E-BOOK

How it works: Disaster recovery in Google Cloud with Cloud Volumes ONTAP

NetApp | Google Cloud

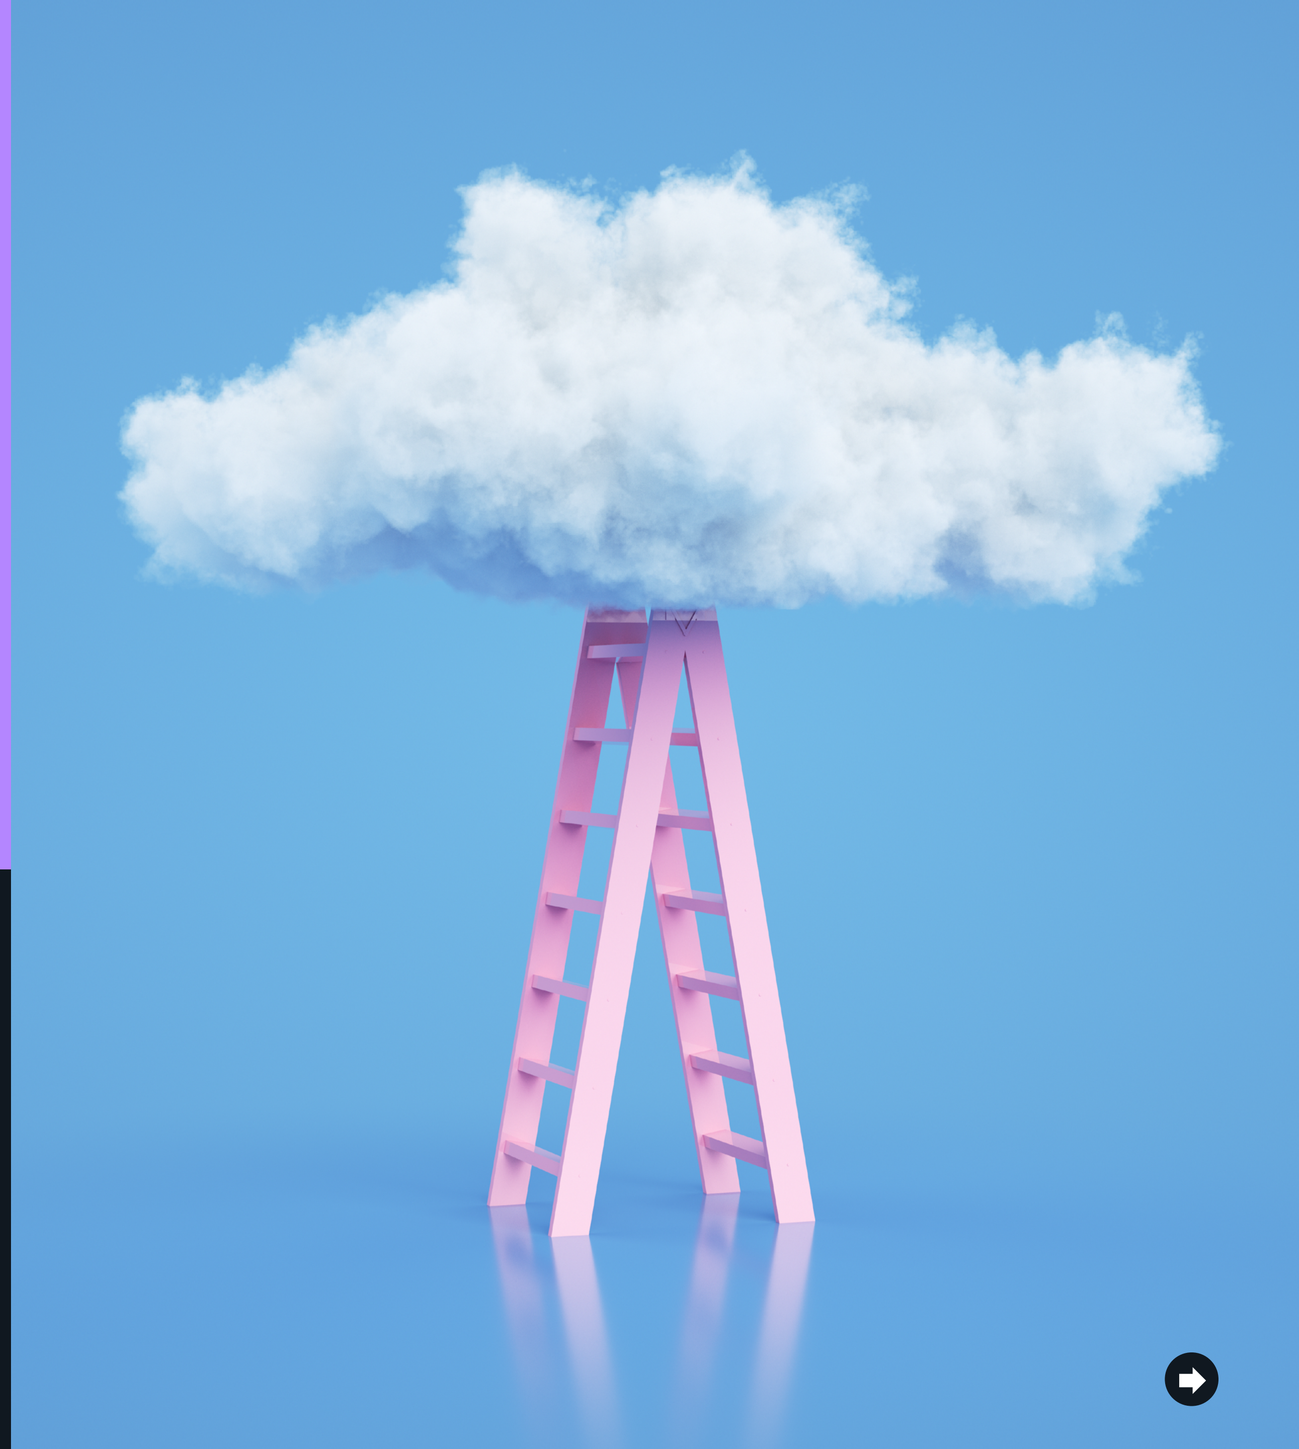

Contents

Introduction

Organizations are facing an increasing number of risks to business continuity. Alongside traditional triggers like natural disaster and equipment failure, there's a new category of risk: Cyberattacks, and specifically ransomware, have taken center stage.

As criminal organizations continue to refine their tools and tactics; companies must be equally purposeful in their approach to disaster recovery.

Storage infrastructures and application architectures today are more complex and more frequently hybridized, which magnifies the risk of data loss and lengthens recovery times. And that complexity can mean an overwhelming expense.

All enterprise platforms require a disaster recovery (DR) environment to ensure that business-critical services and applications can continue to operate in the event of a major infrastructure outage or ransomware attack. Companies must be able to bring services back online rapidly, which requires having standby failover systems in a secondary location.

Whether you're already up and running in Google Cloud, or just finding your way there, NetApp brings disaster recovery and data management into a unified control plane. With our storage and data protection services, you can reduce outages by 99% and drastically accelerate recovery time objectives in NetApp® ONTAP® environments.

Cloud Volumes ONTAP provides comprehensive data protection within a zone, across zones, or across regions, all managed from a single control plane. NetApp enables you to automatically fail over your operations to secondary (and even tertiary) sites, and later recover and fail back to your primary copy reliably, while paying less for data storage.

Now how is that possible?

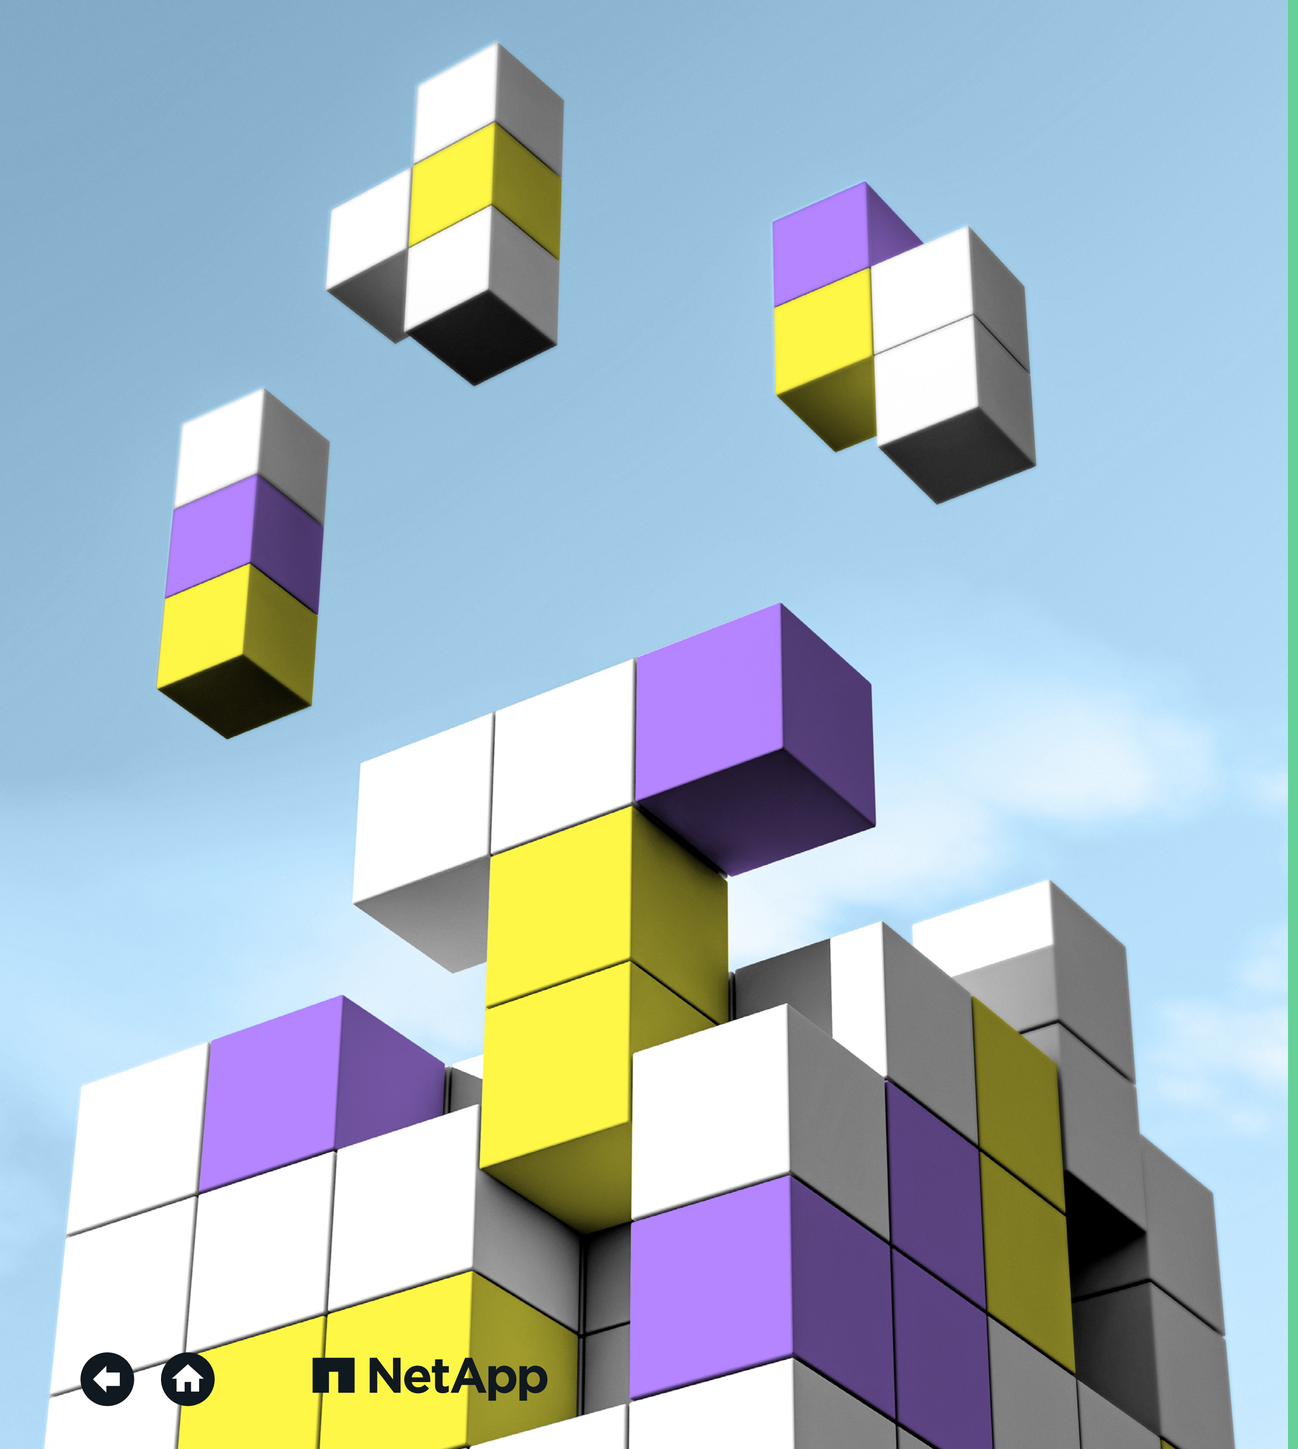

Why you need a DR solution

Data is constantly growing, and organizations must be poised to adapt to increasing demands on their primary and secondary systems.

Deploying DR environments is challenging for most organizations because of the need 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 most of the time.

The cloud can help lower the barrier to entry by providing scalable infrastructure-as-a-service solutions that can be used to build DR environments. After building out all DR services, the challenge becomes synchronizing data from the production environment and keeping it synchronized going forward.

This e-book examines in detail the challenges involved in setting up a DR environment, discusses the available services in Google Cloud that can be used to build DR solutions, and looks at how NetApp Cloud Volumes ONTAP provides cost-effective enterprise-grade support for data replication and disaster recovery for both existing on-premises NetApp storage systems and for cloud- based deployments on Google Cloud.

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. You need multiple complete, up-to-date copies of your data.

The infrastructure necessary for effective disaster recovery must be planned for and implemented in advance. 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 that it has its own metric: the recovery time objective. RTO is the time it takes to bring services back online after a DR event. Reducing this interval as much as possible correspondingly minimizes loss of revenue.

Setting up a replica DR environment also requires having an independent, up-to-date copy of all enterprise data, including database systems, file services, and iSCSI storage. Because data in the production environment is constantly changing, these changes must be reflected in the DR site to prevent data loss. In some cases, no data loss can be tolerated, which means that your recovery point objective (RPO) will equal zero.

Recovery point objective
The maximum acceptable amount of data, as measured by time, that a company can lose during a disaster event.

Recovery time objective
The amount of time it takes to return to normal operation after a disaster event.

Test your limits

A fully functioning disaster recovery site is useful for both planned and unplanned outages. For example, a DR site can be made live to perform updates to the primary production environment. This scenario requires failover to the DR site and then failback after the updates have been performed. The same scenario applies to unplanned outages: Failover and failback can also be used to make the DR site live after a disaster and to restore services back to the primary site. Being able to perform these operations quickly, or even automatically, can save hundreds of thousands, or even millions, of dollars as a result of averted data loss.

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, so you can quickly verify whether it's operating correctly. However, if the DR site is left unchecked for a long period of time, it may fail to come online when it's needed. Therefore it's essential to perform regular testing of DR services and ensure that they are functioning as expected.

Don't sit idly by

Because your DR site is generally used to store copies of your data, it sits idle most of the time. To reduce the cost of storing backup data, DR systems can be used for peripheral requirements, such as read-only reporting or setting up software development test environments.

In a hybrid cloud environment, using NetApp Cloud Volumes ONTAP or Cloud Volumes Service for Google Cloud to tier your data, the cold data of your DR systems can be kept in a less expensive performance tier to reduce costs.

DR challenges at a glance

Scheduled syncs
Data must be synchronized efficiently and regularly to the secondary location to keep it up to date.

Failover and 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
Make sure that DR systems work as expected.

Controlling costs
DR compute and storage resources should remain cost effective even though the system is not normally in active use.

Building cloud based DR environments in Google Cloud

Disaster recovery environments must implement redundancy at the compute, network, and storage layers. This section looks at how this redundancy can be accomplished using Google Cloud services.

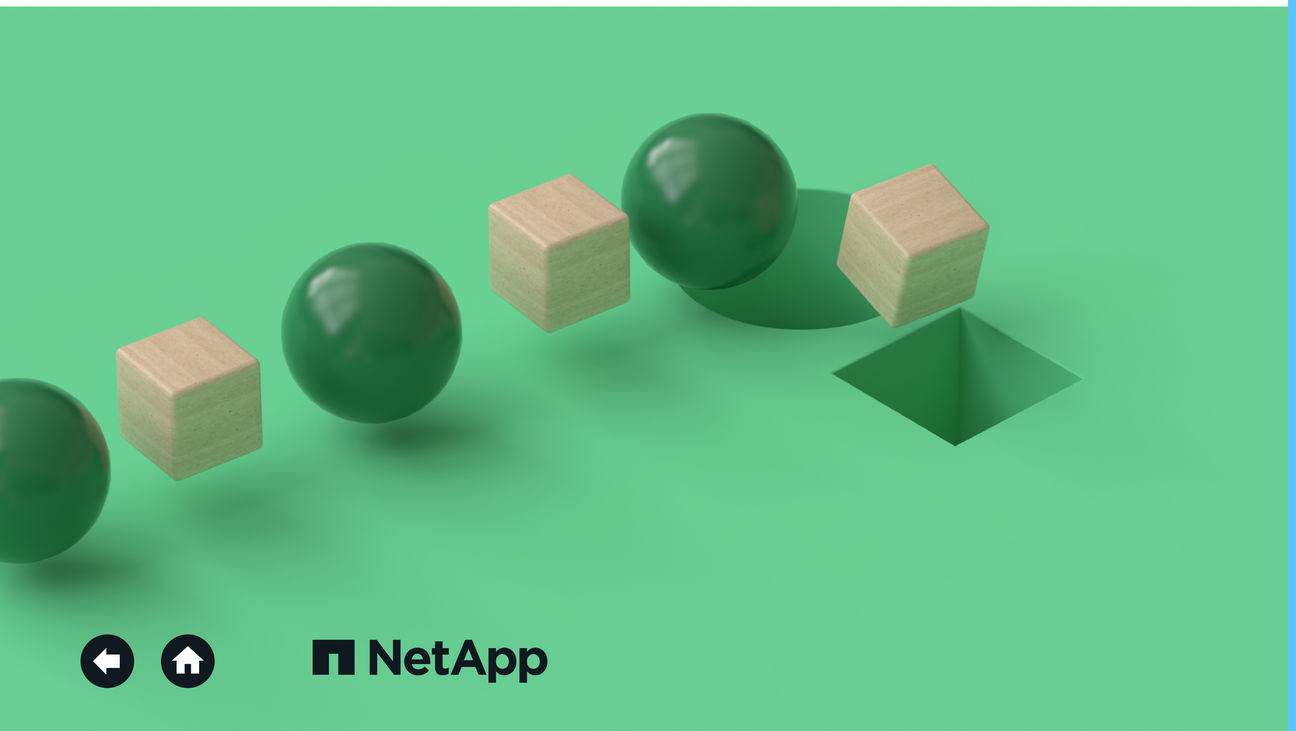

Cloud DR environment building blocks

Compute

The CPU processing power that runs applications.
The Google Cloud Compute Engine provides 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, Google Cloud offers Google Cloud Run plus native Kubernetes services, known as Google Kubernetes Engine (GKE) and Google Cloud Anthos.

Some DR deployments use a “pilot light” architecture, whereby 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. Google Cloud Deployment Manager makes it possible to recreate compute and other cloud resources from a predefined template.

Cloud DR environment building blocks

Network

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

If 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 action is usually performed through DNS, which allows a network name to be repointed to a different resource without requiring any client-side changes. Google Cloud DNS can be used to manually fail over services to a DR site when necessary. Google Cloud Traffic Director takes this process a step further, allowing automatic failover when the primary site is deemed to be unhealthy.

Storage

The repository for all the data, optimized for usage and costs.
Google Cloud provides 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 a single zone or an entire region, as with Google Persistent Disk, or across multiple regions, with services such as Google Filestore and Google Cloud Storage.

Each solution in use at the primary site needs to be managed separately, and that could require end-user administrators to set up additional processes and workflows. For example, Cloud Compute Engine instances using Cloud Persistent Disk would require the data stored at the primary site to be available at the DR site as well. Cloud Persistent Disk snapshots could be used to create a solution for this situation; however, the actual failover and failback processes would need to be manually developed and tested. That can be a difficult, risky, time-consuming, and costly process to carry out and maintain.

The next section looks at how NetApp Cloud Volumes ONTAP helps solve those problems.

Cloud-based DR with Cloud Volumes ONTAP

Cloud Volumes ONTAP is NetApp's solution for enterprise data management in Google Cloud. 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. All of this data can be efficiently replicated from on-premises NetApp ONTAP systems, or from Google Cloud-based deployments of Cloud Volumes ONTAP, using NetApp SnapMirror® replication technology.

Cloud Volumes ONTAP provides SnapMirror as a solution for block-level data replication that keeps the destination up to date through incremental updates. For example, users can set a synchronization schedule of every minute or every hour, to specify when data changes from the source are transferred over.

Creating a new SnapMirror relationship is easy: Simply drag and drop the source system onto the destination. The SnapMirror wizard walks you through this process from start to finish. NetApp BlueXP™ lets you manage your data estate across hybrid multicloud environments. From the BlueXP unified control plane, you can deploy Cloud Volumes ONTAP and manage hybrid cloud ONTAP storage environments. BlueXP delivers a simplified hybrid cloud and cloud-native experience for storage and data protection services across on-premises and Google Cloud environments.

Existing on-premises and cloud-based ONTAP deployments, including Cloud Volumes Service, can be discovered and added to the main dashboard, making it possible to set up replication relationships in any direction. BlueXP gives you the ability to fail over data storage to a SnapMirror destination, and it facilitates efficient failback to the source through a reverse resynchronization operation.

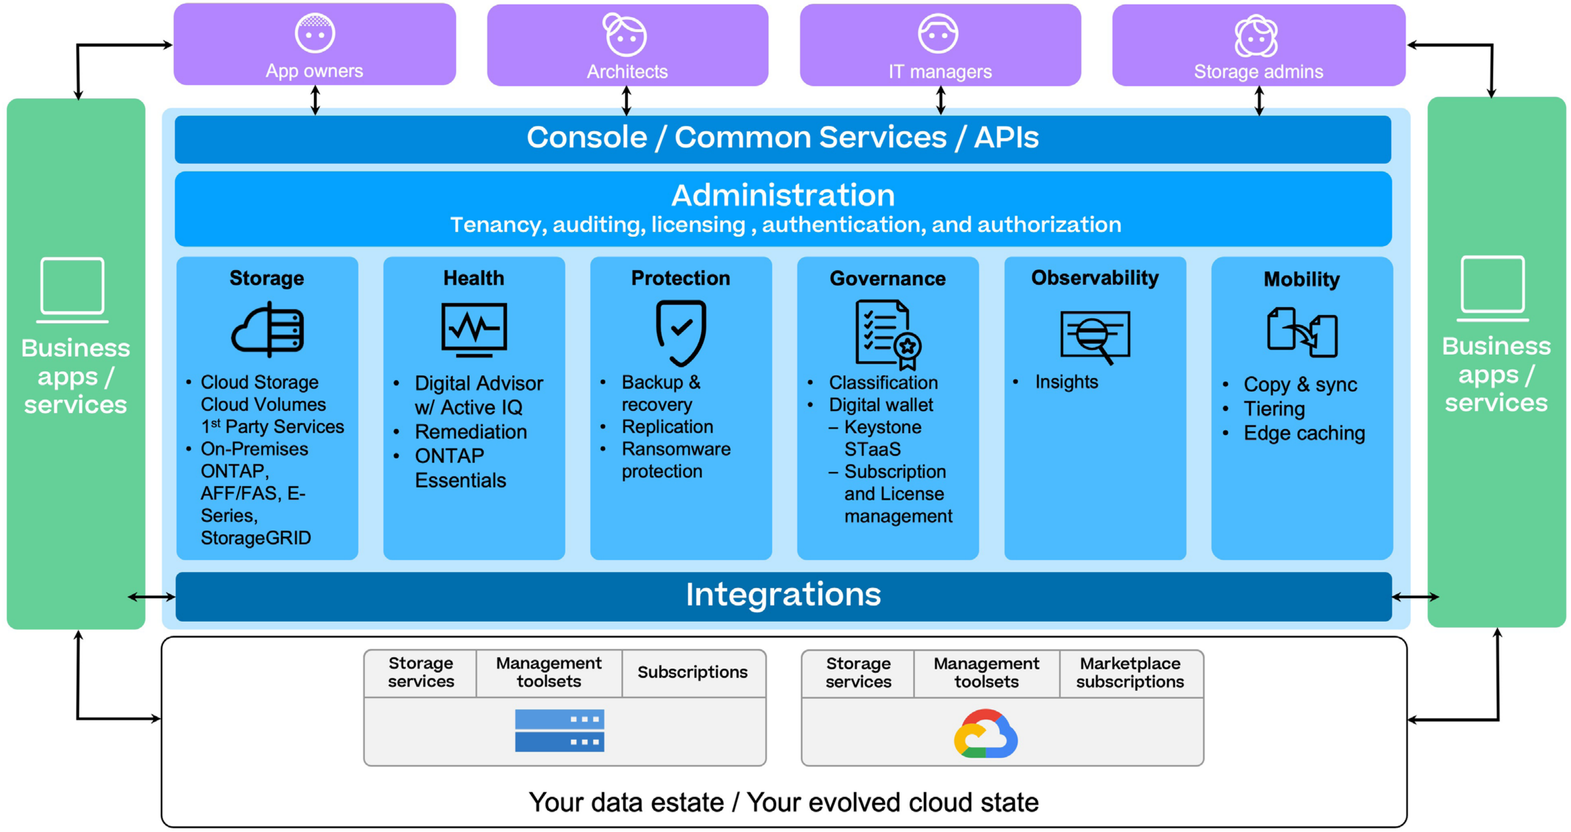

Cloud Volumes ONTAP features for DR

Storage tiering
Automatically shift DR environments to performant disks only when needed, reducing costs.

SnapMirror
Data replication technology keeps DR sites up to date.

FlexClone data clones
Thin-cloning technology for fast and space-efficient DR testing.

BlueXP
Easily manage all primary and DR systems.

Storage efficiencies
Reduce overall storage space to cut DR costs.

Benefits of using Cloud Volumes ONTAP

Reliable data protection

Cloud Volumes ONTAP provides reliable data protection in the cloud within a zone, across zones, or across regions by using SnapMirror replication. This comprehensive capability is easily accessible from BlueXP, which reduces the complexity of protecting cloud, hybrid cloud, and multicloud storage environments.

Cost efficiency: Save space, save costs

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

Data tiering is another compelling storage efficiency feature of Cloud Volumes ONTAP that automatically and seamlessly shifts data between performance and capacity tiers as required. The capacity tier uses Google Cloud Storage object storage, 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 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.

Benefits of using Cloud Volumes ONTAP

Seamless failover and failback

When a disaster occurs, storage administrators need a rapid, easy, and reliable process for bringing storage online at a DR site. SnapMirror provides that ability 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 reestablished without requiring a full baseline copy of the data to be copied over. BlueXP offers 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 improves storage utilization.

Efficient and instant testing environment through data cloning

Software applications are continuously being developed and updated, which necessitates releasing new versions to production, as well as to DR environments. Because a DR site is not normally active, this creates a risk of latent problems in the deployment that may not manifest themselves until after a failover, which would seriously affect the time taken to recover from an outage. Cloud Volumes ONTAP helps tackle these issues through its NetApp FlexClone® functionality, which can be used to instantly create zero-cost, writable clones of a SnapMirror destination volume of any size. These volume clones can be used to execute DR platform test suites that mutate the data they operate on without interrupting the active replication of data from source systems. Clones can also be used to gain greater benefit from idle resources at the DR site by providing data for software development test environments or DevOps continuous integration and continuous deployment (CI/ CD) pipelines.

Orchestration and automation with one click

With BlueXP you can deploy, discover, manage, and optimize data and infrastructure via a unified control plane with flexible consumption capabilities. BlueXP provides a modern, easy-to-use GUI interface for managing Cloud Volumes ONTAP, which includes setting up SnapMirror replication and creating FlexClone volumes. All of these tasks can also be performed through the RESTful API, which allows the tasks to be automated or performed as part of a broader disaster recovery orchestration plan.

Your next steps to cyber resilience

DR is daunting, but it's crucial for data protection and the continued operation of software applications and services when infrastructure has been seriously compromised by server failures, power outages, security threats such as ransomware, and natural disaster. Cloud Volumes ONTAP provides a reliable, cost-effective, and flexible solution for both on-premises and cloud-based ONTAP storage environments to leverage Google Cloud for disaster recovery—whether from ransomware, fire, flood, or data misuse.

Learn how to become more cyber resilient with NetApp and Google Cloud, or speak to a specialist to learn how NetApp can help your company build a foolproof data protection strategy.

Watch the brief-but-powerful INSIGHT session to learn your three first steps to cyber resilience with NetApp for disaster recovery.

About NetApp
In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services, and applications to the right people—anytime, anywhere.

NetApp

+1 877 263 8277