



Technical Report

Commvault IntelliSnap for NetApp

Solution Overview

Subhash Athri, NetApp
July 2017 | TR-3920

Abstract

NetApp® Commvault IntelliSnap for NetApp management software is changing today's backup and recovery landscape. Commvault IntelliSnap for NetApp software combines simplified manageability, power, and flexibility for virtual environments with full support for enterprise database applications. Commvault IntelliSnap for NetApp integrates with NetApp Snapshot® technology in a virtually seamless way for fast and efficient backup operations and with NetApp SnapVault® and NetApp SnapMirror® software to support content cataloging and data movement to tape-based media.

This document is an introduction to the Commvault IntelliSnap for NetApp solution. It contains an overview of the technology and describes the basic configuration options for getting started with Commvault IntelliSnap for NetApp.

TABLE OF CONTENTS

1	Introduction	5
1.1	Commvault IntelliSnap for NetApp Solution Components	5
1.2	Commvault IntelliSnap for NetApp Basic Terminology	7
2	Clustered Data ONTAP Overview	7
2.1	Clustered Data ONTAP Basic Terminology	7
2.2	Clustered Data ONTAP Replication Prerequisites	8
3	Commvault IntelliSnap for NetApp Technical Overview	9
3.1	Basic Functions	9
3.2	Understanding the Backup Workflow	25
3.3	Understanding the Restore Workflow	31
3.4	Data Cloning	31
3.5	OSDP for Data ONTAP 7-Mode Target	32
3.6	OSDP for Clustered Data ONTAP Target	33
4	Application Data	34
5	Virtualization Data	36
5.1	VMware and Applications	37
6	Commvault IntelliSnap for NetApp Network and Security Considerations	38
6.1	Firewall Considerations	38
7	Commvault IntelliSnap for NetApp for E-Series	38
8	Role-Based Access Control for Commvault IntelliSnap for NetApp	39
9	Commvault IntelliSnap for NetApp for MetroCluster	40
9.1	Commvault IntelliSnap for NetApp Backup Configuration for MetroCluster	40
10	Commvault IntelliSnap for NetApp Disaster Recovery	41
10.1	Disaster Recovery Backups	41
10.2	Disaster Recovery Failover	41
10.3	Alternate Disaster Recovery Method	42
11	Summary	42

Resources.....	42
Version History	42

LIST OF TABLES

Table 1) Commvault IntelliSnap for NetApp provisioning policies preconfigured in OnCommand Unified Manager 5.x.	12
Table 2) Default storage provisioning policies in OnCommand Unified Manager 6.0.....	14
Table 3) Enterprise example of scheduling and retention.	30
Table 4) Application support, iDAs, and restore granularity.....	35

LIST OF FIGURES

Figure 1) Commvault IntelliSnap for NetApp software overview.....	5
Figure 2) Unified management for Data ONTAP 7-Mode, clustered Data ONTAP, and OSDP.....	6
Figure 3) Distinction between node, HA pair, cluster, and SVM.	8
Figure 4) NetApp NAS NDMP iDA.	10
Figure 5) Windows file system iDA.....	10
Figure 6) OnCommand handling of storage provisioning and replication.	12
Figure 7) Default storage provisioning policies in OnCommand 6.0 GUI.	14
Figure 8) NAS volume mirroring.	15
Figure 9) Fan-in topology for vaulting NAS qtrees in Data ONTAP 7-Mode.....	16
Figure 10) Fan-in topology for vaulting NAS volumes and qtrees in Data ONTAP 7-Mode.....	16
Figure 11) Poor storage layout of multiple clients with LUN data on a common volume.	17
Figure 12) Better storage layout of multiple clients with LUN data on separate volumes.	18
Figure 13) Efficient storage layout of a single client with LUN data on a common volume.....	18
Figure 14) Flexibility of replication combinations.	21
Figure 15) GUI path for adding AltaVault to Commvault IntelliSnap for NetApp as a disk library.	24
Figure 16) AltaVault specified as a disk library in the Add Disk Library dialog box.....	24
Figure 17) Snapshot management options for AltaVault.	25
Figure 18) Commvault IntelliSnap for NetApp example workflow: vault > mirror > tape.	26
Figure 19) Commvault IntelliSnap for NetApp for clustered Data ONTAP backup workflow.	27
Figure 20) Example schedule and retention at specific times.....	28
Figure 21) OSDP backup workflow.	33
Figure 22) VSA layers.	36
Figure 23) Datastores in separate subclients with different storage policies.	37
Figure 24) Datastores in the same subclient with fan-in.	37
Figure 25) Capabilities to associate with clients and entities.	39
Figure 26) GUI path for integrating CommServe with Active Directory users and enabling SSO.	39
Figure 27) Cluster failover.	40


1 Introduction

NetApp is an industry leader in array-based data protection. The efficiencies of Snapshot copy technology and data replication have changed the way the industry looks at backup and recovery and at disaster recovery strategies. The explosion of data with which enterprises are dealing makes the need to achieve faster SLAs and meet backup windows a constant challenge. Data center consolidation through virtualization has created additional challenges for data protection. Disk-to-disk data protection solutions have become more widely accepted for both backup and recovery and disaster recovery.

NetApp data protection solutions offer speed and flexibility while reducing storage capacity requirements through the use of efficient array-based technologies. The result is a simplified approach that reduces costs and administrative effort.

NetApp Commvault IntelliSnap for NetApp management software offers enterprise-class management for backup and recovery in the data center. The Commvault IntelliSnap for NetApp software manages Snapshot copies on NetApp primary storage and data replication to secondary and tertiary storage as well as tape creation. Regardless of whether you are protecting application data, file data in network-attached storage (NAS), file data in LUNs, or data in virtualized environments, the Commvault IntelliSnap for NetApp solution provides the management, the storage provisioning, the cataloging, and the granular recoverability that are required for seamless operation. Figure 1 shows the basic flow of the Commvault IntelliSnap for NetApp solution.

Figure 1) Commvault IntelliSnap for NetApp software overview.




1.1 Commvault IntelliSnap for NetApp Solution Components

The Commvault IntelliSnap for NetApp solution has the following high-level components:

- **Commvault IntelliSnap for NetApp server.** Functions as a single interface for unified management (Figure 2 shows the Commvault IntelliSnap for NetApp unified management console).
- **NetApp OnCommand® Unified Manager.** Enables automated storage provisioning.

- **FAS storage system (NetApp controllers).** Provides Snapshot backups and integration with Data ONTAP® data protection technologies.
- **Third-party storage or direct-attached storage (DAS).** Can be backed up to NetApp storage through Commvault IntelliSnap for NetApp for open systems (open system data protection [OSDP]).
- **iData agents (iDAs).** Provide client communication and application consistency.

Figure 2) Unified management for Data ONTAP 7-Mode, clustered Data ONTAP, and OSDP.



Commvault IntelliSnap for NetApp software can be used to protect the following applications hosted on NetApp primary storage by using their respective iDAs:

- Microsoft Exchange Server (including DAG configuration)
- Microsoft SQL Server (including availability group deployments)
- Microsoft Office SharePoint Server
- Oracle Database (RAC configurations)
- DB2
- SAP for Oracle
- Lotus Domino

In addition, the Commvault IntelliSnap for NetApp solution supports the following virtualization products:

- VMware vSphere
- Microsoft Hyper-V

For a complete list of supported platforms and product versions, refer to the NetApp [Interoperability Matrix Tool](#).

1.2 Commvault IntelliSnap for NetApp Basic Terminology

The following specific components work together to create a full Commvault IntelliSnap for NetApp solution:

- **Backup set.** A layer of management within iDAs for grouping subclients.
- **Clients.** Hosts that run iDAs for which data is protected.
- **CommCell.** A single instance of a Commvault IntelliSnap for NetApp environment.
- **CommCell console.** The Commvault IntelliSnap for NetApp management interface.
- **CommServe.** The master server in a Commvault IntelliSnap for NetApp environment. This server uses a SQL Server database; therefore, it must be a Windows system (Windows Server 2003 or 2008).
- **Disk library.** A storage resource with an associated mount path that is used to store backups of index information in the Commvault IntelliSnap for NetApp solution. These backups are also targets for backups to cloud storage (NetApp AltaVault® appliance [formerly known as SteelStore]) and act as backup targets for streaming backups.
- **iDAs.** The agents that control data consistency during backup operations.
- **MediaAgent.** A media server in a Commvault IntelliSnap for NetApp environment. MediaAgents have broad operating system support, including Windows, Linux, and UNIX options.
- **NetApp management console (NMC).** The NMC is an interface used for creating resource pools and provisioning policies within the OnCommand framework. The NMC should be installed on a separate system from the OnCommand Unified Manager server.
- **NetApp primary storage.** The production NetApp storage array.
- **NetApp secondary storage.** The secondary NetApp storage array used as a destination for replication.
- **NetApp tertiary storage.** A third NetApp storage array used for replicating previously replicated data.
- **OnCommand Unified Manager server.** A server running OnCommand Unified Manager software. The OnCommand server and the CommServe server are typically separate systems. The OnCommand server is mandatory for replication (auxiliary copy) and adds additional capabilities to the system, such as monitoring, alerting, and reporting.
- **SnapMirror.** A NetApp replication technology used for disaster recovery. In the Commvault IntelliSnap for NetApp solution, a mirror copy uses SnapMirror.
- **Snapshot copy.** A NetApp array-based, point-in-time copy used for recovering data.
- **SnapVault.** A NetApp replication technology used for backup and recovery. In the Commvault IntelliSnap for NetApp solution, a vault copy uses SnapVault.
- **Storage policy.** A logical object through which a subclient is protected. The storage policy defines how data is backed up and replicated; it also defines retention requirements.
- **Subclient.** A layer of management within a backup set. A client can have multiple subclients, each of which can be associated with different source data.

2 Clustered Data ONTAP Overview

2.1 Clustered Data ONTAP Basic Terminology


The following clustered Data ONTAP concepts are relevant for the Commvault IntelliSnap for NetApp solution:

- **Cluster.** One or more storage nodes that are interconnected and managed as a single system.
- **Clustered Data ONTAP.** The Data ONTAP operating mode that supports the interconnection of nodes into a cluster.

- **Cluster interconnect.** A dedicated high-speed, low-latency, private network used for communication and replication between nodes in the same cluster.
- **Data network.** The network used by clients to access data.
- **HA interconnect.** The dedicated interconnect between two nodes in one high-availability (HA) pair.
- **HA pair.** Two nodes configured in a pair for high availability.
- **Intercluster LIF.** A LIF used only for intercluster replication, assigned to only one node.
- **Intercluster network.** The network used for communication and replication between different clusters.
- **Interface group.** A collection of combined physical ports to create one logical port that is used for link aggregation; an integration group.
- **Logical interface (LIF).** A logical interface that is assigned an IP address that provides an Ethernet access point to a particular node in the cluster.
- **Management network.** The network used for administration of the cluster, storage virtual machines (SVMs; formerly called Vservers), and nodes.
- **Node.** A single NetApp controller or one of the controllers in an HA pair.
- **Port.** A physical port, such as e0e or e0f, or a logical port, such as a virtual LAN (VLAN) or an interface group.
- **SVM.** A logical storage server that provides data access to LUNs, a NAS namespace, or both from one or more LIFs.

Figure 3 shows how the cluster, HA pairs, nodes, and SVMs are related in the clustered Data ONTAP architecture.

Figure 3) Distinction between node, HA pair, cluster, and SVM.



2.2 Clustered Data ONTAP Replication Prerequisites

To provide cross-site disaster recovery (DR) capabilities to data centers, clustered Data ONTAP 8.1 and later allows replication between different clusters or different SVMs. Replication involves the following clustered Data ONTAP concepts:

- **Cluster peering.** The act of connecting two clusters to allow replication to occur between them.
- **Intercluster LIFs.** Logical network interfaces used for intercluster communication.
- **Intercluster ports.** Ports dedicated to intercluster replication.
- **SVM peering.** SVM associations defined in OnCommand Unified Manager.

For replication between clusters, clusters must be joined in a peer relationship before replication between them is possible. Cluster peering is required for replication because it defines the network on which

replication between different clusters can occur. Cluster peering is a one-time operation that must be performed by the cluster administrator. To configure replication between different clusters, the cluster administrator must complete the following steps:

1. Understand the prerequisites for cluster peering.
2. Determine whether to share ports for data access and intercluster replication.
3. Designate ports for intercluster replication, if using dedicated ports.
4. Create intercluster LIFs on each node in the clusters.
5. Peer the clusters together.

For replication between SVMs, SVMs must be joined in a peer relationship before replication between them is possible. SVM peering became a prerequisite for replication between SVMs starting with clustered Data ONTAP 8.2. In clustered Data ONTAP 8.1, any SVM could replicate data to any other SVM in the same cluster or in a peer cluster. Control of replication security could be maintained only at the clusterwide level. Clustered Data ONTAP 8.2 offers more granularity in replication security because replication permissions must be defined by peering the SVMs together.

Before any replication relationships can be created between a pair of SVMs, the SVMs must be in an SVM peering relationship. The SVMs can be local (intracluster) or remote (intercluster). SVM peering is a permission-based mechanism and one-time operation that must be performed by the cluster administrator. To configure replication between different SVMs, the cluster administrator must complete the following steps:

1. Understand the prerequisites for SVM peering.
2. Peer the SVMs together.
3. Create NetApp SnapMirror and SnapVault relationships between different SVMs.

Note: Before you can run replication workflows within Commvault IntelliSnap for NetApp, you must configure cluster peering and SVM peering through the OnCommand Unified Manager interface.

Best Practice

The best practice for naming SVMs is to name each SVM with a unique fully qualified domain name (FQDN). For example, you can name an SVM as `dataNasSVM.HQ` or as `mirrorDisasRecovSVM.Offsite`. SVM peering requires unique SVM names; using the FQDN naming convention makes the task of making sure of name uniqueness much easier.

3 Commvault IntelliSnap for NetApp Technical Overview

This section covers the technical details of the Commvault IntelliSnap for NetApp software and how the components work together.

3.1 Basic Functions

The Commvault IntelliSnap for NetApp solution delivers several basic functionalities to create a simplified user experience:


- Snapshot copy creation
- Cataloging and indexing
- Storage provisioning through OnCommand Unified Manager
- Data replication through SnapVault, SnapMirror, or both
- Data movement to tape
- Data movement to the cloud

Snapshot Copy Creation

The Commvault IntelliSnap for NetApp software creates Snapshot copies on the NetApp primary storage as the first backup copy. Snapshot technology allows backups to complete very quickly. Primary Snapshot copy creation is handled differently for different types of data.


For NAS data, the NetApp primary storage system is treated as a client. The system is associated with an iData agent (iDA), called NetApp NAS NDMP. Subclients within this iDA are configured and associated with the NetApp data that requires protection. When a backup for the subclient runs, Commvault IntelliSnap for NetApp creates Snapshot copies for the volumes in that subclient. Figure 4 illustrates this process for a NetApp primary storage system and its iDA.

Figure 4) NetApp NAS NDMP iDA.



For LUN data hosted on NetApp primary storage, the host accessing the data is treated as the client. The attached drive on the client is associated with a subclient within either the file system iDA or the associated application iDA. On Windows clients, to make sure that the data within the file system is consistent, the file system iDA calls Microsoft Volume Shadow Copy Service (VSS). With application-integrated Snapshot copies, to enforce backup consistency, the application iDA calls VSS (on a Windows client) or places the database in hot-backup mode (on a UNIX or Linux client). Commvault IntelliSnap for NetApp then creates the Snapshot copy for the volume containing the LUN on the NetApp primary system. Figure 5 illustrates this process for a Windows client and its file system iDA.

Figure 5) Windows file system iDA.



Cataloging and Indexing

Commvault IntelliSnap for NetApp software catalogs (that is, indexes) the data that it backs up. This functionality is a core value of the Commvault IntelliSnap for NetApp backup solution. Commvault IntelliSnap for NetApp uses different strategies to index data:

- For basic NAS data, Commvault IntelliSnap for NetApp directly indexes the contents of the Snapshot copies that it creates.
- For LUN data, Commvault IntelliSnap for NetApp indexes the contents of the LUN with the help of LUN clones and of the file system iDA for the client's operating system (Linux or Windows).

- Application-specific indexing for Exchange Server, SQL Server, Oracle Database, VMware vSphere, and the other applications supported by Commvault IntelliSnap for NetApp is handled by the iDAs for each application.

During the index-creation operation, Commvault IntelliSnap for NetApp uses a temporary working area, called index cache, to keep metadata. The index cache directory is the location where a single index-creation job places its working files while the job is running. These files are not deleted immediately when the job completes, but they may be aged out to free up more working space if needed. By default, Commvault IntelliSnap for NetApp creates the index cache under the software installation directory of the MediaAgent during installation.

When a backup job runs, the index data is written to two places:

- During the course of the job, it is written to the index cache on the data mover MediaAgent.
- During the archive index phase, it is written to the disk library that is part of the storage policy.

Best Practice

When configuring a location for the index cache, observe the following recommendations:

- For best backup performance and reliability, place the index cache on a local disk of the MediaAgent computer.

Note: Changing the index cache directory is not allowed in Commvault IntelliSnap for NetApp 10 SP9 and later versions.

- For optimal performance, use solid-state drive technology for the local index cache disk. This is particularly important when you:
 - Run NDMP backups by using the NAS iDA
 - Back up large file servers
 - Have any situation in which performance is critical

The indexing process can also be deferred until after the backup has been successfully transferred to the secondary storage. This option allows you to decouple backup and indexing schedules and have further control over processing cycle utilization. If the need arises for you to run a restore before indexing is performed, you can use the live browse feature to browse Snapshot backups instantly and list files and folders for the restore operation. The live browse feature is helpful in scenarios that involve tight backup windows on production controllers or a lot of files; it can also be used when you want to complete the backup job as quickly as possible. In these cases, you specify in the backup policy the primary or secondary copy (vault or mirror Snapshot copy) on which Commvault IntelliSnap for NetApp should run the indexing operation. Regardless of when and where you choose to index the backups, after the indexing operation completes, you have a full backup catalog that is browsable and searchable.

Note: The live browse feature is available for the NAS and VMware iDAs.

Indexes are stored in disk libraries. NetApp recommends creating disk libraries with paths that point to NetApp primary storage.

Note: For NAS data, iDA full indexing is performed only once for a subclient. Subsequent full indexes are the copy of the last full index plus the delta (incremental indexing).

Storage Provisioning


Storage provisioning is required whenever Snapshot copies on NetApp primary storage must be replicated to NetApp secondary storage and tertiary storage. Before replication can be established, the secondary system must have the appropriate volumes in place along with the correct volume settings. Commvault IntelliSnap for NetApp software takes care of these tasks by using the provisioning services of OnCommand Unified Manager.

OnCommand Unified Manager uses policy-based rules that define how storage should be provisioned under different circumstances. The location of the provisioned storage is also flexible because OnCommand uses pools of storage called resource pools: OnCommand containers that point to one or more aggregates within a NetApp storage system. The storage administrator needs only to provide resource pools to the backup administrator. The backup administrator then directs the Commvault IntelliSnap for NetApp software to use these resource pools and provisioning policies for replication purposes. The appropriate storage is provisioned automatically.

Figure 6 shows how OnCommand Unified Manager handles provisioning and replication. When a new replication relationship is configured in Commvault IntelliSnap for NetApp management software, this information is passed on to OnCommand, which creates the datasets that are required to manage the replication requirements.

OnCommand then provisions the necessary volumes, using the resource pools and provisioning policies that were assigned in the Commvault IntelliSnap for NetApp configuration.

Figure 6) OnCommand handling of storage provisioning and replication.



Before the automatic storage provisioning process can be initiated, you must create resource pools manually by using the OnCommand Unified Manager interface. The Commvault IntelliSnap for NetApp software then discovers the resource pools and makes them available in the CommServe console for selection. You do not need to create provisioning policies manually unless custom policies are required. OnCommand Unified Manager 5.x has three preconfigured provisioning policies that can be used with Commvault IntelliSnap for NetApp. These policies are described in Table 1.

Table 1) Commvault IntelliSnap for NetApp provisioning policies preconfigured in OnCommand Unified Manager 5.x.

Policy Name	Availability	Deduplication	Space Thresholds
Commvault IntelliSnap for NetApp_Raid-DP	RAID DP®	No	No

Policy Name	Availability	Deduplication	Space Thresholds
Commvault IntelliSnap for NetApp_Dedupe	RAID DP	On demand	80%, 90%
Commvault IntelliSnap for NetApp_Mirror_Destination	RAID DP	No	80%, 90%

The policies in Table 1 have the following characteristics:

- Because NetApp recommends the use of NetApp RAID DP technology for resiliency, all of the preconfigured provisioning policies enable RAID DP.
- Deduplication can be enabled or disabled for vault copies independently of the deduplication setting on the primary data volume. Deduplication for vault copies uses the on-demand setting and runs automatically as vault copy jobs complete. To enable deduplication for vault secondary storage, you can use the `Commvault IntelliSnap for NetApp_Dedupe` provisioning policy. For vaulting that does not require deduplication on secondary storage, use the `Commvault IntelliSnap for NetApp_RAID-DP` policy.
- Mirror copies inherit the deduplication settings of the primary data volume. Therefore, deduplication is disabled in the `Commvault IntelliSnap for NetApp_Mirror_Destination` provisioning policy. If the primary volume has deduplication enabled, the mirror copy volume is also a deduplicated volume.
- Space thresholds represent the nearly full threshold and full threshold properties that are used when provisioning storage.

If additional settings apart from the ones in the preconfigured policies are required, OnCommand Unified Manager 5.x supports custom provisioning policies created by using the NMC. The Commvault IntelliSnap for NetApp software automatically discovers custom provisioning policies, but their names must begin with the prefix `Commvault IntelliSnap for NetApp_` for Commvault IntelliSnap for NetApp to be able to discover them.

OnCommand Unified Manager 5.x is used to provision storage for Data ONTAP 7-Mode environments. OnCommand Unified Manager 6.0 offers comparable functionality for clustered Data ONTAP environments. Commvault IntelliSnap for NetApp requires an OnCommand 5.x server for Data ONTAP 7-Mode and Commvault IntelliSnap for NetApp for open systems (OSDP) data and an OnCommand 6.0 server for clustered Data ONTAP data.

Like OnCommand Unified Manager 5.x, OnCommand Unified Manager 6.0 has default storage provisioning policies that can be used with Commvault IntelliSnap for NetApp. Figure 7 shows the four default provisioning policies in the OnCommand Unified Manager 6.0 GUI. Table 2 describes these policies.

Figure 7) Default storage provisioning policies in OnCommand 6.0 GUI.




Table 2) Default storage provisioning policies in OnCommand Unified Manager 6.0.

Policy Name	Availability	Description
SnapMirror Destination	RAID DP	Policy on the secondary volume used for mirror replication
SnapVault Destination	RAID DP	Policy on the secondary volume used for vault replication
SnapVault Destination with Deduplication	RAID DP	Policy with deduplication enabled on the secondary volume used for vault replication
SnapVault Destination with Deduplication and Compression	RAID DP	Policy with deduplication and compression enabled on the secondary volume used for vault replication
SnapMirror Destination	RAID DP	Policy on the secondary volume used for mirror replication

Note: Commvault IntelliSnap for NetApp does not support the creation of custom storage provisioning policies for clustered Data ONTAP.

Note: Commvault IntelliSnap for NetApp can enable Data ONTAP storage efficiencies such as deduplication and compression while provisioning volumes for secondary storage (replication targets) through the provisioning policies in Table 2. However, you cannot use the Commvault IntelliSnap for NetApp GUI to enable storage efficiency settings on primary volumes (source data). On primary volumes, you must enable the settings manually on the controller by using the native Data ONTAP commands.

Commvault IntelliSnap for NetApp Subclient Mapping to OnCommand Datasets

When a new replication job is created in Commvault IntelliSnap for NetApp, that task is given to OnCommand Unified Manager for implementation. OnCommand creates the required datasets, provisions the required volumes, and initiates the baseline data transfers. This subsection explains in more detail how source data gets from a subclient to a dataset and then to a destination volume during replication. This process differs depending on the replication type and on the data type.


In all cases, a Commvault IntelliSnap for NetApp subclient has a one-to-one mapping to an OnCommand Unified Manager dataset. In addition, a Commvault IntelliSnap for NetApp subclient cannot span clients.

NAS Data

If a single NetApp primary system is configured as a NAS client, all of the NAS volumes on that primary system can be grouped together in a single subclient. The result is a single dataset in OnCommand Unified Manager. If the storage policy in the Commvault IntelliSnap for NetApp software calls for this subclient to be mirrored, then the dataset creates mirror relationships for each of the volumes in the subclient.

For example, in Figure 8, three NAS volumes are grouped into a single subclient. Creating a mirror copy results in a single dataset in the OnCommand server, and three mirror relationships are established, one for each of the volumes.

Figure 8) NAS volume mirroring.




Vaulting NAS data is slightly different because a vault can be more granular in scope. In Data ONTAP 7-Mode, either an entire volume or individual qtrees in a primary volume can be selected for vaulting purposes.

In Figure 9, single qtrees from three volumes are grouped into a single subclient. Creating a vault copy results in a single dataset in the OnCommand Unified Manager server, and three vault relationships are established, one for each of the qtrees. In this example, OnCommand is configured to allow a fan-in of the vaulted relationships.


Note: To enable fan-in on the OnCommand server, you must set the `dpMaxFanInRatio` parameter on the server. For example, to set the fan-in ratio to 10, run `dfm options set dpMaxFanInRatio=10` on the OnCommand server.

Figure 9) Fan-in topology for vaulting NAS qtrees in Data ONTAP 7-Mode.



If, in contrast, the entire three volumes are vaulted, six vault relationships are established because SnapVault creates a relationship for each volume's nonqtree data and a relationship for each of the volume's qtrees. In Figure 10, the three volumes are grouped into a single subclient and vaulted with fan-in enabled.

Figure 10) Fan-in topology for vaulting NAS volumes and qtrees in Data ONTAP 7-Mode.



In the fan-in scenarios in Figure 9 and Figure 10, multiple SnapVault primary volumes are backed up to one SnapVault secondary volume. The primary use case for this topology is multiple remote sites that back up to one central data center. In Data ONTAP 7-Mode, fan-in means that multiple qtrees can be backed up to the same secondary volume; that is, SnapVault replicates at the qtree level. This topology is referred to as volume-level fan-in.

In clustered Data ONTAP, by contrast, you cannot back up multiple volumes to a single secondary volume because relationships are configured at the volume level. However, you can back up SnapVault

primary volumes from multiple storage virtual machines (SVMs), clusters, or both to separate volumes in a single destination SVM, which can be in a different cluster from the source SVM. This topology is referred to as system-level fan-in.

LUN Data

When working with LUN data, you must follow specific guidelines independently of whether you use an application iDA or the file system iDA. Because Commvault IntelliSnap for NetApp subclients does not span clients, you must lay out primary data with the subclient in mind.


NetApp recommends that all LUNs in the primary volume be protected by a single subclient. For this setup to work, the following conditions must be true:

- All LUN data on the volume must belong to the same client.
- All LUN data on the volume can be protected by the same iDA.

Volumes with LUN data split across multiple subclients result in increased capacity requirements for replication operations. Consider an example in which three clients (F:\) each map to LUNs in a common volume. Three subclients are used to protect these LUNs. If these subclients were mirrored, the result would be three datasets and three baseline copies of the common volume, as shown in Figure 11.


Vaulting would result in the same wasted capacity and duplicated data unless each LUN were located in its own qtree.

Figure 11) Poor storage layout of multiple clients with LUN data on a common volume.




In Figure 12, the LUNs are provisioned on separate volumes. Notice how the replication operation results in an improved storage layout.

Figure 12) Better storage layout of multiple clients with LUN data on separate volumes.



In Figure 13, the LUNs on the common volume are mapped to a single client and grouped in a single subclient rather than spread out in three separate subclients. The results of mirroring this subclient are a single dataset and a single baseline copy for the common volume. This configuration is ideal for maximizing the available storage capacity.

Figure 13) Efficient storage layout of a single client with LUN data on a common volume.



Storage Layout Design Considerations

SVM mapping is both LUN agnostic and NAS agnostic. When you create storage pools in OnCommand Unified Manager, consider the following points as they relate to SVM associations:

- There are two ways of associating SVMs: You can create an association between any source SVM and any destination SVM or between a specified source SVM and a specified destination SVM. After an SVM association is created, it cannot be modified. You must delete it and create a new one.
- By default, SVM associations are not bidirectional. For example, if you created the `svmProdLocA > svmProdLocB` association and now you want to replicate data from `svmProdLocB` to `svmProdLocA`, then you must create a new `svmProdLocB > svmProdLocA` association in OnCommand Unified Manager and define the appropriate SnapVault or SnapMirror relationship for the new SVM association.
- To replicate data between SVMs belonging to different clusters (intercluster replication), use OnCommand System Manager or the Data ONTAP CLI to enable cluster peering between the clusters that contain the SVMs before you create the SVM associations.

- You can create an SVM association with only one destination SVM per cluster. For example, suppose a four-node cluster has three SVMs: `svm1`, `svm2`, and `svm3`. If you already created the associations `svm1 > svm2` (SnapMirror) and `svm1 > svm3` (SnapVault), OnCommand 6.x does not allow `svm1` to be associated with any other SVM within that cluster. However, you can still associate `svm1` to any SVM that belongs to a different cluster as long as you observe the restriction of having only one destination SVM in that cluster as well.

Data Replication Using SnapVault and/or SnapMirror

Two types of replication operations can be configured using the Commvault IntelliSnap for NetApp solution: vault copies and mirror copies. This section explains the difference between the two types of replication.

A vault copy uses NetApp SnapVault software, and a mirror copy uses NetApp asynchronous volume SnapMirror software. SnapVault and SnapMirror replicate data in a similar way in that they both replicate only the blocks that have changed since the last replication operation; a full copy of the data is created only once. From that point on, SnapVault and SnapMirror make block-based incremental backups. This approach to disk-to-disk data protection increases speed, network efficiency, and storage capacity savings.

A vault copy has a certain independence from the primary data, allowing the retention levels of the primary copy and of the secondary copy to be different. For example, a vault copy can have longer term retention than the primary copy.

A mirror copy depends more firmly on the primary data and is an exact mirror of the source volume and its Snapshot copies. A mirror copy cannot have an independent level of retention. Mirror copies are traditionally used in disaster recovery (DR) solutions because they allow failover to the secondary copy. Another use case for mirror copies is for duplication of a vault copy to a remote location.

Commvault IntelliSnap for NetApp allows several combinations of vaulting and mirroring to satisfy many disk-to-disk-to-tape requirements. Clustered Data ONTAP 8.2.x supports cascading SnapMirror and SnapVault relationships. A SnapMirror secondary copy can be the source of a SnapVault relationship (backing up a DR mirror), or a SnapVault secondary copy can be the source of a SnapMirror relationship (protecting a backup). In a SnapMirror to SnapVault relationship, you cannot specify which Snapshot copies are transferred to the SnapVault secondary copy. SnapVault always transfers the Snapshot copy exported by SnapMirror or the base copy of the SnapMirror relationship. This behavior is similar to the SnapMirror base Snapshot only option in Data ONTAP 7-Mode.

SnapVault for Clustered Data ONTAP

SnapVault was rebuilt from the ground up for its debut in clustered Data ONTAP 8.2. Although users will find similarities between the newer version and the 7-Mode version, major enhancements have been made to the clustered Data ONTAP version of SnapVault.

An important architectural change is that SnapVault in clustered Data ONTAP replicates at the volume level, not at the qtree level, as does SnapVault in Data ONTAP 7-Mode. In clustered Data ONTAP, the source of a SnapVault relationship must be a volume, and the volume must replicate to its own volume on the SnapVault secondary copy.

Note: Both the primary and secondary storage systems must be running clustered Data ONTAP 8.2 or later.

SnapVault users can now take advantage of nondisruptive operations, the cornerstone of clustered Data ONTAP. SnapVault administrators can seamlessly rebalance SnapVault primaries and secondaries according to performance or capacity needs, because SnapVault primary and secondary volumes can be moved to different aggregates or nodes within a cluster without disrupting SnapVault operations. If a SnapVault transfer is in progress when a volume is moved by using volume move, the transfer may pause for a few minutes during the volume cutover phase, but it then resumes from the most recent

transfer checkpoint after the volume move operation completes. Administrators never have to reconfigure a SnapVault relationship just because a volume was moved to another node by using the volume move operation.

One major advance in the clustered Data ONTAP version of SnapVault is the ability to preserve the storage efficiencies of the primary volume throughout a SnapVault transfer. If deduplication and compression are enabled on the primary volume, these efficiencies are preserved during SnapVault transfers, resulting in less data being transmitted over the network and leading to shorter backup windows and bandwidth savings. The data is already deduplicated and compressed on the secondary volume after the SnapVault transfer completes, and there is no need to run the deduplication and compression processes on the secondary volume.

For the SnapVault transfer to be successful, it requires at least the same amount of available free space on the secondary volume that would be needed to copy the data from the primary volume in a nondeduplicated and noncompressed format. During the SnapVault transfer, it appears that the expanded size of the primary data has been consumed on the secondary volume. However, as soon as the transfer completes, the storage-efficient dataset size is reflected in the amount of space consumed.

It is possible for deduplication and compression to run on the secondary volume after the SnapVault transfer has completed independently of the storage efficiencies that are configured on the primary volume. However, enabling the compression process to run on the SnapVault secondary volume causes the storage efficiencies present on the primary volume not to be preserved on the SnapVault secondary volume.

SnapMirror for Clustered Data ONTAP

Data replication can be performed within the same cluster or remotely to another cluster. Data ONTAP provides integrated data replication technologies for creating replica copies that can be used for the following purposes:

- Performing disaster recovery
- Offloading tape backup processes from the primary storage
- Distributing datasets to other locations
- Creating read/write clones for test and development environments

SnapMirror is integrated with NetApp Snapshot technology, which provides a method for quickly and efficiently creating on-disk replicas or point-in-time copies of data. SnapMirror replication is efficient because it replicates only the 4KB blocks that have changed or that have been added after the previous update. Additional efficiency is gained when SnapMirror is combined with other NetApp storage efficiency technologies. When FAS deduplication is used on the primary storage, only unique data is replicated to the DR site. If compression is enabled on the source, then compression is maintained on the destination. Data is not uncompressed when it is replicated. Storage efficiency technologies can result in telecommunication savings and significant storage capacity savings.

SnapMirror is an integral part of DR plans. If critical data is replicated to a different physical location, a serious disaster does not result in extended periods of unavailable data. Clients can access replicated data across the network until the damage caused by the disaster is repaired. Application servers at the recovery site can access replicated data and restore operations for business-critical applications for as long as necessary to recover the production site. Recovery operations might include recovery from data corruption, natural disaster at the production site, accidental deletion, and so on.

Note: Commvault IntelliSnap for NetApp offers a single interface for managing and scheduling SnapVault and SnapMirror relationships outside of Data ONTAP. Automatic provisioning of secondary storage or tertiary storage is managed by OnCommand Unified Manager in the back end.

Replication Options


Mirroring and vaulting strategies for replicating data can be architected in various ways. You can use mirroring and vaulting separately, or you can use them together and configure them through storage policies. A storage policy creates two data copies by default: a primary classic copy and a primary snap copy. The primary classic copy is used for tape copies. The primary snap copy relates to the Snapshot copy on the primary system. To vault or mirror that primary data, you must create additional copies in the storage policy:

- To create a mirror of the primary data, create a mirror copy that points to the primary snap copy as its source.
- To create a vault of the primary data, create a vault copy that also points to the primary snap copy as its source.


To vault the mirror copy, the source for the vault is set to the mirror copy. Figure 14 shows some of the replication combinations that can be configured.

Figure 14) Flexibility of replication combinations.

Replication Options



- SnapVault
- SnapMirror (asynchronous volume SnapMirror)
- Combinations (cascades)
- OnCommand handles replication
- Fan-out options



In a fan-out topology, a single primary volume is replicated to multiple destinations. This topology allows a single primary volume to be protected and backed up and provides a read-only copy at a secondary site. Clustered Data ONTAP supports the fan-out topology with a limit of four replication destinations for a single source. The limit of four destinations is shared between SnapMirror and SnapVault relationships: You can have any combination of SnapMirror and SnapVault relationships for a single source as long as the total number of relationships is not greater than four.

Note: Clustered Data ONTAP 8.2.1 supports higher fan-out ratios up to eight relationships.

Note: Cascading Snapshot > SnapVault > SnapMirror is not supported in Commvault IntelliSnap for NetApp.

Selective Copy

A selective copy allows you to use SnapVault to copy data from a specific source copy. The source copy can be either a primary copy or a synchronous copy. Selective copying gives you the freedom to create frequent Snapshot copies on the primary storage but only vault selective ones to a secondary location for backup purposes (for example, you could vault one Snapshot copy per day). This feature helps control secondary storage usage.

Selective copying is applicable only to full backups. You can define a selective copy to be:

- **Time-based.** For example, only the first or last full backup that occurs within the selected week, month, quarter, half-year, or year is copied. Note that, for these selection criteria, SnapVault considers the start time of the backup job.
- **Automatically selected.** For example, if the copy is defined with the setting All Fulls, all full backups from the primary copy are copied during an auxiliary copy operation.
- **Automatically not selected on the primary copy.** In this case, you can manually select the jobs that you want to be copied.

Data Movement to Tape

The Commvault IntelliSnap for NetApp solution supports several methods of replicating backup copies to tape, such as NDMP dump, SMTape, and tape streaming through MediaAgents.

Dump is a Snapshot copy-based backup and recovery solution that backs up files and directories to tape and restores them to a storage system. SMTape is a Snapshot copy-based backup and recovery solution that backs up the required metadata and the data blocks of a FlexVol® volume to tape and restores them to a storage system. You can enable tape backups in the storage policy properties; under the Snapshot tab, select Enable Backup Copy.

Note: Commvault IntelliSnap for NetApp never backs up data directly to tape or to cloud storage. For Commvault IntelliSnap for NetApp to do that, you must create a primary snap copy and then stream it—immediately delete it after the move of snap to media operation is completed, defer the operation to anytime, and source it from any additional snap copies. To stream data directly to tape or any other disk target, enable the Commvault capacity-based licensing model and manage everything from the same GUI

Tape backups can use the primary classic copy in the storage policy. Alternatively, you can create a new primary backup copy for the storage policy. Tape backups can be configured so that any of the Snapshot copy locations can be used as the source for the tape copy. You can select the source for the backup copy in the storage policy properties; under the Snapshot tab, modify Source Snap Copy.

The NetApp NAS NDMP iDA is the only agent that allows NDMP dump and SMTape backups. Tape copies from other iDAs stream through a MediaAgent in a more traditional backup fashion. Backups can be streamed to tape libraries that are connected to either the MediaAgent or the FAS controller.

Clustered Data ONTAP 8.2 has limited SMTape functionality for allowing SnapMirror and SnapVault relationships to be seeded by copying the initial baseline to tape and restoring it to the secondary storage. Seeding relationships is the only supported use of SMTape in clustered Data ONTAP 8.2. Clustered Data ONTAP 8.3 has introduced support for NDMP-based SMTape and SMTape block-level incremental backups. Commvault IntelliSnap for NetApp 10 SP9 and later leverages these new features.

The method to use for replicating full or incremental copies to tape depends on how the Commvault IntelliSnap for NetApp backup was initiated. To run incremental tape jobs, you must create incremental schedules for the Commvault IntelliSnap for NetApp backup. You can select the tape backup jobs to move to tape in the storage policy properties; under the Snapshot tab, modify Selection Rule.

Clustered Data ONTAP 8.2 and later supports SVM-aware NDMP backups. This backup mode optimizes NDMP backup performance by choosing efficient data transfer paths and being fully compatible with the integrated nondisruptive operations and volume mobility capabilities of clustered Data ONTAP. The

cluster-aware backup (CAB) extensions that are required to support this feature are available with the Commvault IntelliSnap for NetApp 10 SP8 release and later versions.

The clustered Data ONTAP CAB extensions define a mechanism and protocol within Data ONTAP for establishing efficient data connections for the backup and recovery operations of a storage cluster. The CAB extensions present details of the storage cluster's resource locations as a means to accomplish the following goals:

- Identify backup resources within the storage cluster:
 - The required backup resources may be located anywhere in the storage cluster.
 - The required backup resources may have been moved since a previous backup.
- Enable a local backup instead of a three-way backup when possible.
- Prevent data flow across the intracluster network.

Choosing an Appropriate Tape Backup Method

Depending on the requirements that you have for your tape backup, either the dump or the SMTape backup engine can offer you the most benefits:

- Use dump backup and restore if you want the following features:
 - Individual file and directory backup and recovery
 - Preservation of backups for several years
 - Exclusion of specific files and directories during a backup
- Use SMTape backup and restore if you want the following features:
 - A DR solution that provides high performance
 - SnapMirror tape seeding
 - Preservation of deduplication and compression settings after backup and restore operations
 - Backup and restore of volumes that have millions of files to achieve desired backup windows
 - Full-volume recovery only

For more information about NDMP configuration, refer to the following documents:

- [Clustered Data ONTAP 8.2 Data Protection Tape Backup and Recovery Guide](#)
- [TR-4330: Cluster-Aware Backup Configuration for Commvault IntelliSnap for NetApp and Simpana](#)

Data Movement to Cloud Using AltaVault

NetApp AltaVault cloud storage appliances can dramatically improve the capabilities of an existing backup infrastructure by leveraging cost-effective cloud-based storage and eliminating tape infrastructure. The AltaVault appliances can reduce the capital and operational costs of managing DR datasets by 30% to 50% and improve RTO/RPO reliability and efficiency.

The AltaVault appliance is a disk-to-disk data storage optimization system with unique cloud storage integration. It easily integrates with Commvault IntelliSnap for NetApp to securely protect critical production data off site without the complexity of using tape management solutions or the cost of using in-house DR sites and services. When you add an AltaVault appliance as a target for your backup infrastructure, the backup server connects to the AltaVault appliance by using the CIFS or the NFS protocol.

For a backup to an AltaVault appliance, the appliance performs in-line byte-level deduplication of the backup data and replicates the data to the cloud. AltaVault appliances use the local disk to store enough data for recovery of recent information. This mechanism provides LAN performance for the most likely restores. The AltaVault appliance then writes the backup data to the cloud storage. AltaVault supports the following cloud storage platforms and services:

- Amazon S3 and Glacier
- AT&T Synaptic Storage as a Service
- HP Cloud Services
- Google Cloud
- Microsoft Azure cloud storage
- Nirvanix Cloud Storage Network
- Rackspace Cloud Files
- General instances of EMC Atmos
- OpenStack (Swift) Object Storage

AltaVault appliances optimize restores from the cloud because they move only deduplicated data over the WAN. AltaVault is configured as a disk library in Commvault IntelliSnap for NetApp. Figure 15 and Figure 16 show the options to select in the Commvault IntelliSnap for NetApp (CommCell) GUI for adding AltaVault as a disk library.

Figure 15) GUI path for adding AltaVault to Commvault IntelliSnap for NetApp as a disk library.




Figure 16) AltaVault specified as a disk library in the Add Disk Library dialog box.





Figure 17 shows the Snapshot management options to configure for AltaVault backups. Select the Enable the Backup Copy option for your storage policy and specify the source copy to back up to the AltaVault appliance.

Figure 17) Snapshot management options for AltaVault.



Commvault IntelliSnap for NetApp can back up data in NetApp storage directly to AltaVault in a remote topology through the two native Data ONTAP NDMP engines, dump and SMTape. In a remote backup configuration, the disk library (AltaVault) is configured on the MediaAgent, and the backup is configured using the NAS iData agent (iDA) on the CommServe server.

For application-aware backups to the AltaVault device, Commvault IntelliSnap for NetApp uses the iDAs for the application in question. These backups are considered streaming backups.

3.2 Understanding the Backup Workflow

Commvault IntelliSnap for NetApp can manage backups for clustered Data ONTAP, Data ONTAP 7-Mode, and heterogeneous arrays from a single interface. Heterogeneous backups are backups from third-party storage to NetApp storage, and they are created by using Commvault IntelliSnap for open systems (OSDP).

To understand how the Commvault IntelliSnap for NetApp software works, it is important to understand each of these backup workflows, starting from the source data and working outward.


Note: The backup workflows are the same for clustered Data ONTAP and Data ONTAP 7-Mode.

Commvault IntelliSnap for NetApp Operations for Clustered Data ONTAP and Data ONTAP 7-Mode

Clients own source data and have specific iDAs, depending on the type of client and on the data being protected. Backup sets and subclients are configured within the iDA, and they group the source data to be protected. For example, if the volume `/vol/datavol` is protected on a client NetApp array, a subclient contains an entry for `/vol/datavol`.

The storage policy determines the behavior of the data protection operations and the retention properties. Each subclient is associated with a storage policy, which contains entries for the various copies in the data protection layout. In the example in Figure 18, the client data is protected by NetApp Snapshot copies. Vaulting is performed for longer term retention. The vaulted data is mirrored for redundancy, and tape copies are created from the mirror copy. The Commvault IntelliSnap for NetApp software orchestrates the operations, passing the vaulting and mirroring job control to the OnCommand Unified Manager server.

Figure 18) Commvault IntelliSnap for NetApp example workflow: vault > mirror > tape.



Commvault IntelliSnap for NetApp for Clustered Data ONTAP


Commvault IntelliSnap for NetApp for clustered Data ONTAP allows you to create Snapshot copies and perform SnapMirror and SnapVault replication from the same interface that you use to manage your 7-Mode and OSDP workflows. For example, you can protect a clustered Data ONTAP volume by making its Snapshot copy and replicating the volume to create a backup copy or a DR copy.

Clustered Data ONTAP supports many protocols (such as NDMP, and NAS and SAN protocols), allowing users to protect volumes, NFS datastores and exports, and applications hosted on LUNs. Clustered Data ONTAP 8.2 added support for the following functionalities:

- SnapVault and SnapMirror workflows for SVM intracluster and intercluster replication
- SVM-aware NDMP backups
- CAB extensions

Figure 19 shows the Commvault IntelliSnap for NetApp workflow for clustered Data ONTAP.

Figure 19) Commvault IntelliSnap for NetApp for clustered Data ONTAP backup workflow.



The Commvault IntelliSnap for NetApp backup workflow in Figure 19 has the following stages:

1. Data is quiesced and protected through application-consistent Snapshot copies.
2. Snapshot copies and clones are used to access data for indexing. You can use the deferred indexing
 - NAS-only option to run indexing on the secondary storage at a later time, after the backup job has completed.
- Note:** To run NAS Snapshot copy indexing in Commvault IntelliSnap for NetApp, you must have clustered Data ONTAP 8.2 P3 or a later version. The NAS live browse feature is not available in earlier versions of clustered Data ONTAP, so if you defer indexing on the secondary storage, you are not able to browse or have a catalog of the primary NAS data.
3. OnCommand Unified Manager 6.x handles the provisioning of the secondary storage, using resource pools and provisioning policies for replication.
4. Disk-to-disk replication: SVM-aware backups are created by using SnapVault or SnapMirror (support for both intracluster replication and intercluster replication).
5. Disk-to-disk-to-tape replication: SVM-aware NDMP backups are moved to tape media (support for CAB extensions).
6. Disk-to-disk-to-cloud replication by using AltaVault: Data is moved to the cloud.

In order for Commvault IntelliSnap for NetApp to discover volumes on an SVM, you must add the cluster management LIF of the cluster that contains the SVM to the CommCell console in the Array Management dialog box.

Note: Make sure that every SVM has a configured LIF for management access.

After the cluster management LIF is added to the SVMs, Commvault IntelliSnap for NetApp can automatically detect SVMs, and you do not need to manually add the SVM management LIF.

Note: You must create resource pools and SVM associations from the Storage menu in the OnCommand Unified Manager 6.0 interface. The other settings can be configured and managed within the Commvault IntelliSnap for NetApp interface.

Scheduling and Retention

You can schedule backups by creating individual schedules or schedule policies. A schedule policy groups various schedules together, each with its own properties. For example, a schedule policy might contain individual schedules for daily, weekly, and monthly backups.

Traditional backup scheduling typically calls for weekly full backups and daily incremental backups. With NetApp Snapshot copy technology, however, the Commvault IntelliSnap for NetApp model consists almost entirely of full backups. An exception to this scenario is NAS data protection. The indexing performance of NAS data backups increases significantly for incremental backups while maintaining seamless searching for single- file recoveries across Snapshot copies. In addition, Data ONTAP 7-Mode systems support the live browse feature for NAS data. Live browsing allows users to opt out of the indexing process (for example, when indexing millions of files would be intensive or lengthy) and simply browse the backups in real time, according to their needs.

Auxiliary copies can be scheduled on specific times, or they can be configured to run automatically.

Figure 20 shows a storage policy and its associated snap copies. In this example, a subclient is scheduled to perform a full backup (a local Snapshot copy) each day at 6 p.m. Retention for these local Snapshot copies is configured in the primary snap copy. The policy establishes a retention model of 10 daily backups and 6 weekly backups.

Mirroring takes place each day at 6:30 p.m. Retention for the mirror copy matches the retention of the primary copy. Vaulting starts when mirroring finishes. Retention for the vault destination is configured in the vault copy. The policy establishes a retention model of 90 daily backups and 52 weekly backups.

Figure 20) Example schedule and retention at specific times.

Storage Policy

Schedule	Snap Copy	Retention
Daily full at 6:00 p.m.	Primary (Snap)	10 days, 6 weeks
Daily at 6:30 p.m.	Mirror Copy	10 days, 6 weeks
	Vault Copy	90 days, 52 weeks

Commvault IntelliSnap for NetApp has two types of retention rules: basic retention rules and extended retention rules. Basic retention rules apply to daily or hourly backups. Extended retention rules apply to longer term retention, such as weekly full backups, monthly full backups, and yearly full backups. These rules are configured in the storage policy and can be set for the primary snap copy, vault copies, and tape copies. Mirror copies do not allow specific retention settings because they inherit the retention settings of the primary copy.

A backup cycle represents a full backup and the incremental backups that depend on that full backup. In many cases, full backups are used for every backup cycle. However, for NAS data that contains millions of objects, a strategy that includes incremental backups improves indexing performance. In addition, if

backup jobs require moving incremental copies to tape, the Commvault IntelliSnap for NetApp backups on the primary storage must include incremental jobs. In a full-backup-only paradigm, each backup can be considered a backup cycle.

When incremental backups are included in a cycle, all of the Snapshot copies in that cycle are retained until the last incremental backup in the cycle has expired. Performing more frequent full backups reduces the number of Snapshot copies associated with a cycle. Basic retention rules allow retention entries for days and cycles. The default setting is 7 days and 2 cycles.

Extended rules can be applied for longer retention. These rules include options to keep all full backups: weekly full backups, monthly full backups, quarterly full backups, half-year full backups, and yearly full backups. Extended rules are not tied to a particular backup schedule. Rather, they are tied to full backups that start on a particular day of the week or of the month. These days can be chosen as required.

To perform both hourly backups and daily backups, you must create separate backup sets and storage policies. One backup set should include a subclient with the hourly schedule and associated with one storage policy. The other backup set should include a subclient with the daily schedule and associated with the other storage policy. When running hourly backups, you must change the data aging schedule to run hourly instead of the default setting of once per day. The data aging operation is what expires backups and deletes Snapshot copies.

Note: Many of the Commvault IntelliSnap for NetApp application iDAs do not allow separate backup sets.

Enterprise Scheduling Policy Examples

The following examples describe how to keep 6 hourly backups, 30 daily backups, weekly backups for 3 months, and monthly backups for 1 year on the NetApp primary system. These examples assume that only full backup jobs are being scheduled:

- **Hourly backups with 6-hour retention.** Create a daily schedule for the subclient that repeats every hour, then set a basic retention rule in the primary snap copy of the associated storage policy to retain 0 days and 6 cycles. By default, the data aging schedule runs once per day; therefore, to expire backups based on hourly retention, the data aging schedule needs to run hourly.
- **Daily backups with 30-day retention.** Create a daily schedule for the subclient that repeats every day, then set a basic retention rule in the primary snap copy of the associated storage policy to retain 30 days and 30 cycles.
- **Retaining weekly backups for 3 months.** Retain one daily backup every week and keep it for 90 days. This schedule requires an extended retention rule. Set an extended retention rule in the primary snap copy of the associated storage policy to retain weekly full backups for 90 days and set the rule to start on the appropriate day of the week. Every daily full backup created on this day of the week is retained for 90 days.
- **Retaining monthly backups for 1 year.** Retain monthly backups by retaining one daily backup every month and keeping it for 365 days. This schedule requires an extended retention rule. Set an extended retention rule in the primary snap copy of the associated storage policy to retain monthly full backups for 365 days and set the rule to start on the appropriate day of the month. Every daily full backup created on that day of the month is retained for 365 days.

Schedules and retention for replication, tape backups, and cloud backups are also set on the storage policy:

- **Replication (mirroring and vaulting).** You can use methods similar to the ones in the scheduling examples to schedule replication and vault retention. However, scheduling is set on the mirror or vault copy in the storage policy rather than in the subclient; for vaulting, retention is set in the vault copy in the storage policy.

- **Tape backups.** The backup to tape is the last hop in the backup cycle. The source of the tape backup can be either the primary Snapshot copy or the vault or mirror copy. Schedules and retention are set in the storage policy.
- **Cloud backups.** Data is backed up to the AltaVault appliance by Commvault IntelliSnap for NetApp and later moved to cloud storage, such as Amazon storage and Microsoft Azure storage, by AltaVault. Schedules and retention for the data on AltaVault are set in the storage policy.

Note: Commvault IntelliSnap for NetApp sets retention for backup data that is in the AltaVault appliance; the cloud retention is, in turn, set on the AltaVault appliance. You must set the same level of retention on both Commvault IntelliSnap for NetApp and AltaVault. For example, if your Commvault IntelliSnap for NetApp backup retention is set to 1 year and the AltaVault to cloud retention is set to 6 months, then there is a data loss scenario for Commvault IntelliSnap for NetApp.

Table 3 contains an expanded scheduling example in which virtual machines (VMs) in VMware datastores are protected with various requirements. Four different storage policies are required in this example because the datastores have mixed retention requirements.

Table 3) Enterprise example of scheduling and retention.

Backup Set	Data-store	Sub-client	Stor. Policy	Backup Sched.	Local Retention	Mirror Sched.	Mirror Retention	Vault Sched.	Vault Retention
A	DS1	SC1	SP1	Daily full at 6 p.m. ¹	10 days (cycles), 6 weeks, set in primary snap copy	Daily at 6:30 p.m., set in mirror copy schedule	10 days, 6 weeks	After mirror finishes	90 days (cycles), 52 weeks, set in vault copy
A	DS2	SC2	SP2	Daily full at 6 p.m. ¹	30 days (cycles), 8 weeks, set in primary snap copy	Daily at 6:30 p.m., set in mirror copy schedule	30 days, 8 weeks	After mirror finishes	180 days (cycles), 52 weeks, set in vault copy
A	DS3 ²	SC3	SP3	Daily full at 6 p.m. ¹	10 days (cycles), set in primary snap copy	Daily at 6:30 p.m., set in mirror copy schedule	10 days	After mirror finishes	90 days (cycles), 52 weeks, set in vault copy
B	DS3 ²	SC4	SP4	Hourly, except at 6 p.m.	23 hours (cycles), set in primary snap copy	—	—	—	—
Table Notes									

¹ The local Snapshot schedules can be set at the subclient level or at the backup set level. In the example in Table 3, all local Snapshot copies run at 6 p.m. Therefore, a single schedule at the backup set level can be used. If subclients in a backup set require different local Snapshot schedules, then the schedules must be set at the subclient level.

² Datastore DS3 is defined in two backup sets (A and B) because of the need to perform both daily backups and hourly backups for this datastore. Therefore, for the basic retention rules to work, it is necessary to have the same datastore in the two subclients. Using two backup sets in this case enables retention for both hourly and daily backups.

3.3 Understanding the Restore Workflow

With Commvault IntelliSnap for NetApp management software, recovery is simple: You can restore data from virtually any backup copy in a single operation. Restores from recent backups might come from local Snapshot copies (by leveraging optimal single-file Snap Restore® technology), while historical data might come from vault or tape copies.

You can locate the data to be restored either by browsing or by using the find feature. When the data is located, a restore job can be initiated. You can perform restores from any of the backup copies by browsing data from a particular copy. The copy order precedence is defined on the Copy Precedence tab of the storage policy properties.

For volumes and LUNs on NetApp primary storage, it is also possible to revert the data in them to a specific Snapshot copy. This feature uses NetApp SnapRestore data recovery software to revert a volume or LUN back to a particular point in time. This option should be used with caution because a revert operation affects all data in a volume or LUN. To initiate a revert, right-click the subclient and select List Snaps. From the list, right-click a Snapshot copy and select Use Hardware Revert Capability if Available.

Because Commvault IntelliSnap for NetApp uses NetApp Snapshot technology, you can copy data directly from a Snapshot copy by using CIFS or NFS.

Note: The single-file SnapRestore option is valid only for NAS datastores. Hardware reverts should be used with caution because this option completely replaces existing data with the previous Snapshot copy.

3.4 Data Cloning

Commvault IntelliSnap for NetApp software enables administrators to create data clones, which allow read/write access to the backup data. Data clones can be used for a variety of purposes. NAS data can be cloned when you use the NetApp NAS NDMP iDA. This functionality creates a NetApp FlexClone® volume and makes the content of the Snapshot copy accessible.

The file system iDA can also be used to clone LUN data. This functionality, when applied to a primary Snapshot copy, uses LUN clone technology. When LUN data from a Snapshot copy on a secondary or tertiary NetApp system is cloned, a FlexClone volume is created. These FlexClone volumes can be used for test/dev purposes and do not consume any additional space until written to, and, even then, only the deltas occupy additional capacity.

Note: To create data clones, right-click the subclient and select List Snaps. From the list, right-click a Snapshot copy and select Mount.

You can duplicate SQL Server or Oracle databases by cloning the Snapshot copies that were created during the database backup. This operation uses the cloning capability of the storage array hardware, which enables you to duplicate large databases within a short period of time. Database clones can be used for multiple purposes, such as the following:

- In test environments to troubleshoot issues found in the production database

- For quick data retrieval without the need to run resource-intensive restores on the production environment and without consuming additional space on the destination server
- For alleviating the load on production servers for running reports and queries

You can create a clone from any full backup. During the clone creation, you can specify a reservation period. At the end of the reservation period, the system automatically shuts down the clone database and frees up all resources. You can also schedule the cloning operation to run periodically to allow the clones to get refreshed from the latest backups on a regular basis.

The SQL Server and Oracle databases can be cloned to the same instance or to a different instance.

Note: You can create database-consistent clones by right-clicking the subclient (SQL Server or Oracle Database), navigating to All Tasks, and selecting Clone.

3.5 OSDP for Data ONTAP 7-Mode Target

OSDP adds support for backup and granular recovery of DAS and third-party data to NetApp storage. The target for the backup is a Data ONTAP 7-Mode storage system. OSDP has block-level incremental replication capabilities but requires full file system scans. OSDP supports select nonclustered hosts and applications, such as Windows file systems, Linux, Solaris, SQL Server, Exchange Server, and Oracle Database.

OSDP replicates data from UNIX and Windows clients by using source volumes carved out of either disks that are locally attached to the clients or disks that are externally attached to the clients from third-party arrays that are connected to a Data ONTAP 7-Mode destination as a single LUN, with a single partition onto a SnapVault copy. OSDP uses OnCommand Unified Manager to provision the destination volumes to which data is replicated. In Windows, data is replicated as GPT LUNs. In UNIX, a LUN keeps the same format as the source volume LUN.


Running OSDP jobs from the replication workflow enables the use of a single NetApp volume as a fan-in destination for many clients under a single dataset context. It allows a group of clients to be controlled by a single job management context, providing many-to-one control. This group defines the operations job unit that triggers the recovery point for Snapshot copies on the SnapVault destination.

Using a consolidated dataset with SnapVault automatically pulls in any efficiencies that NetApp deduplication can achieve across the collection of replicated disk images that are maintained with OSDP.

Note: A valid IntelliSnap for NetApp controller license is required on the FAS or V-Series controllers running Data ONTAP 7-Mode that are the targets for OSDP backups. If the target for OSDP is the FAS primary system (which is already licensed as primary storage for IntelliSnap for NetApp), then no additional IntelliSnap for NetApp license is required.

Figure 21 shows the workflow for heterogeneous backups created with OSDP software.

Figure 21) OSDP backup workflow.



The workflow for OSDP subclient backup jobs includes the following stages:

1. OSDP creates a software Snapshot copy of the source data by using a native snap engine, such as VSS or QSnap or Logical Volume Management (LVM).
2. The volume is added to the OnCommand Unified Manager dataset that corresponds to the primary snap copy of the storage policy to which the subclient is associated.
3. OnCommand carries out any provisioning tasks and requests for the destination controller to connect back to the client for the OSDP transfer or the OSDP test process.
4. If the volume has never been protected before, the OSDP transfer performs a baseline transfer of all the data blocks in the volume to the destination controller.
5. If the volume has been protected before, the OSDP transfer or OSDP test figures out which blocks have changed in the volume since the last backup and then transfers only the changed blocks to the destination file server. For a given volume, all backups, except for the first one, are incremental.
6. When the OSDP transfers to all of the associated subclients are complete, a destination Snapshot copy is created on the destination vault copy, and this Snapshot copy is registered as the primary data copy for that application against the primary snap copy.
7. The software Snapshot copy that was created at the start of the IntelliSnap for NetApp backup job is deleted after the OSDP transfer of the source Snapshot copy is completed successfully.

3.6 OSDP for Clustered Data ONTAP Target

OSDP for clustered Data ONTAP leverages (NetApp open systems backup [NOSB]) NetApp for open systems is a full-volume, block-level incremental replication feature. It replicates heterogeneous data from source partitions and volumes from UNIX and Windows clients to a NetApp clustered Data ONTAP destination storage system as a single LUN with a single partition onto a destination qtree on a single destination volume. The source volumes could be carved out of disks that are locally attached to the

client or from disks that are externally attached to the client from third-party hardware arrays. It uses NetApp OnCommand Unified Manager to provision the destination volumes to which data is replicated. Data is replicated as GPT LUNs on Windows. On UNIX, a LUN has the same format as the source volume LUN.

Running NetApp for open systems jobs from the replication workflow enables the use of a single destination NetApp volume as a fan-in destination for many clients under a single storage policy context. This process allows a group of clients to be controlled by a single job management context (providing a many-to-one control). This group defines the operations job unit that triggers the recovery point Snapshot copies on the destination.

Using a consolidated destination automatically pulls in any efficiencies that NetApp deduplication can achieve across the collection of the replicated disk images that are maintained with NetApp for open systems.

The NetApp for Open Systems replication occurs within the context of an IntelliSnap for NetApp backup. To carry out the replication transfer, you should run or schedule the replication workflow from the storage policy to which all the clients that need to be protected are associated. The workflow in turn starts individual IntelliSnap for NetApp backup jobs for each of the associated subclients. Each IntelliSnap for NetApp backup job for a subclient performs the following tasks:

- OCUM is used while configuring the storage policy copy to create a single destination volume corresponding to the storage policy to hold all the data for the source clients that are associated to the same storage policy.
- Creates a software Snapshot copy of the source data using the native snap engine such as VSS or LVM.
- The backup process mounts the destination qtree that corresponds to its source volume from the destination volume.
- If the volume has never been protected, then the NetApp for open systems agent performs a baseline transfer, such as transferring all blocks for the volume to the destination storage system.
- If the volume has been protected, then the NetApp for open systems agent determines the blocks that have changed for the volume since the last backup by using an SHA-1 checksum mechanism and a checksum database and transfers only the changed blocks to the destination file server. For a given volume, all backups, except the first, are incremental.
- Data is transferred using the NFS protocol wherein the client transferring the data acts as the NFS client and the storage system as the NFS server.

When NetApp for open systems transfer to all of the associated subclients is complete, a destination Snapshot copy is taken on the destination volume. This snap is registered as the primary data copy for that application against the primary snap copy. The software snap that was created at the start of the IntelliSnap for NetApp backup job is deleted after the NetApp for open systems transfer for source Snapshot copy is successful.

4 Application Data

Commvault IntelliSnap for NetApp can be used to protect applications running on physical servers and hosted on NetApp primary storage. Each supported database application has an associated iData agent (iDA). This iDA must be installed on the client system that is running the application. The iDAs prepare the database applications for backup consistency. In addition, they handle tasks such as log truncation during backup, database storage mapping, and log manipulation during restore.

Table 4 lists the applications that Commvault IntelliSnap for NetApp can protect by using iDAs and the level of restore granularity that it can achieve.

Table 4) Application support, iDAs, and restore granularity.

Application	iDA	Restore Granularity
Protection with Application Agents		
NAS	NetApp NAS NDMP	Qtrees, directories, and files ¹
VMware vSphere	Virtual server	VMs, VMDKs, and files ²
Hyper-V	Virtual server	VMs and files
SQL Server	SQL Server	Database and file groups ³
Exchange Server	Exchange database	Information store and DAG objects ⁴
Oracle Database	Oracle	Database and tables
DB2	DB2	Database
SharePoint	SQL, SharePoint Server	Database and objects ⁵
Active Directory	Active Directory	Objects ⁶
Lotus Domino	LN database, LN document ⁷	Database and database plus transaction logs; documents ⁸
Protection with Virtualization Agents		
Exchange Server	Virtual server	VMs, VMDKs, and files ⁹
SQL Server	Virtual server	VMs, VMDKs, and files ⁹
Table Notes		
<p>¹ The NAS iDA can restore qtrees, directories, and files. In addition, it can use SnapRestore on the primary array to revert an entire volume. The live browse functionality for Snapshot copies is available without indexing.</p> <p>² Single-file restore (from a NetApp Snapshot copy) for VMs works only for Windows VMs. For Linux, the restore works only when the backup is to a disk library or tape library (streamed copy). Using IntelliSnap for NetApp to perform granular indexing or a single-file restore for Linux VMs from a NetApp Snapshot copy is not supported.</p> <p>³ SQL Server supports point-in-time restore with log replay.</p> <p>⁴ For Exchange Server, message-level restores are available through mining operations. You can use the offline mining tool, or you can configure and perform snap mining.</p> <p>⁵ For SharePoint, document-level restores are available through snap mining operations.</p> <p>⁶ Active Directory backups are accomplished by streaming the backup from the Active Directory server. You can also create a system-state backup by streaming the backup through the use of the Windows file system iDA. This operation requires that you disable the system-state backup from the default subclient and enable the system-state backup in a new subclient (with no other content defined).</p> <p>⁷ LN stands for “Lotus Notes.”</p> <p>⁸ The LN database iDA supports restore operations for the database only or for the database plus transaction logs. The LN document iDA supports the restore of documents. For more detailed information about these restore options, refer to the IntelliSnap for NetApp online documentation (Books Online).</p>		

⁹ When Exchange Server or SQL Server is virtualized and the databases are on VMDKs, you can use the virtual server agent (VSA) to back up the VMs and perform database consistency tasks during the VM backup. Recovery works for VMs, VMDKs, or files and is not specific to Exchange Server or SQL Server. The options for Exchange Server allow log truncation during the VM backup. SQL Server logs cannot be truncated when you use the VSA approach. When Exchange Server is protected in this manner (that is, by using VSA), DAGs are supported.


5 Virtualization Data

A key feature of Commvault IntelliSnap for NetApp management software is the ability to protect many VMs very quickly. Commvault IntelliSnap for NetApp can index the contents of each VM and allows different levels of recoverability, including single-file recovery.

Commvault IntelliSnap for NetApp software is flexible and enables you to establish discovery rules so that new VMs can be automatically added to a subclient and protected. For example, the Datastore Affinity discovery rule automatically protects new VMs on specific datastores.

Commvault IntelliSnap for NetApp software uses the VSA to perform data protection operations for virtual environments. The VSA is installed on a system configured as a MediaAgent. Within the VSA, Commvault IntelliSnap for NetApp creates instances that define the type of virtualization solution being used. For example, in a VMware environment, Commvault IntelliSnap for NetApp creates a VMware vCenter instance under the VSA. Within the instance, a backup set contains the subclients. Figure 22 shows the VSA layout.


Figure 22) VSA layers.



Note: Because of the advantages that the VMware HotAdd transport mode can provide during restores, NetApp recommends installing the VSA on a virtualized MediaAgent. This virtualized MediaAgent should run on a VMware ESXi host that has access to production datastores, such as an ESXi proxy host.


In the example shown in Figure 23, multiple datastores are grouped into a single backup set. However, because the datastores have different scheduling and retention requirements, they are separated into their own subclients, and each subclient is associated with a different storage policy. The datastores are mirrored and then vaulted.

Figure 23) Datastores in separate subclients with different storage policies.



In Figure 24, the datastores are grouped into a single backup set and a single subclient. In this example, the datastores have the same scheduling and retention requirements. The datastores are mirrored and then vaulted. Because the datastores are grouped into the same subclient, it is possible to do a fan-in on the vault copy.

Figure 24) Datastores in the same subclient with fan-in.



Backup settings can have different levels of granularity for restore operations. During restores, data for the VMs can be browsed and recovered based on the selected recovery type. A container restore can be performed to recover an entire VM. Individual files can also be restored for Windows VMs and may leverage the single-file SnapRestore option if the files reside on primary storage.

5.1 VMware and Applications

When run inside a VM, Exchange Server and SQL Server have integration with the Volume Shadow Copy Service (VSS), which maintains database consistency during the backup of the VM. The file system iDA and the VSS provider must be installed on the guest operating system for this functionality to be available.

Note: To enable these application-consistent backups, verify that the Application Aware Backup for Granular Recovery checkbox is selected under the Commvault IntelliSnap for NetApp Operations

tab for the subclient. Exchange Server backups offer the additional option to perform log truncation as part of the backup operation. To enable this option, select Truncate ExDB Logs.

You can perform consistent out-of-place restores of SQL Server and Exchange Server databases by restoring the flat database files. The Exchange Offline Mining tool is a standalone utility that allows the restore of individual messages from a backup copy of the Exchange Server database.

6 Commvault IntelliSnap for NetApp Network and Security Considerations

All computers in a Commvault IntelliSnap for NetApp environment—the CommServe server, the MediaAgents, the OnCommand Unified Manager server, and the client machines—must be connected through a TCP/IP network.

In a Commvault IntelliSnap for NetApp deployment, the bulk of the backup data is processed at the storage level. Therefore, defining preferred interface options for replication at the storage level isolates the backup traffic from production networks. The backup VLAN in such deployments is used to transfer index information to MediaAgents, backup-to-tape use cases, and a few of the restore workflows.

CommCell components separated by a firewall must be configured to reach each other through the firewall by using connection routes. After they are configured, they can communicate to perform data management operations such as backups, browsing, and restores.

6.1 Firewall Considerations

A Commvault IntelliSnap for NetApp deployment has the following port requirements:

- **Connections to OnCommand Unified Manager 6.x:**
 - HTTP port 80: Unified Manager web UI (redirects to HTTPS port 443)
 - HTTPS port 443: Unified Manager web UI and programs using APIs
 - SSH/SFTP port 22: maintenance console
 - MySQL port 3306: OnCommand workflow automation and OnCommand report access
- **Connections from OnCommand Unified Manager 6.x:**
 - HTTPS port 443: storage systems
 - HTTPS port 443: AutoSupport™ server
 - LDAP port 389: authentication requests for users and groups
 - SMTP port 25: alert notification e-mails
 - SNMPv1 or SNMPv3 port 162 (UDP): alert notifications (SNMP traps)
 - NTP port 123 (UDP): time synchronization

For more information about firewall requirements and configuration between clients, MediaAgents, and the CommServe host, refer to [Firewall: Getting Started](#) in the Commvault IntelliSnap for NetApp online documentation.

7 Commvault IntelliSnap for NetApp for E-Series


Commvault IntelliSnap for NetApp supports Snapshot integration for E-Series storage systems, but open systems data protection (OSDP) can be used to back up E-Series workloads to a Data ONTAP 7-Mode target and clustered Data ONTAP targets, go [here](#) for more details

E-Series storage systems can also be configured as backup targets in Commvault IntelliSnap for NetApp with software compression enabled. These backups are considered streaming backups.

8 Role-Based Access Control for Commvault IntelliSnap for NetApp


You can associate new users added under security in Commvault IntelliSnap for NetApp with customized user groups and the appropriate privileges and entities to manage the backup infrastructure. For example, an Oracle database administrator can be associated with that administrator's clients, storage policies, and disk library so that the administrator is isolated from the enterprise backup and recovery window. Figure 25 shows the capabilities that can be associated with clients and entities in Commvault IntelliSnap for NetApp.

Figure 25) Capabilities to associate with clients and entities.



The CommServe server can be configured to integrate with existing Active Directory and Domino DNS services with SSO enabled. Figure 26 shows the GUI path for setting these configurations.

Figure 26) GUI path for integrating CommServe with Active Directory users and enabling SSO.



9 Commvault IntelliSnap for NetApp for MetroCluster

NetApp highly available pairs couple two controllers for protection against single controller failures. NetApp disk shelves have built-in physical and software redundancies, such as dual power supplies and RAID DP (double parity) technology. NetApp HA pairs and shelves protect against many data center failures but cannot guarantee extreme levels of availability.


NetApp MetroCluster™ technology layers additional protection onto existing NetApp HA pairs to provide extreme levels of availability. Synchronous data mirroring enables zero data loss, and automatic failover enables nearly 100% uptime. Therefore, MetroCluster provides a zero RPO and a near-zero RTO.

NetApp HA pairs leverage takeover functionality, otherwise known as cluster failover, for protection against controller failures. When one of the controllers in an HA pair fails, the surviving controller takes over the failed controller's data-serving operations while continuing its own data-serving operations.

Controllers in an HA pair use the cluster interconnect to monitor partner health and to mirror NVLOG information composed of recent writes not propagated to disk.

Figure 27 illustrates cluster failover for a MetroCluster configuration.

Figure 27) Cluster failover.



9.1 Commvault IntelliSnap for NetApp Backup Configuration for MetroCluster

The cluster failover scenario exemplified in Figure 27 has the following characteristics:

- Site A is the active site for the VS1 SVM, and site B hosts the mirrored VS1-MC SVM that is locked in normal mode.
- Site A and site B have different cluster management interfaces.
- The VS1 and VS1-MC SVMs have the same interfaces (management and data).

In site A, assume that the backup configuration is primary Snapshot copy > mirror > vault. The mirror destination is SVM (secondary) VS-Mirror, and the vault destination is SVM (tertiary) VS-Vault.

For this scenario, configure OnCommand Unified Manager for MetroCluster in the following way:

1. Define the following SVM associations for both `VS1` and `VS1-MC` in OnCommand Unified Manager:
 - `VS1` > `VS-Mirror (SnapMirror)`
 - `VS1` > `VS-Vault (SnapVault)`
 - `VS1-MC` > `VS-Mirror (SnapMirror)`
 - `VS1-MC` > `VS-Vault (SnapVault)`
2. Make the resource pool configuration the same as the configuration for a normal CommVault IntelliSnap for NetApp backup workflow.

Configure the CommServe server and client for MetroCluster in the following way:

3. Add the cluster in site A in the Array Management dialog box. Make sure that the client name is resolvable to the IP address of the cluster in site A.
4. Add the cluster in site A as a client and configure the NDMP details.
5. Detect the SVMs through the Storage Virtual Machine tab in the client properties. These detected SVMs are added in the Array Management dialog box and have tunneling enabled by default.
6. In the switchover scenario, make sure that the client name of the cluster in site A is resolvable to the new IP address of the cluster in site B from the CommCell environment (CommServe server, client, and MediaAgent). The NDMP details should be updated with the information for the cluster in site B.
7. In the switchover scenario, no changes to the SVM names are necessary.

In the normal state, all configured backups point to `VS1` in site A; during switchover, the backups point to `VS1-MC` in site B. The switchover to site B is seamless and transparent to CommVault IntelliSnap for NetApp and does not require that you perform a rebaseline of relationships.

10 CommVault IntelliSnap for NetApp Disaster Recovery

Different methods are available to keep the databases on the CommServe hosts synchronized. Based on the recovery point objective of your organization, you can use one of following disaster recovery (DR) methods.

10.1 Disaster Recovery Backups

In this method, you can run DR backup jobs to protect the production CommServe database. The DR backups include the CommServe database and other databases hosted by CommServe, such as the SRM database (if metrics reporting is enabled) and the Workflow Engine database. By default, a full DR backup is run for these databases every day. These DR backups are stored in an export location or moved to a media backup for later retrieval.

When the production CommServe host is not accessible, you can restore the most recent DR backup of the production CommServe database to an alternate or standby CommServe host and resume CommCell operations by using the Disaster Recovery tool.

10.2 Disaster Recovery Failover

In this method, a standby CommServe host is set up in a remote location, and the production databases hosted by CommServe are replicated to the standby CommServe host at regular intervals. When the production CommServe host goes offline, you can immediately fail over the CommServe functionality and resume CommCell operations on the standby CommServe host. This failover process can be initiated manually or automatically.

10.3 Alternate Disaster Recovery Method

Optionally, you can use the SQL Server database mirroring feature to keep the databases of the production CommServe host and the standby CommServe host in a near-synchronized state. Committed SQL Server transactions on the production CommServe host are immediately applied to the standby CommServe host. The standby CommServe host can be activated by automatic or manual failover or by breaking the mirror relationship and bringing the standby CommServe host databases online.

For detailed instructions about how to configure disaster recovery, refer to [CommCell Disaster Recovery](#).

11 Summary

NetApp Commvault IntelliSnap for NetApp management software is changing today's backup and recovery landscape. Commvault IntelliSnap for NetApp software combines simplified manageability, power, and flexibility for virtual environments with full support for enterprise database applications. Commvault IntelliSnap for NetApp integrates with NetApp Snapshot technology in a virtually seamless way for fast and efficient backup operations and with NetApp SnapVault and SnapMirror software to support content cataloging and data movement to tape-based media.

This document covered an introduction to the Commvault IntelliSnap for NetApp solution by providing an overview of the technology and describing the basic options that you must configure to get started with Commvault IntelliSnap for NetApp. Commvault IntelliSnap for NetApp offers single-interface management for backup and recovery workflows for Data ONTAP 7-Mode, clustered Data ONTAP, limited third-party storage to NetApp storage (OSDP), and much more. Because it centralizes all of these functions and offers policy-based management and granular recovery across all supported workloads, Commvault IntelliSnap for NetApp is a compelling enterprise backup and recovery solution.

Resources

- For more detailed installation information about Commvault IntelliSnap for NetApp, refer to the Commvault IntelliSnap for NetApp online documentation (Books Online)
http://documentation.commvault.com/commvault/v11/article?p=products/netapp/c_netapp_overview.htm
- NetApp partners can find additional Commvault IntelliSnap for NetApp information, ranging from training presentations, best practice guides, , datasheets, technical reports, or FAQ documents, on Field Portal:
<https://fieldportal.netapp.com/Modules/FieldPortal/Binders/Content.aspx?contentID=344605>

Version History

Version	Date	Document Version History
Version 4	March 2015	<p>Document updated with information about the following topics:</p> <ul style="list-style-type: none">• Clustered Data ONTAP concepts• IntelliSnap for NetApp cataloging and indexing• SnapMirror for clustered Data ONTAP• IntelliSnap for NetApp integration with AltaVault• Networking considerations for IntelliSnap for NetApp• E-Series and IntelliSnap for NetApp• RBAC for IntelliSnap for NetApp• IntelliSnap for NetApp for MetroCluster• Disaster recovery methods

Version	Date	Document Version History
Version 5	February 2016	<ul style="list-style-type: none"> • Changed all references of SteelStore to AltaVault. • Made branding name change: SnapProtect to Commvault IntelliSnap for NetApp. • Added licensing information. • Updated all graphics.
Version 6	July 2017	<ul style="list-style-type: none"> • Minor update.

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 © 2004–2017 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.