



Technical Report

NetApp HCI Data Protection

Erik Kemp, NetApp
May 2020 | TR-4641

Abstract

This document provides an overview of the native and partner integrated data protection capabilities of NetApp® HCI.

TABLE OF CONTENTS

1	Introduction	4
2	Data Protection Features	4
2.1	Enterprise Reliability and High Availability	4
2.2	Protection Domains	6
2.3	Real-Time Replication	6
2.4	Snapshot Copies and Clones	6
3	Disaster Recovery and Business Continuity	9
3.1	Disaster Recovery Failover and Failback—HCI to HCI	10
3.2	Disaster Recovery with SnapMirror	10
3.3	Backup of and Restore of ONTAP Select to Cloud Volumes ONTAP	12
3.4	VMware Site Recovery Manager	13
3.5	HCI to HCI with VMware SRM	13
4	Integrated Backup and Restore	14
5	Data Fabric	15
5.1	NetApp HCI Data Fabric Integration Points	15
6	StorageGRID	16
6.1	Use Cases	16
7	Partnered Solutions	17
7.1	Veeam	17
7.2	VMware Site Recovery Manager	18
7.3	Commvault	20
8	Data Protection for Containers	21
8.1	Kubernetes Pods, Volumes, and Projects	21
8.2	Kubernetes Container Storage Interface Volume Cloning	22
8.3	Trident for Kubernetes	22
8.4	Third-Party Data Protection Tools for Containers	24
9	Conclusion	25
	Where to Find Additional Information	25
	Version History	25

LIST OF TABLES

Table 1)	High-availability event impact	5
----------	--------------------------------	---

Table 2) Disaster recovery characteristics by feature.....	10
Table 3) Primary use cases.....	19

LIST OF FIGURES

Figure 1) NetApp HCI data protection overview.	4
Figure 2) NetApp HCI SnapMirror to Cloud Volumes ONTAP.....	12
Figure 3) NetApp HCI SnapMirror to ONTAP.....	13
Figure 4) NetApp HCI integrated backup and restore.	14
Figure 5) NetApp Fabric Orchestrator.	15
Figure 6) StorageGRID on NetApp HCI.....	16
Figure 7) Element Snapshot integration on NetApp HCI with Veeam Backup & Replication.....	18
Figure 8) VMware vCenter SRM with NetApp HCI.	19
Figure 9) VMware SRM with NetApp HCI and Element SRA.	19
Figure 10) IntelliSnap with NetApp HCI overview.....	21

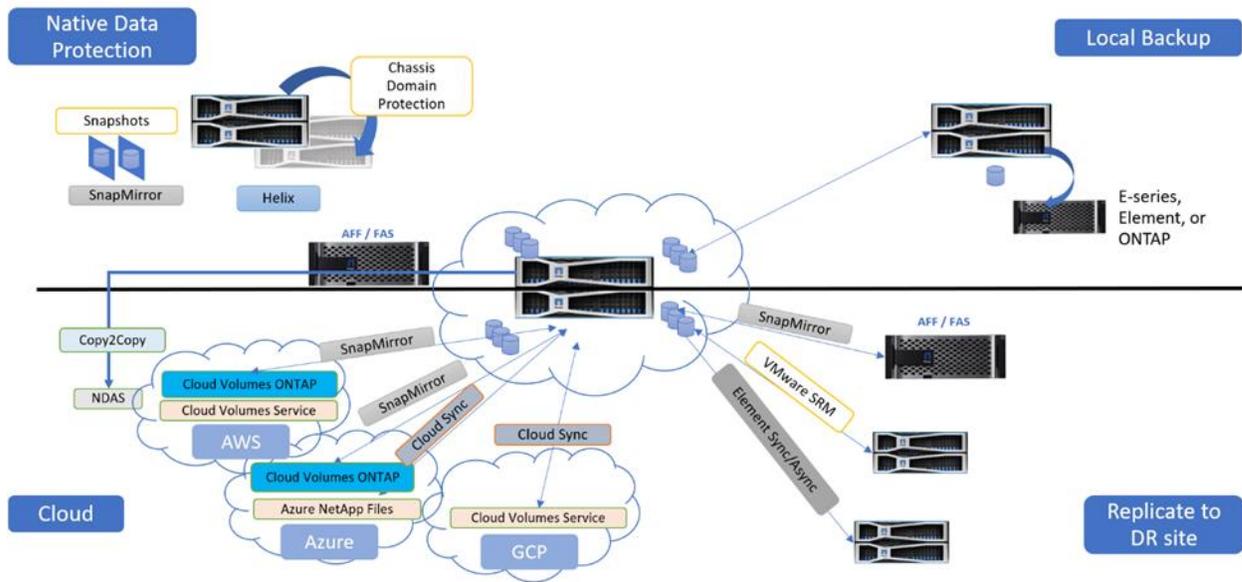
1 Introduction

NetApp HCI Element® software is designed to meet the high-availability standards and data protection that are expected from enterprise-class arrays. It is designed to be ready for tomorrow's applications and data demands. As protecting data becomes ever more important, the native data protection features found in Element are essential to all infrastructures.

NetApp HCI customers can run multiple applications with guaranteed performance and industry-leading quality of service (QoS) to confidently deploy resources across the entire data center. This architecture allows you to deploy your infrastructure with simplified management and independently scaled compute and storage nodes, further ensuring that resources are deployed when and where they need to be. NetApp HCI is data-fabric ready for easy access to all your data across public, private, or hybrid clouds. By moving to NetApp HCI, IT organizations can transform their data center, driving operational efficiencies and reducing costs.

The data fabric is a software-defined approach from NetApp for data management that enables businesses to connect disparate data management and storage resources, whether on-premises or in the public cloud providers. NetApp HCI can streamline data management between the locations for enhanced data portability, visibility, and protection (Figure 1).

Figure 1) NetApp HCI data protection overview.



2 Data Protection Features

NetApp Element software enables you to protect your data in a variety of ways with capabilities such as the creation of Snapshot copies for individual volumes or groups of volumes. You can also perform replication between clusters and volumes running Element and replicate data to NetApp ONTAP® systems.

2.1 Enterprise Reliability and High Availability

NetApp HCI Helix data protection is a RAID-less data protection solution designed to maintain both data availability and performance regardless of the failure condition. Helix data protection is a distributed replication algorithm that spreads at least two redundant copies of data across all drives in the system.

This approach allows the system to absorb multiple failures across all levels of the storage solution while maintaining data redundancy and QoS settings.

Failure Prevention

The NetApp HCI shared-nothing storage architecture has no single point of failure and eliminates common failure scenarios in which a loss of availability or data might occur. NetApp HCI also benefits from a fully redundant storage architecture. NetApp HCI automatically rebuilds redundant data across the remaining storage nodes in minutes to maintain high availability with a minimal effect on performance upon disk or controller failures.

Data Protection

The NetApp HCI clustered architecture allows nondisruptive software upgrades on a rolling node-by-node basis. Upgrades can be performed during production hours with little to no effect on workload performance.

NetApp HCI has a self-healing architecture. A NetApp HCI system can recover from failures in minutes and is fully automated. In a failure event, each drive in the system redistributes a small percentage of its data (usually 1% to 2%) in parallel to the free space on all remaining drives. Failure recovery requires no operator intervention, eliminating the fire drills common with traditional RAID-based architectures.

NetApp HCI automatically rebuilds redundant data across the remaining nodes in minutes to maintain high availability with a minimal effect on performance. In contrast, RAID-based systems suffer significant performance degradation upon disk or controller failures, taking hours or days to restore redundancy and equally long to replace failed hardware.

No matter the failure mode—drive, node, backplane, network failure, or software failure—the recovery process is the same. Because the recovery workload is distributed across all nodes in the cluster, redundancy is restored quickly, and no single node (or application workload) takes a performance hit. The more nodes in the cluster, the faster recovery occurs and the lower the overall effect on the storage system.

NetApp provides the following services to further protect your HCI storage system:

- Secure assist
- 24/7/365 worldwide availability
- Expert tier-3 support engineers
- NetApp Active IQ monitoring

Table 1 shows various events that do not cause downtime and their associated effects on the NetApp HCI cluster.

Table 1) High-availability event impact.

Event	Causes Downtime	Impact
SSD failure	No	Fully automated
Node failure	No	Fully automated
HW upgrades or replacement	No	Seamless
Hardware expansion	No	Instant resource availability
Software upgrades	No	Rolling or online

2.2 Protection Domains

NetApp Element software supports Protection Domain functionality, which optimizes data layout on storage nodes for optimal data availability. A Protection Domain is a node, or a set of nodes grouped together such that any part or even all the domain can fail, without affecting data availability. To use this feature, you should split storage capacity evenly across three or more NetApp H-series chassis for optimal storage reliability. With three chassis, the storage cluster automatically enables Protection Domains.

A Protection Domain layout assigns each node to a specific Protection Domain. Two different Protection Domain layouts are supported, and they are referred to as Protection Domain levels.

- At the node level, each node is in its own Protection Domain. This was the only level supported before Element 11.0.
- At the chassis level, only nodes that share a chassis are in the same Protection Domain.

The chassis-level Protection Domain Layout is automatically determined from the hardware when the node is added to the cluster.

Note: If each node in a cluster is in a separate chassis, then these two levels are functionally identical.

Protection Domains and the vCenter Plug-In

The following Protection Domain monitoring parameters are available for the cluster:

- **Selected Monitoring Level.** The Protection Domain resiliency levels selected by the user: either chassis or node. Green indicates that the cluster is capable of the selected monitoring level. Red indicates that the cluster cannot perform the selected monitoring level and corrective action is needed.
- **Remaining Block Capacity.** Indicates the remaining block capacity that can be used while maintaining the selected resiliency level.
- **Metadata Capacity.** Indicates if there is enough metadata capacity to heal from failure while also maintaining undisrupted data availability. Normal (green) indicates that the cluster has enough metadata to maintain the selected monitoring level. Full (red) indicates that the cluster is no longer capable of the selected monitoring level and corrective action is needed.

2.3 Real-Time Replication

With the built-in synchronous, asynchronous, and snapshot replication of NetApp HCI, you can quickly copy data between multiple sites, regardless of where your clusters sit. Synchronous replication confirms data synchronization for your mission-critical data in near real time. Asynchronous replication protects against hardware failure and natural disaster incidents over long distances through bidirectional replication, allowing each replication partner to fail over or fail back. With snapshot replication, changed data is replicated at specific points of time to a remote cluster. Only those snapshots that are created on the source cluster are replicated. Active writes from the source volume are not. No additional licenses or support are required for these capabilities, so your costs are reduced. Plus, by finding additional uses for your remote copies, you can drive more revenue from your data. The simplified management of NetApp HCI facilitates flexible replication to multiple locations from a single cluster.

2.4 Snapshot Copies and Clones

NetApp HCI enables you to get the most out of the storage layer by allowing for the creation of instant replication-ready copies, zero-RTO snapshot backups and restores, and highly efficient multi-use clones. When combined with native NetApp HCI features such as per-volume QoS controls, these features enable advanced topologies with the economics of shared infrastructure. Together, these features unlock architectural patterns and flexibility that no other storage architecture can provide.

Using NetApp HCI volume snapshots is a simple process that can drastically reduce TCO for backups due to the simplification of maintenance tasks and the minimization of storage overhead. Backups on NetApp HCI can be accompanied with no system effects or maintenance window. In addition, backups can be automated using the NetApp HCI API.

Similarly, creating a clone with NetApp HCI is as easy as taking a snapshot backup. By using the clone-from-snapshot functionality to copy the volume metadata to a different account, you can then mount the volume to another instance mount point. Because NetApp HCI deduplicates system wide, the majority of the initial data does not require additional capacity utilization. The ability to clone from snapshots reduces both capital (storage) and operational (administration) outlays, driving TCO from both sides.

Snapshot Copies

There is no performance effect from a snapshot operation. Snapshot copies are crash-consistent, point-in-time references of a volume. If you want to make a Snapshot copy of a live file system, consider suspending writes on the client beforehand.

You can replicate Snapshot copies to a remote NetApp HCI cluster and use them as a backup copy for the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot. You can also create a clone of a volume from a replicated snapshot.

You can create a new volume from a snapshot of a volume. When you create a volume that way, the system uses the snapshot information to clone a new volume by using the data that is contained on the volume at the time that the snapshot was created. This process also stores information about other snapshots of the volume in the newly created volume.

You can use the integrated backup feature to back up a volume snapshot. You can back up snapshots from a NetApp HCI cluster to an external object store or to another NetApp HCI cluster. When you back up a snapshot to an external object store, you must have a connection to the object store that allows read/write operations.

Using Volume Snapshots for Data Protection

A volume snapshot is a point-in-time copy of a volume. You can take a snapshot of a volume and use the snapshot later if you need to roll a volume back to the state it was in at the time the snapshot was created.

Snapshots are like volume clones. However, snapshots are simply replicas of volume metadata, so you cannot mount or write to them. Creating a volume snapshot also takes only a small amount of system resources and space, which makes snapshot creation faster than cloning.

You can take a snapshot of an individual volume or a set of volumes. Optionally, you can replicate snapshots to a remote cluster and use them as a backup copy of the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot. Alternatively, you can create a clone of a volume from a replicated snapshot.

Snapshot Replication Rules

The following rules govern the behavior of snapshot replication:

- Only snapshots that have the attribute `enableRemoteReplication=true` are replicated.
- Snapshot replication only occurs when the replication state of the volume is Active, unless the mode is SnapshotsOnly.
- The target replicates snapshots in the order of newest to oldest.
- The target does not replicate a snapshot if its volume already has that snapshot.
- The target replicates a snapshot if another volume has that snapshot.
- Users can pause or cancel in-progress replication.

- Users can delete snapshots or modify snapshot attributes on either side without affecting the remote snapshot.

Group Snapshots

You can create a group snapshot of a related set of volumes to preserve a point-in-time copy of the metadata for each volume. You can use the group snapshot in the future as a backup or rollback to restore the state of the group of volumes to a previous state.

Group Snapshot Replication Rules

The following rules govern the behavior of group snapshot replication:

- If a group snapshot has the attribute `enableRemoteReplication=true`, then the target cluster creates the group snapshot after it completes replication of all constituent members.
- The group snapshot and all constituent snapshots must have the attribute `enableRemoteReplication=true`. Otherwise, only those constituent snapshots with `enableRemoteReplication=true` are replicated.
- The source volumes involved in the group snapshot must be correctly paired with target volumes sharing the same cluster pair. Remote replication must not be paused. Otherwise, only the constituent snapshots on valid pairings are replicated.
- Group snapshot replication is prioritized over replication of individual snapshots.
- If a user deletes a constituent snapshot on the source after it has been replicated to the target, but group snapshot replication is still in progress, then the target cluster creates the group snapshot over all the replicated constituent snapshots. The same group snapshot UUID exists on the source and the target, but they will have different contents.
- If the user deletes a replicating constituent snapshot on the source, then the snapshot is removed from the group on the source and the incomplete snapshot is cleaned up on the target. The group is replicated, and both the source and the target have a group snapshot with the same contents.
- If the user deletes a constituent snapshot on the target while replication of a group snapshot is in progress, then the target cluster creates the group snapshot with the other replicated constituent snapshots. The same group snapshot UUID exists on the source and the target, but they will have different contents.
- If the user deletes a constituent snapshot on either the source or target after replication of a group snapshot is complete, then that cluster removes the snapshot from the group snapshot. The same group snapshot UUID exists on the source and the target, but they will have different contents.
- If a group snapshot is replicated to the target cluster and then a constituent source volume is paired to a different target volume in the same target cluster, then the target cluster replicates the constituent snapshot and replaces the original constituent snapshot in the group snapshot.

Replication of Existing Group Snapshots

The Element software started supporting snapshot replication in version 8.0. At that time, if a group snapshot was created with `enableRemoteReplication=true`, then all constituent members were replicated. Each snapshot on the target cluster received the `groupSnapshotUUID` attribute from the source. However, the group snapshot object itself was not replicated. There was no automatic way to perform a group snapshot rollback on the target side.

The Element release 11.3 adds support for replicating the group snapshot object. This allows all group snapshot operations, such as rollback, to work on the target.

All existing group snapshots are automatically replicated from the source cluster to the target cluster after both clusters are upgraded to version 11.3 or later.

Snapshot Backup Operations

You can back up Element volume snapshots to external object stores that are compatible with Amazon S3. You can also back up Element snapshots to secondary object stores that are compatible with OpenStack Swift. Finally, you can back up volume snapshots residing on a NetApp HCI cluster to a remote NetApp HCI cluster.

Volume Snapshot Restore Operations

You can perform the following volume snapshot restore operation in NetApp HCI:

- Restore a volume from backup on an Amazon S3 object store
- Restore a volume from backup on an OpenStack Swift object store
- Restore a volume from backup on a NetApp HCI cluster

Clones Rules

The following rules govern the behavior of clones:

- There is no performance effect from the cloning operation.
- Clones have the same QoS level as the parent volume.
- Cloned volumes have no explicit link to their parent, so you are free to delete the parent volume at any time without affecting any clones.
- Volumes can be resized (grown or shrunk) during the cloning process. NetApp recommends not shrinking a volume. It is a best practice to create a smaller volume, copy the data over from the file system, and then delete the original volume.
- Cloned volumes are crash-consistent, point-in-time copies of a volume. If you want to make a snapshot of a live file system, consider first suspending writes on the client before making the copy.
- Only two active clone requests can be running at one time for a single volume. All subsequent clone requests are queued for later processing.
- Only a total of eight active volumes can be cloned at one time. All subsequent clone requests are queued for later processing.
- A clone is a readable and writable full copy of the original data, but it often takes little space due to deduplication. Clones also are point-in-time copies.
- Cloned volumes can be changed from read/write to read-only or can be set as locked.

The following characteristics are shared by snapshots and clones:

- Inherently thin, deduped, and compressed (like all data inside Element)
- Until block-data changes, both are simply metadata operations
- Very fast
- Data efficient

3 Disaster Recovery and Business Continuity

Table 2 describes the disaster recovery features of the various replication methods available with NetApp HCI.

Table 2) Disaster recovery characteristics by feature.

Feature	RPO	RTO	Retention	Functionality	Typical Use Case
Asynchronous replication	Seconds to minutes	Short	No limit	Writes are acknowledged only after they are written on the source cluster.	Recommended if there are high latencies (> 5ms) between the source cluster and the target cluster.
Synchronous replication	0	Short	No limit	Writes are acknowledged only after they are written on the target cluster.	Recommended when the RPO requirement is zero
Snapshot replication	5mins+	Minutes+	32 per volume	Snapshots are replicated at discrete intervals. New snapshots include new writes in their images.	Recommended when there is low bandwidth and a shared network between the source and the target. Use to revert a volume to a previous state.
Integrated Backup & Restore	5mins+	Minutes+	No limit	Volume level replication through HTTP/HTTPS	<ul style="list-style-type: none"> • Backing up a volume snapshot to an Amazon S3 object store. • Backing up a volume Snapshot copy to an OpenStack Swift object store. • Backing up a volume Snapshot copy to a NetApp HCI cluster.
NetApp SnapMirror®	5mins+	Minutes+	30 per volume	Snapshot replication between Element and ONTAP	Designed for failover from primary storage to secondary storage at a remote site.
Storage Replication Adaptor (SRA) for VMware Site Recovery Manager (SRM)	0 (Sync) seconds to minutes (Async)	User configurable	No limit	Element-based replication (asynchronous/synchronous)	<ul style="list-style-type: none"> • Virtual machine (VM) failover • Failback for disaster recovery • Business continuity between sites

3.1 Disaster Recovery Failover and Failback—HCI to HCI

For clusters running Element software, real-time replication enables the quick creation of remote copies of volume data. You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

3.2 Disaster Recovery with SnapMirror

Systems running NetApp Element software support SnapMirror functionality to copy and restore Snapshot copies with NetApp ONTAP systems.

Systems running Element can communicate directly with SnapMirror on ONTAP systems 9.3 or later. The NetApp Element API provides methods to enable SnapMirror functionality on clusters, volumes, and Snapshot copies. Additionally, the Element UI includes all necessary functionality to manage SnapMirror relationships between Element software and ONTAP systems. For more information, see the [NetApp SolidFire and Element Documentation Center page on this subject](#).

If your organization is ready to add NetApp HCI to your environment, you can still use your existing investments in NetApp AFF and FAS systems as secondary storage for multiple use cases. You can use AFF or FAS systems to protect against storage outages, with easy disaster recovery relationship setups. The storage-efficient transport mechanism transfers only changed data to the secondary site, which reduces recovery time and enables more frequent replication. Because SnapMirror stores data in its native format, it maximizes your investment in disaster recovery infrastructure, while your disaster recovery site copies can help accelerate development.

You can further leverage your data protection copies by creating zero-effect copies of the replicated data for development and testing. Data that is replicated from a SolidFire system to AFF or FAS can benefit from the rich data management capabilities of ONTAP. NetApp SnapMirror technology provides an agile, flexible, and secure way to transport data between storage systems, protecting your data and improving accessibility across the data fabric.

The primary use case is disaster recovery of NetApp HCI to ONTAP. Endpoints include ONTAP, ONTAP Select, and Cloud Volumes ONTAP. SnapMirror supports applications failing over to a secondary volume and then continuing operation. SnapMirror can also fail back applications to the primary location later. This capability is sometimes referred to as disaster recovery.

Block snapshots can be replicated from an Element source volume to an ONTAP destination volume and back.

NetApp converged infrastructure and NetApp HCI are interoperable. Various use cases have been developed for data movement between on-premises systems for NetApp converged infrastructure and NetApp HCI, including backup, disaster recovery failover and failback, and migration.

Disaster Recovery

In the event of a problem with a volume or cluster running NetApp Element software, use the SnapMirror functionality to break the relationship and failover to the destination volume. For more information, see [TR-4651: NetApp SolidFire SnapMirror Architecture and Configuration](#).

Backup and Restore—HCI Storage to Cloud Volumes ONTAP

Keeping local snapshots of volumes in your NetApp HCI deployment is the first line of protection against unexpected data loss. However, the cost of storing many idle snapshots on more expensive SSDs can quickly become prohibitive. NetApp SnapMirror allows you to transfer data from the on-premises Element cluster in your NetApp HCI environment to a remote ONTAP system in the public cloud. This relationship can be configured as a NetApp SnapVault® volume for archival purposes to store these snapshots for an extended period on Cloud Volumes ONTAP. This enables the removal of the snapshots on the source NetApp HCI system, resulting in saving storage space and infrastructure costs. After snapshots are moved to the cloud, the ability to make copies instantly, operate without retention limits, and support single-file recovery are differentiators for Cloud Volumes ONTAP.

NetApp SnapMirror integration with NetApp HCI and Cloud Volumes ONTAP allows a variety of use cases that add value for customers with respect to hybrid multicloud deployments. These use cases include the following:

- Backup and restore of HCI storage to Cloud Volumes ONTAP
- Disaster recovery to a public cloud
- Development and test in a public cloud

- Backup and restore to a public cloud
- ONTAP Select to Cloud Volumes ONTAP

Disaster Recovery to the Public Cloud

By offering efficient data replication and a rich set of data management tools and partnerships, NetApp helps businesses recover from unplanned IT outages. NetApp SnapMirror enhances disaster recovery options by using converged infrastructure and NetApp HCI. SnapMirror improves business continuity on an Element system by replicating snapshots of an Element volume to a Cloud Volumes ONTAP instance when an on-site application is not accessible. If there is a disaster at the Element site, you can serve data to clients from the Cloud Volumes ONTAP instance and then reactivate the Element system when service is restored.

Development and Test in the Public Cloud

SnapMirror technology can be used to distribute large amounts of data throughout the enterprise and to the public cloud, enabling access to data at remote locations. Having datasets in the public cloud allows clients in the cloud to access copies of the production workload.

The replication of datasets enables the efficient and predictable use of network and server resources, because SnapMirror operations can be scheduled at a predetermined replication time. Storage administrators can replicate production data at a specific time to minimize overall network usage. After data is in the public cloud, development and testing of applications can occur on cloud resources. The applications can then be mirrored back to the on-premises data center and moved into production when ready. Cloud Volumes ONTAP delivers the storage efficiencies of ONTAP, including thin provisioning, deduplication, and compression, which allows you to consume fewer cloud resources during development.

3.3 Backup of and Restore of ONTAP Select to Cloud Volumes ONTAP

In addition to the Element operations described previously, file services in NetApp HCI can be protected with Cloud Volumes ONTAP. Volumes created on the ONTAP Select system, running as a virtual guest in a NetApp HCI solution, can establish a replication relationship with another ONTAP system provisioned on the premises or in the cloud. The volumes can be a destination for either mirror or vault purposes, enabling data to be easily and rapidly transferred between the two environments (Figure 2).

Figure 2) NetApp HCI SnapMirror to Cloud Volumes ONTAP.

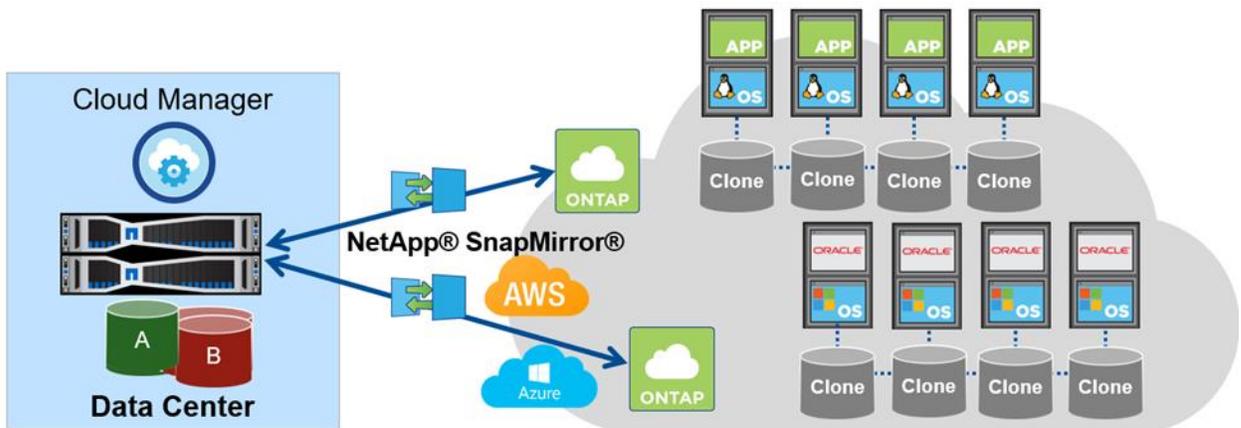
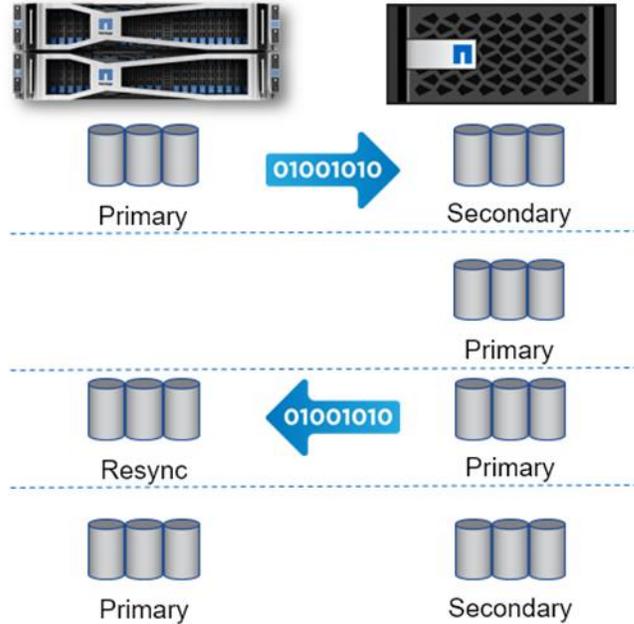


Figure 3) NetApp HCI SnapMirror to ONTAP.



To enable data transfer between two different environments, complete the following high-level steps:

1. Establish the link and replicate.
2. Promote FAS volumes.
3. Replicate back to NetApp HCI.
4. Fail back to the NetApp HCI volumes.

Performing a Transfer or One-Time Migration from ONTAP to Element

Typically, when you use SnapMirror for disaster recovery from Element to ONTAP, Element is the source and ONTAP is the destination. However, in some cases, the ONTAP storage system can serve as the source and Element can serve as the destination.

Migration can only be performed with iSCSI LUNs in a process that is primarily driven by the ONTAP CLI or the Element API. For more information, see the [SolidFire and Element Documentation page on this subject](#).

3.4 VMware Site Recovery Manager

vSphere replication, an optional feature, can be added to Site Recovery Manager (SRM) or run as a stand-alone application used as an alternative to array-based replication to provide replication at the vSphere datastore level.

3.5 HCI to HCI with VMware SRM

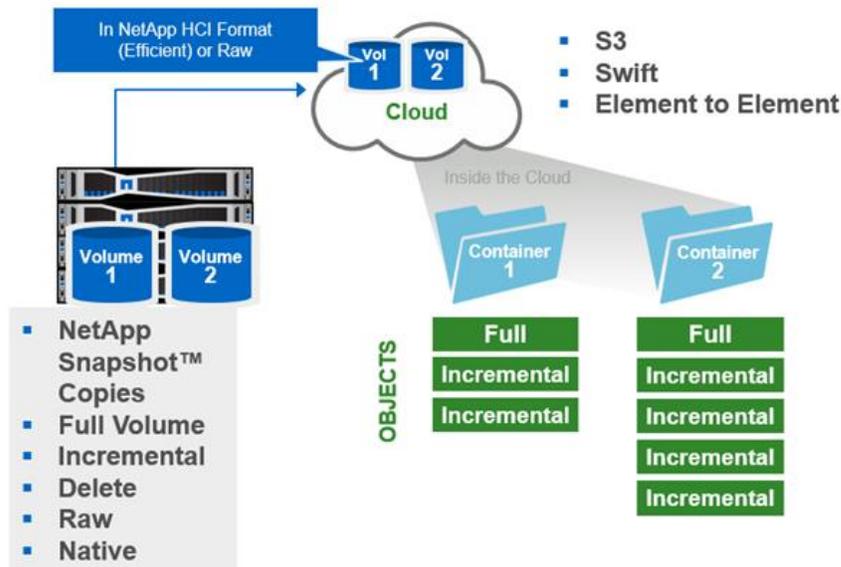
Array-based replication consists of peered storage arrays at the protected and recovery sites. This configuration allows the SRM application to communicate directly with the Element storage arrays, improving the replication of virtual environments from array to array. In addition, SRM servers can monitor and control array functions related to migrations, failovers, reprotectons, failback, and test scenarios. A Storage Replication Adaptor (SRA) must be installed on the VM hosting the SRM application. The NetApp

SolidFire SRA package is available through the VMware and NetApp download sites. For more information, see [TR-4795: NetApp HCI Disaster Recovery and Replication](#).

4 Integrated Backup and Restore

The native data protection in NetApp HCI extends past the cluster to any Amazon S3 or Swift-compatible system with integrated backup and restore. API-driven backup makes it easy to manage at scale, and it automatically resumes backup after an interruption. The incremental backup approach reduces network traffic, and effects on the host are reduced through direct transfer of data. The system maintains data efficiencies during backup, reducing network traffic, shortening backup windows, and using less target storage.

Figure 4) NetApp HCI integrated backup and restore.



You can configure the system to back up and restore the contents of a volume to and from an object store container that is external to NetApp HCI storage using Amazon S3 or OpenStack Swift. You can also backup and restore data to and from remote NetApp HCI storage systems.

- Data format:
 - **Native.** A compressed format readable only by NetApp HCI storage systems.
 - **Uncompressed.** An uncompressed format compatible with other systems.
- Automation:
 - Automatically resume backup after interruption
 - API-driven backup makes it easy to manage at scale
- Economical:
 - Allows for data protection without requiring third-party tools
 - Easily incorporates object storage into your data protection strategy
 - Offloads infrequently used point-in-time copies of data from SSD
- Efficient:
 - Uses incremental backup approach to reduce network traffic
 - Reduces effects on the host with direct transfer of data

- Maintains data efficiencies during backup, reducing network traffic, shortening backup windows, and using less target storage

Note: You can run a maximum of two backup or restore processes at a time on a volume.

For more information, see the [NetApp SolidFire and Element Documentation Page on this subject](#).

5 Data Fabric

Enterprises are under tremendous pressure to harness today’s wealth of data and apply it to create new value and competitive advantage, all with limited time, skills, and budget. The data fabric is the NetApp vision for the future of data management. NetApp HCI is NetApp data fabric ready, making sure that you can leverage the full potential of your data to be unleashed, whether on the premises or in a public or hybrid cloud. For more information on this subject, see [TR-4748: Build Your Data Fabric](#).

5.1 NetApp HCI Data Fabric Integration Points

Robust integrations provide additional services through the data fabric and third parties, including the following:

- File services using ONTAP Select
- Object services using NetApp StorageGRID® technology
- Replication with NetApp SnapMirror
- Backup and recovery using data protection partners such as Commvault and Veeam
- Orchestration and disaster recovery using VMware vRealize Automation and VMware Site Recovery Manager

Figure 5) NetApp Fabric Orchestrator.

Provider	Labels	Version	Model	Capacity	Site	Actions
Boulder	environment:production	NKS 1.12.9	HCI	10 worker nodes	Boulder	Available
Boulder	environment:production	Cloud Volumes on HCI	SDE	8.7 TB / 15 TB: 6%	Boulder	Available
East US	environment:production	Azure NetApp Files	CVS	309 TB / 400 TB: 8%	US West 2	Available
ontap-c01	dataset:training-data	ONTAP v.9.5	A800	190 TB / 204 TB: 9%	Seattle	Available
ontap-cvo2	environment:production	ONTAP v.9.5	CD-A300	128 TB / 204 TB: 6%	US East	Available
Seattle	environment:production	Cloud Volumes Service	CVS	19 TB / 204 TB: 9%	Seattle	Available
sgri4.inimtech.software	data:artifacts_dev:build	StorageGRID v1.0.1	SGWS	203 TB / 274 TB: 7%	Boulder	Available
US East	environment:production	Cloud Volumes Service	CVS	200 TB / 200 TB: 10%	US East	Available
net us-west-1	environment:development	Cloud Volumes Service	CVS	Provisioned 200 TB	us-west-1	Available

For more information about the NetApp Fabric Orchestrator, see the [NetApp Cloud Central page on this subject](#).

6 StorageGRID

The combination of NetApp HCI and StorageGRID provides customers with the speed and simplicity of deploying compute and storage along with massive S3 object storage.

NetApp HCI and StorageGRID integrate in the following ways:

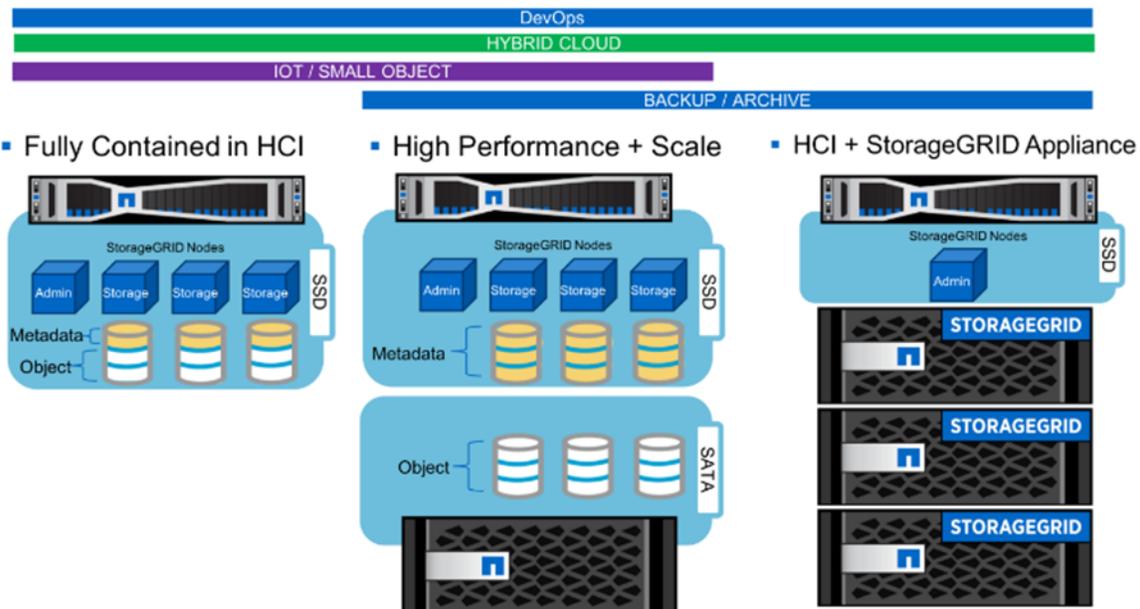
- Direct backup and restore from NetApp HCI using the Integrated Backup & Restore feature, with either S3 or SWIFT to StorageGRID.
- By using the cloud connector of one of our data protection alliance partners to StorageGRID.
- Running fully virtualized on NetApp HCI, in combination with HCI resources for performance and scale, or with the StorageGRID appliance, as shown in Figure 6.

6.1 Use Cases

StorageGRID is an optimal solution for the following use cases:

- On-premises, S3-compatible object storage on NetApp HCI.
- Hybrid cloud workflows.
- Shared image repositories requiring simultaneous read/write capabilities. Examples include the OpenShift container registry and the Docker registry.
- Workloads that ingest large amount of unstructured data for which managing directory structures and filesystems usage can be cumbersome. Examples include the internet of things and deep learning workloads.
- Adhering to data governance and regulatory requirements while ensuring simplicity of data access.

Figure 6) StorageGRID on NetApp HCI.



For more information on this subject, see [TR-4734: StorageGRID on NetApp HCI](#).

7 Partnered Solutions

7.1 Veeam

With Veeam Backup & Replication v9.5 Update 4, Veeam and NetApp HCI offer native storage integration with Element Software, the storage operating system that powers NetApp HCI.

Backup from Element Storage Snapshots

Veeam's Backup from Storage Snapshot technology is designed to dramatically reduce the performance effect typically associated with traditional API-driven VMware backups on primary hypervisor infrastructure. This process dramatically improves backup performance, with the added benefit of reducing the performance effect on production VMware infrastructure.

Granular Application Item Recovery from Element Storage Snapshots

Veeam makes the process of recovering individual Windows files, Linux files, or application items from an Element storage snapshot fast, easy, and painless. With the integration of Element snapshots and Veeam, you can quickly recover application items such as the following directly from snapshots:

- Individual Windows or Linux guest files
- Exchange items
- MS SQL databases
- Oracle databases
- Microsoft Active Directory items
- Microsoft SharePoint items

This functionality works with Element snapshots created by Veeam on NetApp HCI. The only requirement is that VMs must be in the Virtual Machine Disk (VMDK) format.

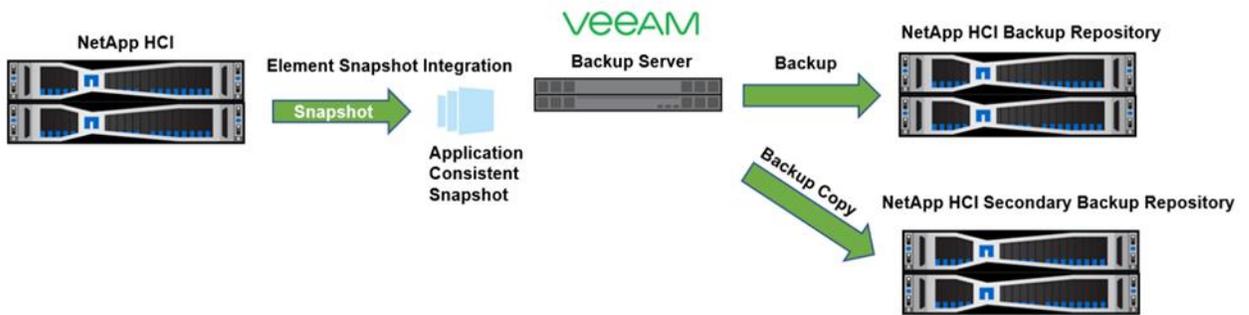
Hyper-Available VMs with Instant VM Recovery from Element Snapshots

Instant VM Recovery technology from Veeam leverages Element snapshots for NetApp HCI. With this capability, you can recover an entire VM, no matter the size, in a very short timeframe. Not only is this process extremely fast, there is no performance loss during this process, because, after it is recovered, the VM is running from your primary production NetApp HCI cluster.

Element Snapshot Orchestration for Better RPO

Most organizations use a nightly or twice-daily backup schedule. The problem with this strategy is that it exposes an organization to 12 to 24 hours of data loss. The amount of acceptable data loss is your recovery point objective (RPO) and reducing your RPO to the lowest level possible makes good business sense. With Veeam and Element snapshot management, we can supplement an off-array backup schedule with more frequent storage array-based snapshots. One common example would be taking hourly storage-based snapshots in between nightly off-array Veeam backups. When a restore event happens, you now have hourly snapshots or a Veeam backup to choose from when performing the recovery operation (Figure 7).

Figure 7) Element Snapshot integration on NetApp HCI with Veeam Backup & Replication.



For more information on this subject, see [TR-4634: NetApp HCI Reference Architecture with Veeam Backup and Replication 9.5](#).

7.2 VMware Site Recovery Manager

VMware Site Recovery Manager is an automation software solution that integrates with an underlying replication technology. It provides policy-based management, automated orchestration of recovery plans to minimize downtime in case of disasters, and non-disruptive testing of your DR plans. It is designed for VMs and is scalable to manage all applications in a VMware vSphere environment. To deliver flexibility and choice, VMware Site Recovery Manager integrates natively with vSphere Replication as well as NetApp HCI replication through our storage replication adapter.

- NetApp HCI integrates seamlessly with VMware SRM, providing availability and performance control for disaster recovery.
- VMware SRM pairs well with the NetApp HCI QoS architecture and integrated remote replication functionality, providing superior performance control for disaster recovery.
- Only NetApp HCI with VMware allows IT management to set and enforce fine-grained QoS policies for each virtual disk in the NetApp HCI system and the disaster recovery site (Figure 8).

Figure 8) VMware vCenter SRM with NetApp HCI.

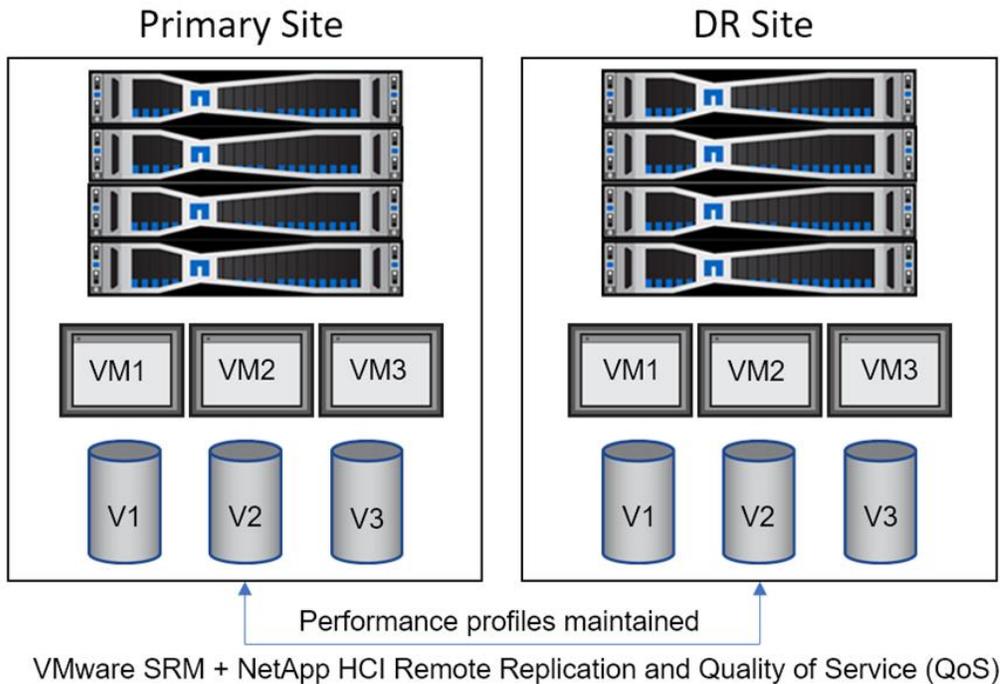
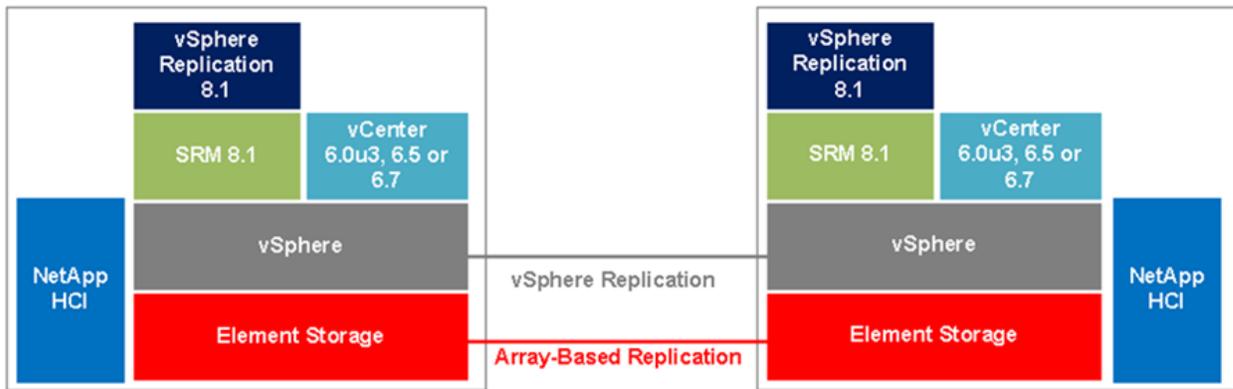


Figure 9) VMware SRM with NetApp HCI and Element SRA.



The combination of VMware SRM on NetApp HCI adds disaster recovery and planned migration capabilities. The use cases listed in Table 3 provide information on how customers use SRM in a VMware environment.

Table 3) Primary use cases.

Use case	Description
Disaster recovery	VMware SRM combined with the Element SRA plug-in provides NetApp HCI to NetApp HCI recovery against site failure by automating and orchestrating the recovery of critical systems.

Use case	Description
Planned migration	VMware SRM can be used for workload rebalancing from datacenter to datacenter. It can also be used to effectively plan migration during scheduled site maintenance. With vSphere replication, you can plan migration of workloads from your current converge infrastructure to a high-performance NetApp HCI infrastructure.
Disaster avoidance	Being proactive is key to preventing downtime. Disaster avoidance is used for a situation in which disasters are detected ahead of time. Such as a planned power outage or natural disasters.
Upgrade and patch testing	VMware SRM can be used to conduct VM and application upgrade and patch testing by creating a copy of the production environment in a test environment with isolated networking. This process is useful for dry runs before maintenance operation procedures against production components.

7.3 Commvault

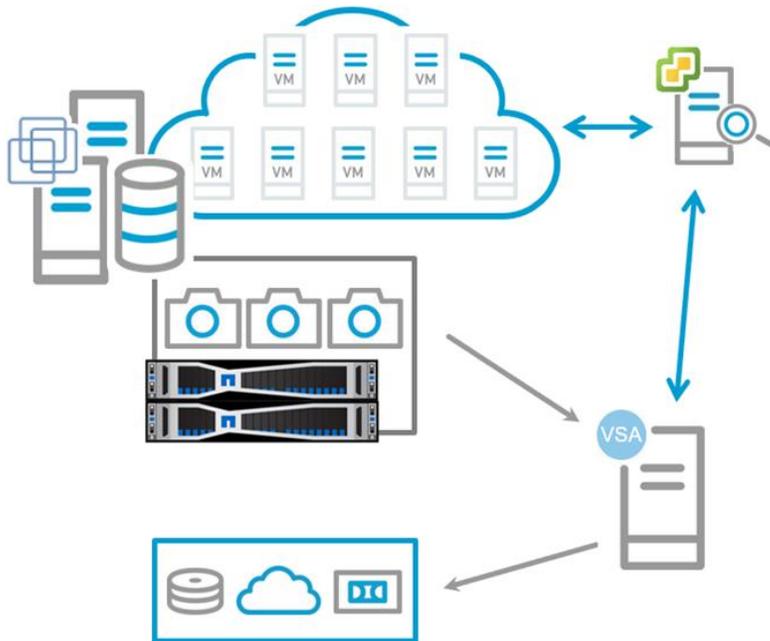
The Commvault data platform allows you to maximize your investment in NetApp HCI by extending the value of hardware snapshots with Commvault IntelliSnap, providing for faster data management and recovery and comprehensive data protection across the entire NetApp portfolio.

Benefits of Commvault IntelliSnap Technology

NetApp HCI includes hardware-based snapshots, which allow for the creation and deletion of point-in-time views of data with scheduling and retention. IntelliSnap technology significantly extends the value of hardware snapshots. IntelliSnap technology adds application awareness and consistency to automate and orchestrate the creation, retention, and access of NetApp HCI snapshots. IntelliSnap software maintains source application context to provide simple, granular recovery.

The Commvault platform provides centralized management to manage protection, retention, search, and reporting for multiple applications, heterogeneous storage platforms, locations, and environments from a single, web-based console.

Figure 10) IntelliSnap with NetApp HCI overview.



The IntelliSnap Operation Workflow is as follows:

1. vCenter is integrated for auto discovery of new VMs for protection.
2. Virtual Server Agent (VSA) contacts vCenter and creates a consistent VMware snapshot for all VMs being protected.
3. IntelliSnap software communicates with NetApp HCI to take snapshots.
4. Snapshots are mounted to an ESXi proxy for indexing and LiveBrowse operations, and VSA performs indexing operations.
5. Optionally, data can be streamed off the array to create a long-term retention copy.

For more information on this subject, see [TR-4636: NetApp SolidFire Reference Architecture with Commvault Data Platform v11](#).

8 Data Protection for Containers

NetApp HCI provides robust and secure integrations for containerized services and applications. You can backup data on a NetApp HCI volume by setting a snapshot schedule to a NetApp HCI volume. This takes snapshots of the volume at the required interval. However, it is not possible to set a snapshot schedule to a volume through the solidfire-san driver. This must be set manually using the Element web UI or Element APIs. In the event of data corruption, you can choose a snapshot and rollback the volume to the snapshot manually using the Element Web UI or Element APIs. This reverts any changes made to the volume since the snapshot was created.

Additionally, there are data protection methods within the various container platforms themselves, NetApp Trident, and third-party tools to protect persistent volumes on containers.

8.1 Kubernetes Pods, Volumes, and Projects

In the Kubernetes architecture, a set of containers can be deployed and scaled together. This is achieved by using pods, which are the minimum unit of deployment in a Kubernetes cluster. Each Kubernetes pod

is a set of one or more containers which enables more than one container to share the same resources, such as IP address and file systems. Pods run in several replicas of each service, which creates no single point of failure for critical services.

A Kubernetes volume is storage provisioned directly to a pod. Kubernetes Service supports a wide variety of volume types, including Amazon EBS, Azure Disk Storage, Google Persistent Disk, NFS, and many more. Volumes enable the containers within a pod to share information and are destroyed when their parent pod is deleted.

With Kubernetes, you can use a persistent volume with its own lifetime that exists independently of any specific pod. Persistent volumes can be used to support stateful applications, such as database services, enabling all components of an enterprise solution to be deployed and managed by Kubernetes. Using Trident to manage persistent volume claims (PVCs) insulates the developers creating pods from the lower-level implementation details of the storage that they are accessing.

8.2 Kubernetes Container Storage Interface Volume Cloning

The Container Storage Interface (CSI) Volume Cloning feature adds support for specifying existing PVCs in the `dataSource` field to indicate a user would like to clone a volume.

A Clone is defined as a duplicate of an existing Kubernetes volume that can be consumed as a standard volume. The only difference is that, upon provisioning, rather than creating a new, empty volume, the backend device creates an exact duplicate of the specified volume.

This implementation of cloning, from the perspective of the Kubernetes API, allows you to specify an existing unbound PVC as a `dataSource` during new PVC creation.

You should be aware of the following issues when using this feature:

- Cloning support (`VolumePVCDataSource`) is only available for CSI drivers.
- Cloning support is only available for dynamic provisioners.
- CSI drivers might or might not have implemented the volume cloning functionality.
- You can only clone a PVC when it exists in the same namespace as the destination PVC. The source and destination must be in the same namespace.

For more information on this subject, see the [Github Kubernetes page on volume cloning](#).

8.3 Trident for Kubernetes

You can back up data on an Element volume by setting a snapshot schedule for the volume to make sure that the snapshots are taken at the required intervals. Currently, it is not possible to set a snapshot schedule to a volume through the `solidfire-san` driver. You can set the schedule with the Element web UI or with Element APIs.

In the event of data corruption, you can choose a snapshot and rollback the volume manually using the Element Web UI or Element APIs. This reverts any changes made to the volume since the snapshot was created.

The section “Creating Snapshots of Persistent Volumes” describes a complete workflow for creating volume snapshots and then using them to create PVCs.

For more information on this subject, see the [NetApp Trident page on Kubernetes and Trident objects](#) from the [main NetApp Trident documentation page](#).

VolumeSnapshot

Similar to how the API resources `PersistentVolume` and `PersistentVolumeClaim` are used to provision volumes for users and administrators, the API resources `VolumeSnapshotContent` and `VolumeSnapshot` are provided to create volume snapshots for users and administrators.

A `VolumeSnapshotContent` resource is a snapshot taken from a volume in the cluster that has been provisioned by an administrator. It is a resource in the cluster just like `PersistentVolume` is a cluster resource. A `VolumeSnapshot` resource is a request for snapshot of a volume by a user. It is similar to a `PersistentVolumeClaim`.

Although the `VolumeSnapshots` resource allows a user to consume abstract storage resources, cluster administrators must be able to offer a variety of `VolumeSnapshotContents` volume snapshots without exposing users to the details of how those snapshots are provisioned. For these needs, there is the `VolumeSnapshotClass` resource.

Users should be aware that the API objects `VolumeSnapshot`, `VolumeSnapshotContent`, and `VolumeSnapshotClass` are custom resource definitions (CRDs), not part of the core API. In addition, `VolumeSnapshot` support is only available for CSI drivers.

As part of the deployment process, the Kubernetes team provides a sidecar helper container for the snapshot controller called `external-snapshotter`. It watches `VolumeSnapshot` objects and triggers `CreateSnapshot` and `DeleteSnapshot` operations against a CSI endpoint.

CSI drivers might not have implemented the volume snapshot functionality, and the CSI drivers that do provide support for volume snapshots are likely to use the `external-snapshotter` controller. The CSI drivers that support volume snapshots automatically install the CRDs defined for them.

Lifecycles of Volume Snapshots and Volume Snapshot Content

`VolumeSnapshotContents` are resources in the cluster. `VolumeSnapshots` are requests for those resources. The interaction between `VolumeSnapshotContents` and `VolumeSnapshots` follow this lifecycle:

Provisioning Volume Snapshots

There are two ways snapshots can be provisioned: statically or dynamically. For static provisioning, a cluster administrator creates a number of `VolumeSnapshotContents` objects. They carry the details of the real storage, which is available for use by cluster users. They exist in the Kubernetes API and are available for consumption.

For dynamic provisioning, when none of the static `VolumeSnapshotContents` objects the administrator created match a user's `VolumeSnapshot` object, the cluster might try to dynamically provision a volume snapshot specially for the `VolumeSnapshot` object. This provisioning is based on `VolumeSnapshotClasses`: the `VolumeSnapshot` must request a volume snapshot class and the administrator must have created and configured that class in order for dynamic provisioning to occur.

Binding

A user creates or has already created in the case of dynamic provisioning, a `VolumeSnapshot` object with a specific amount of storage requested and with certain access modes. A control loop watches for new `VolumeSnapshots` objects, finds a matching `VolumeSnapshotContent` object (if possible), and binds them together. If a `VolumeSnapshotContent` object was dynamically provisioned for a new `VolumeSnapshot`, the loop always binds that `VolumeSnapshotContent` object to the `VolumeSnapshot`. Once bound, `VolumeSnapshot` binds are exclusive, regardless of how they were bound. A `VolumeSnapshot` to `VolumeSnapshotContent` binding is a one-to-one mapping.

`VolumeSnapshots` remain unbound indefinitely if a matching `VolumeSnapshotContent` object does not exist. `VolumeSnapshots` are bound as matching `VolumeSnapshotContents` become available.

For more information on this subject, see the [NetApp blog post concerning the use of On-Demand Snapshots with CSI Trident](#).

Etcd Snapshots Using `etcdctl` Command Line Utility

The `etcdctl` command line utility allows you to take a snapshot of an etcd cluster and restore from the previously taken snapshot.

etcdctl Snapshot Backup

The `etcdctl` command `etcdctl snapshot save /var/etcd/data/snapshot.db` enables us to take a point-in-time snapshot of the etcd cluster. NetApp recommends using a script to take timely backups. This command can be deployed from within the etcd container, or the command can be deployed using the `kubectl exec` command directly. Store the periodic snapshots under the persistent Trident NetApp volume `/var/etcd/data` so that snapshots are stored securely and can safely be recovered should the trident pod be lost. Periodically check the volume to be sure it does not run out of space.

etcdctl Snapshot Restore

In the event of the accidental deletion or corruption of Trident etcd data, you can choose the appropriate snapshot and restore it back using the following command:

```
etcdctl snapshot restore snapshot.db --data-dir /var/etcd/data/etcd-test2 --name etcd1
```

Take note to restore the snapshot onto a different folder. This folder is indicated in the example above as `/var/etcd/data/etcd-test2`, which is on the mount inside the Trident NetApp volume. After the restore is complete, uninstall Trident. Take note not to use the `-a` flag during uninstallation. Mount the Trident volume manually on the host and make sure that the current “member” folder under `/var/etcd/data` is deleted. Copy the “member” folder from the restored folder `/var/etcd/data/etcd-test2` to `/var/etcd/data`. After the copy is complete, re-install Trident. Verify if the restore and recovery has been completed successfully by making sure all the required data is present.

8.4 Third-Party Data Protection Tools for Containers

The following third-party tools expand and enhance the data protection of containers:

- **Restic.** An open source backup program. If you are running production workloads in Kubernetes, you might want to make a backup of your disks. Traditional tools are typically too complex to setup and maintain in a dynamic compute environment like Kubernetes. To address this concern, Restic has been designed to be fast, efficient, and secure with few moving parts.
- **Stash by AppsCode.** A Kubernetes operator for Restic. Stash is a CRD controller for Kubernetes built around Restic to address the issues of traditional backup systems. Using Stash, you can backup Kubernetes volumes mounted in following types of workloads:
 - Deployment
 - DaemonSet
 - ReplicaSet
 - ReplicationController
 - StatefulSet

Stash employs the following features to improve Kubernetes backup functionality:

- Fast, secure, efficient backup of any Kubernetes volumes
- Automated configuration of Restic for periodic backup
- Storage of backed up files in various cloud storage provider, including S3, GCS, Azure, OpenStack Swift, and DigitalOcean Spaces
- Easy restoration of backups
- Periodic checks of the integrity of backed up data
- Backups can be taken in offline mode

- Support for a workload initializer for faster backup
- Prometheus-ready metrics for the backup process

For more information on this subject, see the [AppsCode page on Backing up Volumes with Stash](#).

- **Commvault.** Provides data protection and recovery for Docker containers and images. You can perform full, incremental, or synthetic full backups and you can also restore full containers and images. The first backup of a container or image is always a full backup, regardless of the selected backup type. See the [Commvault page on Docker for more information](#).

9 Conclusion

As data centers migrate away from dedicated, siloed platforms and consolidate environments, the need to scale granular components and control the performance of the applications is critical. Data centers are moving away from dedicated platforms and are trying to avoid overprovisioning to increase efficiency and cost. The NetApp HCI solution provides a unique solution with QoS limits, allowing the granular control of every application, eliminating noisy neighbors, and enabling administrators to set and satisfy all performance SLAs. As your environment changes, you can easily add compute and/or storage without overprovisioning any element of the solution.

Your enterprise needs a data protection solution that is reliable, flexible, and easy to use. NetApp HCI delivers self-healing resiliency, continuously accessible data, and a range of backup, restore, and disaster recovery options to best fit your environment’s needs. NetApp HCI provides native protection, SnapMirror replication across the data fabric, and trusted third-party integrations to keep your dynamic data protected and safe. NetApp HCI makes it easy to protect your data, and your protected data remains reliable, flexible, and easy to use on the infrastructure that makes the most sense for your business needs.

Where to Find Additional Information

To learn more about the information that is described in this document, review the following websites:

- NetApp HCI Documentation Center
<https://docs.netapp.com/hci/index.jsp>
- NetApp HCI Documentation Resources
<https://www.netapp.com/us/documentation/hci.aspx>
- NetApp Product Documentation
<https://docs.netapp.com>

Version History

Version	Date	Document Version History
Version 1.0	October 2017	Initial release
Version 2.0	October 2019	–
Version 2.1	May 2020	Updated links

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 © 2020 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.

TR-4641-0520