



Technical Report

# SnapVault Best Practices Guide

## Clustered Data ONTAP

Amit Prakash Sawant, NetApp  
December 2013 | TR-4183

### Abstract

This document helps those who are familiar with NetApp<sup>®</sup> SnapVault<sup>®</sup> operating in 7-Mode to understand the changes that they need to be aware of in the clustered Data ONTAP<sup>®</sup> version of SnapVault. This document is an overview of the changes and contains links to other NetApp documentation for further details.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b> .....	<b>3</b>
<b>2</b>	<b>Clustered Data ONTAP Nondisruptive Operations and SnapVault</b> .....	<b>3</b>
<b>3</b>	<b>Storage Virtual Machines and Cluster Peering</b> .....	<b>3</b>
<b>4</b>	<b>SnapVault and Storage Efficiencies</b> .....	<b>4</b>
<b>5</b>	<b>Fan-In and Fan-Out SnapVault Topologies</b> .....	<b>4</b>
5.1	Fan-In Topologies.....	4
5.2	Fan-Out Topologies.....	4
<b>6</b>	<b>Cascading SnapVault and SnapMirror Relationships</b> .....	<b>5</b>
<b>7</b>	<b>Scheduling Snapshots and SnapVault Transfers</b> .....	<b>5</b>
<b>8</b>	<b>Interoperability with NDMP and SMTape</b> .....	<b>5</b>
8.1	SMTape Seeding.....	5
8.2	NDMP .....	5
<b>9</b>	<b>Management with System Manager</b> .....	<b>6</b>

## 1 Introduction

SnapVault is NetApp's disk-to-disk backup solution that is built into NetApp Data ONTAP. Enabling SnapVault on your NetApp system is as simple as installing a license key; no additional hardware or software must be installed. SnapVault allows you to replicate your data to a secondary volume and to retain the data for a longer period of time than you might on your primary volume.

SnapVault for clustered Data ONTAP is introduced in Data ONTAP 8.2.

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

SnapVault was rebuilt from the ground up for its debut in clustered Data ONTAP, and although former 7-Mode users will find similarities, major enhancements have been made in this version of SnapVault. One major advance is the ability to preserve storage efficiencies on primary data during SnapVault transfers. This technical report covers the details of storage efficiency preservation as well as other topics that have changed from the 7-Mode version of SnapVault.

One important architectural change is that SnapVault in clustered Data ONTAP replicates at the volume level as opposed to the qtree level, as in 7-Mode SnapVault. This means that the source of a SnapVault relationship must be a volume, and that volume must replicate to its own volume on the SnapVault secondary. The implications of this fact are discussed in section 5, "Fan-In and Fan-Out SnapVault Topologies."

## 2 Clustered Data ONTAP Nondisruptive Operations and SnapVault

SnapVault users can now take advantage of nondisruptive operations, the cornerstone of clustered Data ONTAP. SnapVault administrators can seamlessly rebalance SnapVault primaries and secondaries for performance or capacity needs, because it is now possible to move SnapVault primary and secondary volumes 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 vol move, the transfer may pause for a few minutes during the volume cutover phase, but the transfer resumes from the most recent transfer checkpoint after the vol move operation completes. Administrators never have to reconfigure a SnapVault relationship just because a volume was moved to another node by using the vol move operation.

## 3 Storage Virtual Machines and Cluster Peering

Clustered Data ONTAP allows a SnapVault secondary volume to be on the same NetApp Storage Virtual Machine (SVM; formerly known as Vserver) as the primary volume, or on a different SVM than the primary volume. These SVMs can also be in the same cluster or in different clusters. If the primary and secondary volumes are in different SVMs, then in order for SnapVault updates to be successful, the primary and secondary SVMs must be able to communicate with each other.

This requires that the primary and secondary Storage Virtual Machines be peered; and if the SVMs are in different clusters, then the clusters must be peered as well. Before configuring SnapVault relationships, be sure that SVM and cluster peering are configured as required. For information about creating these peer relationships, refer to the "System Administration Guide for Cluster Administrators" and express guides, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

## 4 SnapVault and Storage Efficiencies

A new feature for the clustered Data ONTAP version of SnapVault is the ability to preserve the storage efficiencies of the primary volume throughout a SnapVault transfer. This means that 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. It also means that the data is already deduplicated and compressed on the secondary volume after the SnapVault transfer completes, without having to run the deduplication and compression processes on the secondary volume.

For the SnapVault transfer to be successful, it needs at least the same amount of free space available on the secondary volume, as if it had 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. 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, independent of the storage efficiencies present 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 to not be preserved on the SnapVault secondary volume. For more information about this caveat and other best practices associated with using storage efficiencies with SnapVault, refer to section 14.1 of [TR-3966](#) "NetApp Data Compression and Deduplication Deployment and Implementation Guide." TR-3966 also helps you understand how much disk space you can save by using storage efficiencies with SnapVault.

## 5 Fan-In and Fan-Out SnapVault Topologies

SnapVault in clustered Data ONTAP does support fan-in and fan-out topologies. However, there are some differences from 7-Mode. More information about these topologies is available starting on page 42 of the clustered Data ONTAP Data Protection Guide, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

### 5.1 Fan-In Topologies

A fan-in topology is one where multiple SnapVault primaries back up to one SnapVault secondary. The primary use case for this topology is multiple remote sites that back up to one central data center. In 7-Mode, fan-in meant that multiple qtrees can back up to the same secondary volume. This is referred to as *volume-level fan-in*. In clustered Data ONTAP, because the relationships are configured at the volume level, it is not possible to back up multiple volumes to one secondary volume. However, it is possible to back up SnapVault primary volumes from multiple Storage Virtual Machines and/or clusters to separate volumes in a single destination SVM, which can be in a different cluster. This is referred to as *system-level fan-in*.

One limit to keep in mind is that in order to have SnapVault primaries and secondaries in different clusters, those clusters must be peered; in 8.2 a maximum of 8 clusters can be peered together. This leads to a cluster-level fan-in limit of 7 to 1. Clustered Data ONTAP 8.2.1 supports higher fan-in ratios up to 64 clusters.

### 5.2 Fan-Out Topologies

A fan-out topology is one where a single primary volume is replicated to multiple destinations. This topology allows a single primary volume to be protected and backed up, as well as providing a read-only copy at a secondary site. This is supported in clustered Data ONTAP with a limit of 4 replication destinations for a single source. This limit of 4 destinations is shared between SnapMirror<sup>®</sup> and SnapVault relationships. There can be any combination of SnapMirror and SnapVault relationships for a single

source, as long as the total number of relationships is not greater than 4. Clustered Data ONTAP 8.2.1 supports higher fan-out ratios up to 8.

## 6 Cascading SnapVault and SnapMirror Relationships

Cascading SnapMirror and SnapVault relationships are supported in clustered Data ONTAP 8.2. A SnapMirror secondary can be the source of a SnapVault relationship (backing up a DR mirror); or a SnapVault secondary can be the source of a SnapMirror relationship (protecting a backup). In the case of a SnapMirror to SnapVault relationship, it is not possible to specify which Snapshot™ copies are transferred to the SnapVault secondary; SnapVault always transfers the SnapMirror exported Snapshot copy, or the base copy of the SnapMirror relationship. This is similar to the “snapmirror base snapshot only” option in 7-Mode. The administrator does have the ability to choose when the most recent version of the SnapMirror exported Snapshot copy is transferred to the SnapVault secondary, and how many copies to keep on the SnapVault secondary. More information about cascading is available starting on page 41 of the Data Protection Guide, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

## 7 Scheduling Snapshots and SnapVault Transfers

Setting up SnapVault in clustered Data ONTAP requires several steps. On the primary volume, you must configure a Snapshot policy that specifies what copies to create and when to create them by using cron schedules, and assign labels to specify which copies should be transferred by SnapVault. On the secondary, a SnapMirror policy must be created that specifies the labels of Snapshot copies to be transferred and how many of these copies should be kept on the secondary. After configuring these policies, you can create the SnapVault relationship and establish a transfer schedule. A flow chart describing these steps and detailed setup instructions are available in the SnapVault Express Guide and the Data Protection Guide, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

## 8 Interoperability with NDMP and SMTape

Clustered Data ONTAP 8.2 brings new functionality to allow you to integrate tape into your clustered Data ONTAP backup strategy.

### 8.1 SMTape Seeding

Starting with clustered Data ONTAP 8.2, limited SMTape functionality was introduced to allow SnapMirror and SnapVault relationships to be seeded by copying the initial baseline to tape and restoring it to the secondary. Seeding relationships is the only supported use of SMTape in clustered Data ONTAP 8.2. Detailed information about how to use tape seeding can be found starting on page 64 of the Clustered Data ONTAP 8.2 Data Protection Guide, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

### 8.2 NDMP

Cluster-aware NDMP became available starting in clustered Data ONTAP 8.2. Prior to 8.2, NDMP commands were node scoped and could be executed only by a node or cluster administrator. Also, because the commands were node scoped, if the primary volume of a scheduled NDMP dump backup was moved to a different controller by using vol move, the administrator would have to reschedule the backup on the new node on which the primary volume resides. In 8.2, Storage Virtual Machine scoped commands were introduced. This means that an SVM administrator can execute NDMP commands for any of the volumes in the SVM, regardless of which node they reside on. Detailed information about how to use SVM-aware NDMP can be found starting on page 47 of the Data Protection Tape Backup and Recovery Guide, available on the [NetApp Support site for clustered Data ONTAP 8.2](#).

## 9 Management with System Manager

With the release of OnCommand® System Manager 3.0, System Manager is able to create and manage SnapVault relationships in clustered Data ONTAP. System Manager can carry out the following tasks: create intercluster LIFs on the primary and destination clusters, perform cluster peering, perform Storage Virtual Machine peering, create and apply Snapshot policies, create and modify cron schedules, create and apply SnapMirror policies, create secondary volumes, and create SnapVault relationships. More information about System Manager capabilities can be found on the [NetApp Support site for OnCommand System Manager 3.0](#).

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®