




NetApp®

Clustered Data ONTAP® 8.3

Volume Disaster Recovery Preparation Express Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09376_A0
November 2014

Contents

Deciding whether to use this guide	4
Volume disaster recovery preparation workflow	6
Verifying the cluster peer relationship	7
Creating the SnapMirror relationship	7
Setting up the destination SVM for data access	12
Monitoring the status of SnapMirror data transfers	13
Where to find additional information	15
Copyright information	16
Trademark information	17
How to send your comments	18
Index	19

Deciding whether to use this guide

This guide describes how to quickly protect a source volume on a peered Data ONTAP cluster in preparation for disaster recovery. You should use this guide if you want to configure and monitor SnapMirror relationships between peered clusters for volume disaster recovery and do not need a lot of conceptual background for the tasks.

SnapMirror provides scheduled asynchronous, block-level data protection. SnapMirror replicates Snapshot copies and can replicate NAS or SAN volumes on which deduplication, data compression, or both are run, including volumes containing qtrees and LUNs. SnapMirror configuration information is stored in a database that Data ONTAP replicates to all the nodes in the cluster.

You should use this guide if you want to create SnapMirror relationships for volume-level disaster recovery in the following way:

- You are working with clusters running Data ONTAP 8.3 or later.
- You are a cluster administrator.
- You are using FlexVol volumes and not an Infinite Volume.
- You have configured an authenticated peer relationship between two clusters.
[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)
- You have enabled the SnapMirror license on both the source and the destination clusters.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use OnCommand System Manager, not the command-line interface or an automated scripting tool.

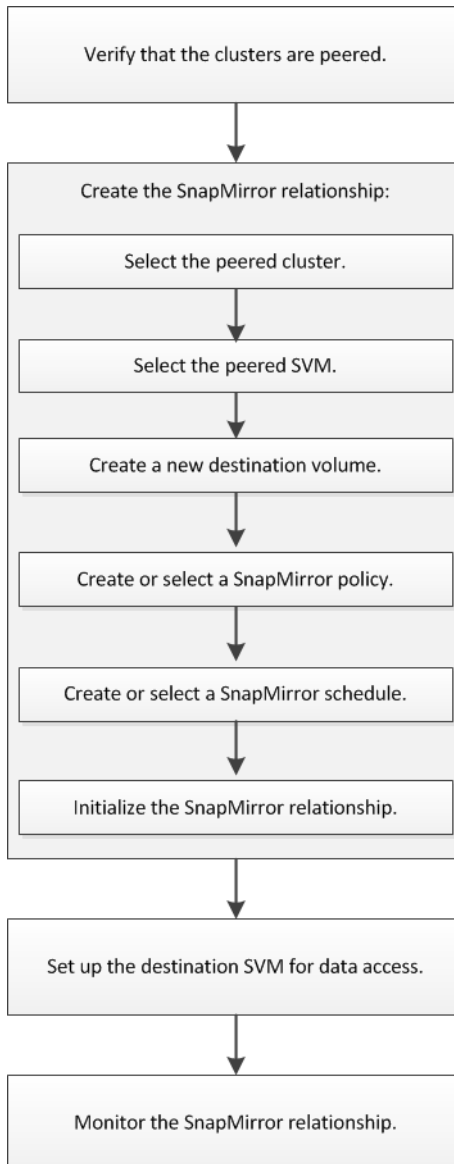
If these assumptions are not correct for your situation, or if you want more conceptual background information, you should see the following resources:

- [Clustered Data ONTAP 8.3 Data Protection Guide](#)
- [NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#)
- [Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)
- [NetApp Documentation: OnCommand Workflow Automation \(current releases\)](#)
OnCommand Workflow Automation enables you to run prepackaged workflows that automate management tasks such as the workflows described in Express Guides.
- [Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

Provides the set of commands for configuring and managing SnapMirror relationships. You might want to use the SnapMirror commands to write a script that creates multiple SnapMirror relationships.

Volume disaster recovery preparation workflow

Preparing volumes for disaster recovery involves verifying the cluster peer relationship, creating the SnapMirror relationship between volumes residing on peered clusters, setting up the destination SVM for data access, and monitoring the SnapMirror relationship periodically.



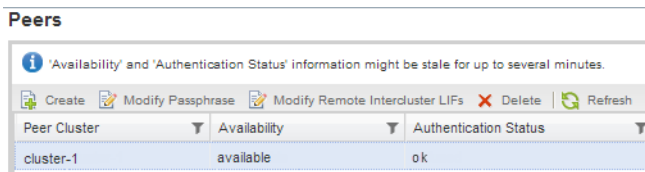
Verifying the cluster peer relationship

Before you set up a volume for disaster recovery, you must verify that the source and destination clusters are peered and are communicating with each other through the peer relationship.

Steps

1. Enter the URL `https://IP-address-of-cluster-management-LIF` in a web browser and log in to System Manager using your cluster administrator credential.
2. Expand the **Cluster** hierarchy in the left navigation pane.
3. Click **Configuration > Peers** to ensure that the peered cluster is authenticated and available.

The destination cluster is cluster-1 and the source cluster is cluster-2.



Peer Cluster	Availability	Authentication Status
cluster-1	available	ok

Creating the SnapMirror relationship

You must create a SnapMirror relationship between the source volume on one cluster and the destination volume on the peered cluster for replicating data for disaster recovery.

Before you begin

- You must have the cluster administrator user name and password for the destination cluster.
- The destination aggregate must have available space.
- Both the clusters must be configured and set up appropriately to meet the requirements of your environment for user access, authentication, and client access.

About this task

You must perform this task from the **source** cluster.

Steps


1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. Select the source SVM that contains the volume you want to protect, and then select **Storage > Volumes**.
3. Select the volume you want to protect from the Volumes list, and then click **Protect by > Mirror**.

The Create Mirror Relationship window is displayed.

4. In the **Destination Volume** section, select the peered cluster.
5. Specify the SVM for the destination volume:

If the SVM is...	Then...
Peered	Select the peered SVM from the list.
Not peered	<ol style="list-style-type: none"> a. Select the SVM. b. Click Authenticate. c. Enter the cluster administrator's credentials of the peered cluster, and then click Create.

6. Create a new destination volume:
 - a. Select the **New Volume** option.
 - b. Use the default volume name or enter a new volume name.
 - c. Select the destination aggregate.



7. In the **Configuration Details** section, specify the SnapMirror policy:

If you want to...	Then...
Assign an existing policy	Select a SnapMirror policy from the list.

If you want to...	Then...
Create a new policy	<ol style="list-style-type: none"> a. Click Create Policy. b. In the Create Mirror Policy window, specify a policy name and set the schedule transfer priority. Normal priority transfers are scheduled before low-priority transfers. By default, the priority is set to Normal. c. Select the Transfer All Source Snapshot Copies check box to include the “all_source_snapshots” rule to the mirror policy, which will enable you to back up all the Snapshot copies from the source volume. d. Select the Enable Network Compression check box to compress the data that is being transferred. e. Click Create.

The screenshot shows a 'Create Mirror Policy' dialog box with the following details:

- Destination Cluster: cluster-1
- Destination Storage Virtual Machine: dest_SVM
- Policy Name: policy
- Transfer Priority: Normal
- Transfer All Source Snapshot Copies
- Enable Network Compression
- [Add Comments](#)
- Buttons: Create, Cancel

8. Specify the SnapMirror schedule:

If you want to...	Then...
Assign an existing schedule	From the list of schedules, select an existing schedule.
Create a new schedule	<ol style="list-style-type: none"> a. Click Create Schedule. b. In the Create New Schedule window, choose one of the following options: <ul style="list-style-type: none"> • Basic to specify recurring days and recurring schedule details • Advanced to specify the advanced cron options in months, days, week days (if applicable), hours, and minutes c. Click Create.

Create New Schedule

Destination Cluster: cluster-1

Schedule Name: custom_schedule1

Basic ? Advanced (cron style)

Repeat: Daily

Time: 12:00AM

Create Cancel

9. Ensure that the **Initialize Relationship** check box is selected and then click **Create**.

Initializing the SnapMirror relationship ensures that the destination volume has a baseline to start protecting the source volume.

Configuration Details

? Mirror Policy: DPDefault [Create Policy](#)
SnapMirror labels: sm_created

Create version flexible mirror relationship. ?

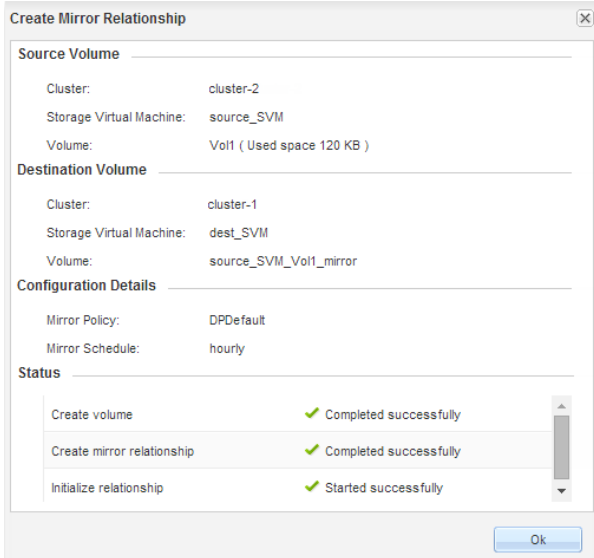
? Mirror Schedule: hourly [Create Schedule](#)
Every hour at 05 minute(s)
 None

Initialize Relationship

Create Cancel

The wizard creates the relationship with the default **DPDefault** mirror policy and the **hourly** schedule. The relationship is initialized by starting a baseline transfer of data from the source volume to the destination volume.

The initialization operation might take some time. The Status section shows the status of each job.



10. Verify the relationship status of the SnapMirror relationship:

- a. Select the volume from the Volumes list, and then click **Data Protection**.
- b. In the **Data Protection** tab, verify that the SnapMirror relationship you created is listed and the relationship state is **snapmirrored**.

Volumes

Name	Aggregate	Status	% Used	Available Space	Total Space	Storage Efficiency
Vol1	aggr3	Online	5	18.89 MB	20 MB	Disabled
Vol11	aggr3	Online	6	18.67 MB	20 MB	Disabled
Vol2	aggr3	Online	5	18.85 MB	20 MB	Disabled
dest_SVM_source_S...	aggr3	Online	5	18.87 MB	20 MB	Disabled
dest_SVM_vol_040_...	aggr1	Online	5	42.62 MB	45 MB	Disabled

Destination Storage Virtual Machine	Destination Volume	Is Healthy	Relationship State	Transfer Status	Type	Lag Time	Policy
dest_SVM	source_SVM_Vol1_mirror	Yes	Snapmirrored	Idle	Mirror	None	DPDefault

Details | Space Allocation | Snapshot Copies | Storage Efficiency | **Data Protection**

After you finish

You must make a note of the source volume settings such as thin provisioning, deduplication, compression, and autogrow. You can use this information when you break the SnapMirror relationship and verify the destination volume settings.

Setting up the destination SVM for data access

You can minimize data access disruption when activating the destination volume by setting up required configurations such as LIFs, CIFS shares, and export policies for the NAS environment, and LIFs and initiator groups for the SAN environment on the SVM containing the destination volume.

About this task

You must perform this task on the **destination** cluster for the SVM containing the destination volume.

Choices

- NAS environment:
 1. Create NAS LIFs.
 2. Create CIFS shares with the same share names that were used on the source.
 3. Create appropriate NFS export policies.
 4. Create appropriate quota rules.
- SAN environment:
 1. Create SAN LIFs.
 2. Optional: Configure portsets.
 3. Configure initiator groups.
 4. For FC, zone the FC switches to enable the SAN clients to access the LIFs.

After you finish

If any changes were made on the SVM containing the source volume, you must replicate the changes manually on the SVM containing the destination volume.

Related information

[NetApp Documentation: Clustered Data ONTAP Express Guides](#)

Monitoring the status of SnapMirror data transfers

You should periodically monitor the status of the SnapMirror relationships to ensure that the SnapMirror data transfers are occurring as per the specified schedule.

About this task

You must perform this task from the **destination** cluster.

Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. Select the Storage Virtual Machine (SVM) that contains the destination volume, and then click **Protection**.
3. Select the SnapMirror relationship between the source and the destination volumes, and then verify the status in the **Details** bottom tab.

The Details tab displays the health status of the SnapMirror relationship and shows the transfer errors and lag time.

- The Is Healthy field must display **Yes**.
For most SnapMirror data transfer failures, the field displays **No**. In some failure cases, however, the field continues to display **Yes**. You must check the transfer errors in the Details section to ensure that no data transfer failure occurred.
- The Relationship State field must display **snapmirrored**.
- The Lag Time must be no more than the transfer schedule interval.
For example, if the transfer schedule is hourly, then the lag time must not be more than an hour.

You should troubleshoot any issues in the SnapMirror relationships.

[NetApp Technical Report 4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#)

14 | Volume Disaster Recovery Preparation Express Guide

Source Location:	source_SVM:Vol1	Is Healthy:	● Yes	Transfer Status:	Idle
Destination Location:	dest_SVM:source_SVM_Vo...	Relationship State:	Snapmirrored	Current Transfer Type:	None
Source Cluster:	cluster-2	Network Compression Ratio:	Not Applicable	Current Transfer Error:	None
Destination Cluster:	cluster-1			Last Transfer Error:	None
Transfer Schedule:	hourly			Last Transfer Type:	Initialize
Data Transfer Rate:	Unlimited			Latest Snapshot Timestamp:	09/16/2014 23:42:24
Lag Time:	None			Latest Snapshot Copy:	snapmirror.3e61ed5f-31a3-11e4-98c7-005056974d2d_2147484686.2014-09-16_233529

Where to find additional information

Additional documentation is available to help you activate the destination volume to test the disaster recovery setup or when a disaster occurs. You can also learn more about how to reactivate the source volume after the disaster.

Express guide

- [*Clustered Data ONTAP 8.3 Volume Disaster Recovery Express Guide*](#)
Describes how to quickly activate a destination volume after a disaster and then reactivate the source volume in clustered Data ONTAP 8.3.

Comprehensive guides

- [*Clustered Data ONTAP 8.3 Data Protection Guide*](#)
Provides detailed conceptual and CLI-based task information about disaster recovery by using SnapMirror technology.
- [*OnCommand Unified Manager 6.1 Administration Guide*](#)
Provides information about performing OnCommand Unified Manager tasks using the web UI and information about troubleshooting, as well as providing in-depth conceptual information.

Copyright information

Copyright © 1994–2014 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, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- C**
- cluster peer relationships
 - verifying the status [7](#)
 - clusters
 - requirements for using Express Guide to configure volume disaster recovery [4](#)
 - comments
 - how to send feedback about documentation [18](#)
- D**
- data access
 - setting up the destination SVM [12](#)
 - data protection
 - where to get additional information about [15](#)
 - destination volumes
 - creating for disaster recovery [7](#)
 - setting up for data access from [12](#)
 - disaster recovery
 - creating SnapMirror relationships between volumes on different clusters for [7](#)
 - requirements for using Express Guide to configure mirror relationships for [4](#)
 - where to get additional information about volumes [15](#)
 - workflow for volume [6](#)
 - documentation
 - how to send feedback about [18](#)
- E**
- express guides
 - requirements for using to configure SnapMirror relationships for disaster recovery [4](#)
- F**
- feedback
 - how to send comments about documentation [18](#)
 - flowcharts
 - volume disaster recovery preparation workflow [6](#)
- I**
- information
 - how to send feedback about improving documentation [18](#)
 - intercluster
 - creating SnapMirror relationships between volumes for disaster recovery [7](#)
- L**
- lag times
 - monitoring SnapMirror relationship health [13](#)
- M**
- mirror relationships
 - creating between volumes on different clusters for disaster recovery [7](#)
 - initializing [7](#)
 - monitoring health of [13](#)
 - requirements for using Express Guide to configure, for disaster recovery [4](#)
- P**
- peer relationships
 - verifying the status of cluster [7](#)
 - policies
 - SnapMirror, creating for disaster recovery [7](#)
 - protection
 - where to get additional information about data [15](#)
- R**
- relationships
 - creating SnapMirror, between volumes on different clusters for disaster recovery [7](#)
 - monitoring health of SnapMirror [13](#)
 - requirements for using Express Guide to configure SnapMirror, for disaster recovery [4](#)
 - verifying cluster peer [7](#)

S

- schedules
 - SnapMirror, creating for disaster recovery [7](#)
- setting up
 - destination volume [12](#)
- SnapMirror
 - workflow for volume disaster recovery preparation [6](#)
- SnapMirror policies
 - creating for disaster recovery [7](#)
- SnapMirror relationships
 - See* mirror relationships
- SnapMirror schedules
 - creating for disaster recovery [7](#)
- suggestions
 - how to send feedback about documentation [18](#)

V

- verifying

- cluster peer relationships [7](#)
- volume disaster recovery preparation
 - where to get additional information about [15](#)
- volumes
 - requirements for using Express Guide to configure disaster recovery between [4](#)
 - setting up data access from destination [12](#)
 - SnapMirror for disaster recovery preparation workflow [6](#)

W

- workflows
 - volume disaster recovery preparation [6](#)