Blob, which is extremely cost effective for data that is not currently in active use, such as that of a DR environment. Cloud Volumes ONTAP provides fast on-demand access to this data by automatically bringing it back into the performance tier on Amazon EBS,

Google Persistent Disk, or Azure disks when it needs to be accessed, such as in a DR scenario. As with syncs, SnapMirror integrates with data tiering by sending data received at the destination directly to the capacity tier.


Seamless Failover and Failback

When a disaster actually occurs, storage administrators need a quick, easy, and reliable process for bringing storage online at a DR site, and SnapMirror provides this through intrinsic support for failing over to destination volumes. If the primary site is later recovered successfully, the new data created in the DR storage volumes can be efficiently synchronized back to the source volumes, which enables the normal flow of data replication between source and destination to be re-established without requiring a full baseline copy of the data to be copied over. Cloud Manager provides an easy-to-use graphical user interface for performing these failover and failback operations.



Thin Provisioning

Allocates storage only as it needs to be used, not ahead of time.




Compression

Compresses block groups to reduce the amount of storage space used.



Deduplication

Reduces storage space by automatically removing duplicate blocks.



Compaction

Consolidates data from blocks that are not full, which drives up storage utilization.



Efficient and Instant Testing Environment through Data Cloning

Cloud Volumes ONTAP helps testing and development of the DR environments using NetApp FlexClone® technology. FlexClone creates instant, zero-capacity cost, writable clones of data volumes of any size which can be used to execute DR platform test suites.

Orchestration and Automation with a Click

Cloud Manager provides a modern, easy-to-use GUI interface for managing Cloud Volumes ONTAP, which includes setting up SnapMirror replication policies and creating FlexClone volumes. All of these tasks can also be performed through Cloud Manager's RESTful API, which allows for them to be automated or performed as part of a wider disaster recovery orchestration plan.


Conclusion

DR environments are crucial for ensuring data protection and the continued operation of software applications and services when physical infrastructure has been seriously compromised, such as from server failures, power outages, security threats, and natural disasters, to name a few. As shown in this guidebook, Cloud Volumes ONTAP provides a reliable, cost-effective, and flexible solution for both on-premises and cloud-based ONTAP storage environments to leverage AWS, Google Cloud, and Azure for disaster recovery.

Find out more about the Cloud Volumes ONTAP enterprise [DR solution](#) and learn from our customers how they leveraged [Cloud Volumes ONTAP for DR in these case studies](#).

**Start a free
30-day trial
today with
Cloud Volumes
ONTAP**

Start now





Not many partnerships are as tried-and-true as ours.

No need to start from scratch when you choose ePlus and NetApp for your data management and integrity needs. From ransomware resiliency and data protection to maximizing the value of hybrid cloud and integrating new technologies, we'll bring broad experience mixed with strategic vision that will help you build upon and extend your success.

masterclass.eplus.com

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NA-000-0121