

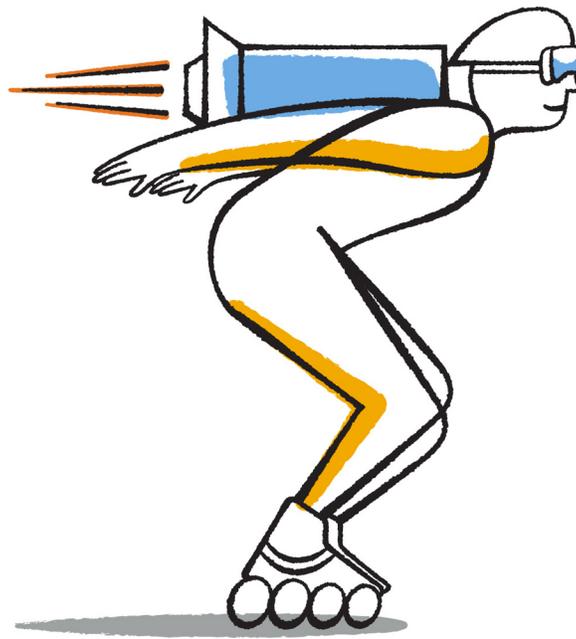


NetApp[®]

New for 8.3.1

Clustered Data ONTAP[®] 8.3

SVM Disaster Recovery 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: 215-10150_B0
August 2015

Contents

| | |
|--|-----------|
| Deciding whether to use this guide | 4 |
| SVM disaster recovery workflow | 5 |
| Quiescing SnapMirror transfers | 5 |
| Aborting any ongoing SnapMirror transfers | 6 |
| Breaking the SVM disaster recovery relationship | 6 |
| Stopping the source SVM | 7 |
| Starting the destination SVM | 7 |
| Source SVM reactivation workflow | 9 |
| Reactivating the source SVM | 9 |
| Creating the new source SVM | 10 |
| Setting up the existing source SVM | 12 |
| Stopping the destination SVM | 14 |
| Updating the SnapMirror relationship | 15 |
| Breaking the SVM disaster recovery relationship | 15 |
| Starting the source SVM | 16 |
| Resynchronizing the destination SVM from the source SVM | 17 |
| Where to find additional information | 19 |
| Copyright information | 20 |
| Trademark information | 21 |
| How to send comments about documentation and receive update notifications | 22 |
| Index | 23 |

Deciding whether to use this guide

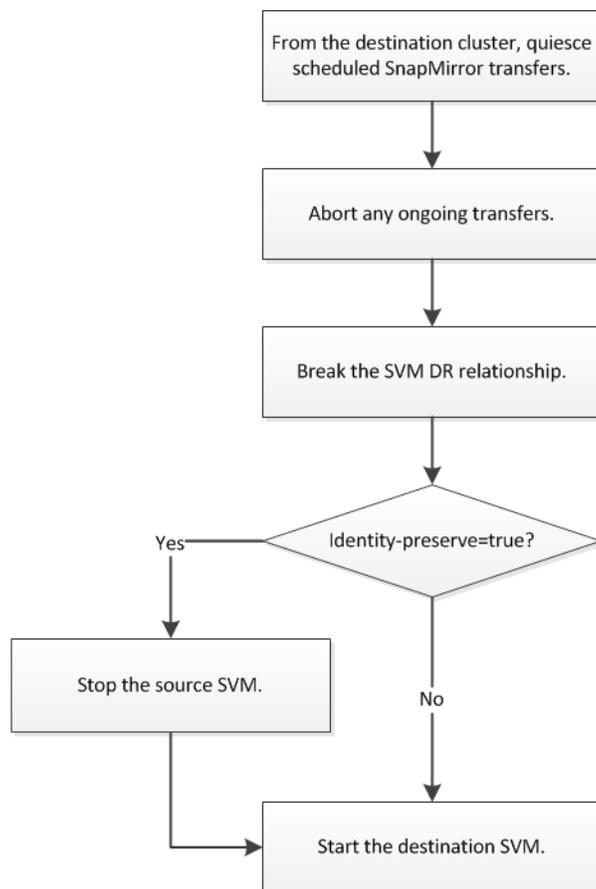
This guide describes how cluster administrators quickly activate a destination Storage Virtual Machine (SVM) after a disaster and then reactivate the source SVM. You can use the procedures listed in this guide to test the SVM disaster recovery solution.

You should use this guide if you want to activate the destination SVM and then reactivate the source SVM in the following way:

- You are a cluster administrator.
- You are working with SVMs with FlexVol volumes on clusters running Data ONTAP 8.3.1 or later.
- You have prepared the source SVM for disaster recovery by configuring the destination SVM by using the *SVM Disaster Recovery Preparation Express Guide*.
[Clustered Data ONTAP 8.3 SVM Disaster Recovery Preparation Express Guide](#)
- You are using the Data ONTAP command-line interface.
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.

SVM disaster recovery workflow

To recover from a disaster, you must activate the destination SVM. Activating the destination SVM involves quiescing the scheduled SnapMirror transfers and any ongoing transfers, breaking the SVM disaster recovery relationship, stopping the source SVM, starting the destination SVM, and verifying the status of the destination SVM.



Quiescing SnapMirror transfers

Before activating the destination Storage Virtual Machine (SVM), you must quiesce the SVM disaster recovery relationship to stop scheduled SnapMirror transfers from the source SVM.

About this task

You must perform this task from the destination cluster.

Steps

1. Stop the scheduled SnapMirror transfers by using the `snapmirror quiesce` command.

Example

```
destination_cluster::> snapmirror quiesce -destination-path dvs1:
```

2. Verify that the SnapMirror relationship between the source and the destination SVMs is in the Quiescing or Quiesced state by using the `snapmirror show` command.

For viewing the detailed status of the relationship, you can use the `-instance` option.

Example

```
destination_cluster::> snapmirror show
```

| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Healthy | Progress Last Updated |
|-------------|------|------------------|--------------|---------------------|----------------|---------|-----------------------|
| vsl: | DP | dvs1: | Snapmirrored | Quiesced | - | true | - |

Aborting any ongoing SnapMirror transfers

You must abort any ongoing SnapMirror transfers or any long-running quiesce operations before breaking the SVM disaster recovery relationship.

About this task

You must perform this task from the destination cluster.

Steps

1. Abort any ongoing SnapMirror transfers by using the `snapmirror abort` command.

Example

```
destination_cluster::> snapmirror abort -destination-path dvs1:
```

2. Verify that the SnapMirror relationship between the source and destination SVMs is in the `idle` state by using the `snapmirror show` command.

For viewing the detailed status of the relationship, you can use the `-instance` option.

Example

```
destination_cluster::> snapmirror show
```

| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Healthy | Progress Last Updated |
|-------------|------|------------------|--------------|---------------------|----------------|---------|-----------------------|
| vsl: | DP | dvs1: | Snapmirrored | Idle | - | false | - |

Breaking the SVM disaster recovery relationship

You must break the SnapMirror relationship created between the source and the destination SVMs for disaster recovery before activating the destination SVM.

About this task

You must perform this task from the destination cluster.

Steps

1. Break the SVM disaster recovery relationship by using the `snapmirror break` command.

Example

```
destination_cluster::> snapmirror break -destination-path dvs1:
```

2. Verify that the SnapMirror relationship between the source and destination SVMs is in the Broken-off state by using the `snapmirror show` command.

Example

```
destination_cluster::> snapmirror show
```

| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Progress Last Updated | Healthy |
|-------------|------|------------------|--------------|---------------------|----------------|-----------------------|---------|
| vsl: | DP | dvs1: | Broken-off | Idle | - | - | true |

The subtype changes from `dp-destination` to `default`. The type of the volumes in the destination SVM changes from `DP` to `RW`.

Stopping the source SVM

If you chose to set `identity-preserve` to **true** SVM configuration or if you want to test the SVM disaster recovery setup, you must stop the source SVM before activating the destination SVM.

Before you begin

If the source SVM is available on the source cluster, then you must have ensured that all clients connected to the source SVM are disconnected.

About this task

You must perform this task from the source cluster.

Steps

1. Stop the source SVM by using the `vserver stop` command.

Example

```
source_cluster::> vserver stop -vserver vs1
```

2. Verify that the source SVM is in the stopped state by using the `vserver show` command.

Example

```
source_cluster::> vserver show
```

| Vserver | Type | Subtype | Admin State | Operational State | Root Volume | Aggregate |
|---------|------|---------|-------------|-------------------|-------------|-----------|
| vs1 | data | default | stopped | stopped | rv | aggr1 |

Starting the destination SVM

In case of a disaster or while testing the SVM disaster recovery setup, you must activate the destination SVM to provide data access from the destination SVM.

Before you begin

The source SVM must be in the **stopped** state.

About this task

You must perform this task from the destination cluster.

Steps

1. If you chose to set the `-identity-preserve` option to **true** and the source and destination SVMs are in different subnets, perform the following steps:
 - a. To serve data, set the `-status-admin` option of the replicated LIFs that belong to the subnet of the destination SVM to **up** by using the `network interface modify` command.

Example

```
destination_cluster::> network interface modify -vserver dvs1
-lif LIF1a -status-admin up
```

- b. Set the `-status-admin` option of the replicated LIFs that belong to subnet of the source SVM to **down** by using the `network interface modify` command.

Example

```
destination_cluster::> network interface modify -vserver dvs1
-lif LIF1 -status-admin down
```

2. Start the destination SVM by using the `vserver start` command.

Example

```
destination_cluster::> vserver start -vserver dvs1
[Job 30] Job succeeded: DONE
```

3. Verify that the destination SVM is in the **running** state and the subtype is **default** by using the `vserver show` command.

Example

```
destination_cluster::> vserver show
```

| Vserver | Type | Subtype | Admin State | Operational State | Root Volume | Aggregate |
|---------|------|---------|-------------|-------------------|-------------|-----------|
| dvs1 | data | default | running | running | voll | aggr1 |

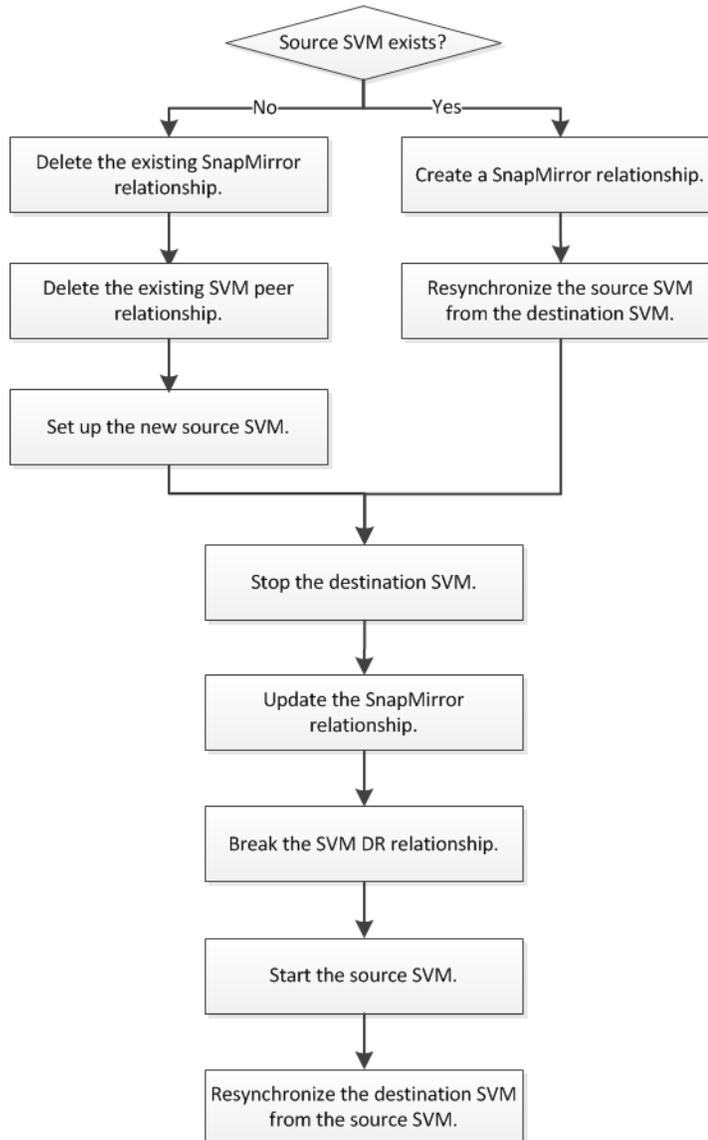
4. If you chose to set the `-identity-option` option to **true**, verify that the LIFs are correctly placed on the network of destination SVM by pinging from the LIF to its gateway or a host in the same subnet.

After you finish

The clients might have to remount or rescan per the protocol requirements for accessing the data from the destination SVM.

Source SVM reactivation workflow

If the source SVM exists after a disaster, you can reactivate it and protect it by re-creating the SVM disaster recovery relationship between the source and the destination SVMs. If the source SVM does not exist, you must create and set up a new source SVM and then reactivate it.



Reactivating the source SVM

Depending on whether the source SVM exists after a disaster, you can either use the existing source SVM or create a new source SVM for reactivation.

Choices

- [Creating the new source SVM](#) on page 10

- [Setting up the existing source SVM](#) on page 12

Creating the new source SVM

If the source SVM does not exist, you must delete the SnapMirror relationship between the source and destination SVMs, delete the SVM peer relationship, and create and set up a new source SVM to replicate the data and configuration from the destination SVM.

Steps

1. [Deleting the SnapMirror relationship](#) on page 10
2. [Deleting the SVM peer relationship](#) on page 10
3. [Setting up the new source SVM](#) on page 11

Deleting the SnapMirror relationship

If the source SVM no longer exists, you must delete the SnapMirror relationship between the source and the destination SVMs before setting up a new source SVM.

Steps

1. From the destination cluster, identify the SnapMirror relationship between the source SVM that no longer exists and its destination SVM by using the `snapmirror show` command.

Example

```
destination_cluster::> snapmirror show
```

| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Healthy | Progress Last Updated |
|-------------|------|------------------|--------------|---------------------|----------------|---------|-----------------------|
| vs1: | DP | dvs1: | Broken-off | Idle | - | true | - |

2. Delete the SnapMirror relationship by using the `snapmirror delete` command.

Example

```
destination_cluster::> snapmirror delete -destination-path dvs1:
```

3. Verify that the SnapMirror relationship is deleted by using the `snapmirror show` command. The deleted SnapMirror relationship entry is no longer displayed in the output.

Deleting the SVM peer relationship

If the source SVM no longer exists, you must delete the SVM peer relationship between that source SVM and its destination SVM before you create and configure a new source SVM.

Steps

1. From the destination cluster, identify the SVM peer relationship between the source SVM that no longer exists and its destination SVM by using the `vserver peer show` command.

Example

```
destination_cluster::>vserver peer show
```

| Vserver | Peer Vserver | Peer State | Peering Applications |
|---------|--------------|------------|----------------------|
| dvs1 | vs1 | peered | snapmirror |

2. Delete the SVM peer relationship by using the `vserver peer delete` command.

Example

```
destination_cluster::> vserver peer delete -vserver dvsl -peer-vserver vs1
Info: [Job 47] 'vserver peer delete' job queued
```

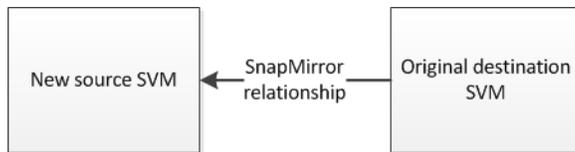
Setting up the new source SVM

You must set up the new source SVM by creating the SVM disaster recovery relationship to replicate the data and configuration from the destination SVM.

About this task

You must set up the disaster recovery relationship by using the same method and configuration you used to set up the SnapMirror relationship before the disaster. For example, if you chose to replicate data and all the configuration details when creating the SnapMirror relationship between the original source and destination SVMs, you must choose to replicate data and all configuration details when creating the SnapMirror relationship between the new source SVM and the destination SVM.

You must perform this task from the new source cluster.



You can follow the detailed steps provided in the *SVM Disaster Recovery Preparation Express Guide* to set up the new source SVM.

[Clustered Data ONTAP 8.3 SVM Disaster Recovery Preparation Express Guide](#)

Steps

1. Prepare the new source cluster:
 - a. Verify that the cluster peer relationship is healthy.
 - b. Install all the required feature licenses and protocols.
 - c. Create the required custom schedules.
 - d. Ensure that a non-root aggregate with minimum free space of 10 GB exists.
2. Create the new source SVM of subtype `dp-destination`, which is now the destination SVM.
3. Create the SVM peer relationship between the new source SVM and original destination SVM.
4. Create a SnapMirror relationship between the new source SVM and the original destination SVM.
5. For CIFS: If you chose to replicate data and a subset of SVM configuration, create a CIFS server.
6. For different subnet: If you chose to set the `-identity-preserve` option to `true`, create LIFs on the original destination SVM for the new source SVM in the subnet of the new source SVM.
7. Initialize the new source SVM.
8. If you chose to replicate data and a subset of SVM configuration, configure network and protocol access on the new source SVM for data access.

Setting up the existing source SVM

If the source SVM exists after a disaster, you must create the SVM disaster recovery relationship between the destination and source SVMs and resynchronize the data and configuration from the destination SVM to the source SVM.

Steps

1. [Creating a SnapMirror relationship](#) on page 12
2. [Resynchronizing the source SVM from the destination SVM](#) on page 13

Creating a SnapMirror relationship

You must create a SnapMirror relationship between the existing source and the destination SVMs for replicating the data and configuration details from the destination SVM.

Before you begin

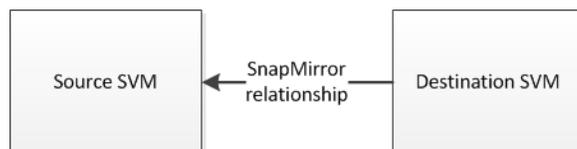
- The destination cluster must have at least one non-root aggregate with a minimum free space of 10 GB for configuration replication.
The best practice is to have at least two non-root aggregates with a minimum free space of 10 GB.
- The source and destination clusters must be peered.
- Any custom schedules that are being used by the destination SVM must be created on the source SVM.
- The existing source SVM and the destination SVM must be peered.

About this task

If you are not reactivating the source SVM for the first time, then the SnapMirror relationship between the source and destination SVMs exists in `Broken-off` state. In such cases, you do not have to perform this task.

You must set up the disaster recovery relationship by using the same method and configuration you used before the disaster. For example, if you chose to replicate data and all the configuration details when creating the SnapMirror relationship between the original source and destination SVMs, you must choose to replicate data and all configuration details when creating the SnapMirror relationship between the new source and destination SVMs.

You must perform this task from the new source cluster.



Steps

1. Create a SnapMirror relationship between the source and destination SVMs by using the `snapmirror create` command:

You can either specify the source and the destination SVMs as paths or SVM names. If you want to specify the source and destination as paths, then the SVM name must be followed by a colon.

- Replicate the data and all of the configuration information by setting the `-identity-preserve` option to `true`.

The following command creates the SnapMirror relationship with SVM names as the `-destination-path` and `-source-path` parameters:

```
source_cluster::> snapmirror create -source-path dvs1: -destination-path vs1: -type DP
-throttle unlimited -policy DPDefault -schedule hourly -identity-preserve true
```

The following command creates the SVM disaster recovery relationship with SVM names as the `-destination-vserver` and `-source-vserver` parameters:

```
destination_cluster::> snapmirror create -source-vserver dvs1 -destination-vserver vs1
-type DP -throttle unlimited -policy DPDefault -schedule hourly -identity-preserve true
```

- Replicate the data and a subset of the configuration information by setting the `-identity-preserve` option set to **false**.

The following command creates the SnapMirror relationship with SVM names as the `-destination-vserver` and `-source-vserver` parameters:

```
source_cluster::> snapmirror create -source-path dvs2: -destination-path vs2: -type DP
-throttle unlimited -policy DPDefault -schedule hourly -identity-preserve false
```

The following command creates the SnapMirror relationship with SVM names as the `-destination-vserver` and `-source-vserver` parameters:

```
source_cluster::> snapmirror create -source-vserver dvs2 -destination-vserver vs2 -
type DP -throttle unlimited -policy DPDefault -schedule hourly -identity-preserve false
```

2. Verify that the SnapMirror relationship is established and is in the Broken-off state by using the `snapmirror show` command.

Example

```
destination_cluster::> snapmirror show
```

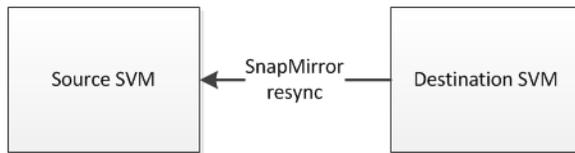
| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Progress Last Updated | Healthy |
|-------------|------|------------------|--------------|---------------------|----------------|-----------------------|---------|
| dvs1: | DP | vs1: | Broken-off | Idle | - | - | true |

Resynchronizing the source SVM from the destination SVM

Before activating the source SVM, you must resynchronize the data and configuration details from the destination SVM to the existing source SVM for data access.

Before you begin

- The SVM root volume must not contain any other data apart from metadata because the other data is not replicated.
Root volume metadata such as volume junctions, symbolic links, directories leading to junctions symbolic links are replicated.
- The source SVM must not contain load-sharing mirrors apart from the load-sharing mirror created for the SVM root volume protection.
- The source SVM must not contain any new volumes that are not replicated on the destination SVM.
You must delete such volumes on the source SVM to prevent resynchronization failure.

About this task**Steps**

1. From the source cluster, resynchronize the source SVM from the destination SVM by using the `snapmirror resync` command.

Example

```
source_cluster::> snapmirror resync vs1:
```

2. Verify that the resynchronization operation is complete and the SnapMirror relationship is in the `Snapmirrored` state by using the `snapmirror show` command.

For viewing the detailed status of the relationship, you can use the `-instance` option.

Example

```
source_cluster::> snapmirror show
```

| Source Path | Destination Path | Mirror State | Relationship Status | Total Progress | Healthy | Progress Last Updated |
|-------------|------------------|--------------|---------------------|----------------|---------|-----------------------|
| dvsl: | DP vs1: | Snapmirrored | Idle | - | true | - |

```
source_cluster::> snapmirror show -instance
```

```

Source Path: dvsl:
Destination Path: vs1:
Relationship Type: DP
Relationship Group Type: vserver
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
Mirror State: Snapmirrored

.....
.....

Total Transfer Bytes: -
Total Transfer Time in Seconds: -
```

After the resynchronization, you can only promote load-sharing mirrors and cannot delete them from the source SVM.

Stopping the destination SVM

If you chose to set `identity-preserve` to `true`, you must stop the destination SVM before starting the source SVM to prevent any data corruption.

Before you begin

You must have ensured that all clients of the destination SVM are disconnected.

Steps

1. From the destination cluster, stop the destination SVM by using the `vserver stop` command.

Example

```
destination_cluster::> vserver stop -vserver dvs1
```

2. Verify that the destination SVM is in the `stopped` state by using the `vserver show` command.

Example

```
destination_cluster::> vserver show
Vserver  Type      Subtype  Admin  Operational  Root  Aggregate
-----  -
dvs1     data     default  stopped  stopped      rv    aggr1
```

Note: You must not perform any configuration changes on the destination SVM.

Updating the SnapMirror relationship

You must update the SnapMirror relationship to replicate the changes from the destination SVM to the source SVM since the last resynchronization operation.

Steps

1. From the source cluster, perform a SnapMirror update by using the `snapmirror update` command.

Example

```
source_cluster::> snapmirror update -destination-path vs1:
```

2. Verify that the SnapMirror update operation is complete and the SnapMirror relationship is in the `Snapmirrored` state.

For viewing the detailed status of the relationship, you can use the `-instance` option.

Example

```
source_cluster::> snapmirror show
Source  Destination  Mirror  Relationship  Total  Progress  Healthy  Last  Updated
Path   Type   Path   State         Status   Progress   Healthy   Last  Updated
-----  -
dvs1:  DP     vs1:   Snapmirrored  Idle    -          true     -
```

Breaking the SVM disaster recovery relationship

You must break the SnapMirror relationship created between the source and the destination Storage Virtual Machines (SVMs) for disaster recovery before reactivating the source SVM.

About this task

You must perform this task from the source cluster.

Steps

1. Break the SVM disaster recovery relationship by using the `snapmirror break` command.

Example

```
source_cluster::> snapmirror break -destination-path vs1:
```

2. Verify that the SnapMirror relationship between the source and the destination SVMs is in the Broken-off state by using the `snapmirror show` command.

Example

```
source_cluster::> snapmirror show
```

| Source Path | Type | Destination Path | Mirror State | Relationship Status | Total Progress | Progress Healthy | Last Updated |
|-------------|------|------------------|--------------|---------------------|----------------|------------------|--------------|
| dvsl: | DP | vs1: | Broken-off | Idle | - | true | - |

The source SVM continues to be in the Stopped state and the subtype changes from **dp-destination** to **default**. The state of the volumes in the source SVM changes from **DP** to **RW**.

Starting the source SVM

For providing data access from the source SVM after a disaster, you must reactivate the source SVM by starting it.

Before you begin

The destination SVM must be in the **stopped** state.

Steps

1. On the destination cluster, verify that the destination SVM is in the **stopped** state by using the `vserver show` command.

Example

```
destination_cluster::> vserver show
```

| Vserver | Type | Subtype | Admin State | Operational State | Root Volume | Aggregate |
|---------|------|---------|-------------|-------------------|-------------|-----------|
| dvsl | data | default | stopped | stopped | voll | aggr1 |

2. If you chose to set the `-identity-preserve` option to **true** and the source and destination SVMs are in different subnets, perform the following steps:
 - a. To serve data, set the `-status-admin` option of the replicated LIFs that belong to the subnet of the source SVM to **up** by using the `network interface modify` command.

Example

```
source_cluster::> network interface modify -vserver dvsl
-lif LIF1 -status-admin up
```

- b. Set the `-status-admin` option of the replicated LIFs that belong to the subnet of the destination SVM to **down** by using the `network interface modify` command.

Example

```
source_cluster::> network interface modify -vserver dvsl
-lif LIF1a -status-admin down
```

- From the source cluster, start the source SVM by using the `vserver start` command.

Example

```
source_cluster::> vserver start -vserver vs1
[Job 30] Job succeeded: DONE
```

- Verify that the source SVM is in the **running** state by using the `vserver show` command.

Example

```
source_cluster::> vserver show
-----
```

| Vserver | Type | Subtype | Admin State | Operational State | Root Volume | Aggregate |
|---------|------|---------|-------------|-------------------|-------------|-----------|
| vs1 | data | default | running | running | voll | aggr1 |

```
-----
```

After you finish

The clients might have to remount or rescan as per the protocol requirements for accessing data from the source SVM.

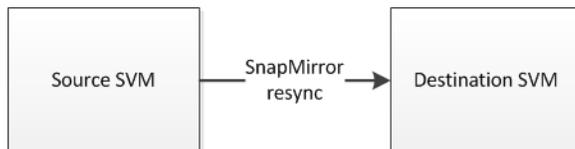
Resynchronizing the destination SVM from the source SVM

You can protect the reactivated source SVM by resynchronizing the data and configuration details from the source SVM to the destination SVM.

Before you begin

- The SVM root volume must not contain any other data apart from metadata because the other data is not replicated.
Root volume metadata such as volume junctions, symbolic links, directories leading to junctions symbolic links are replicated.
- The destination SVM must not contain load-sharing mirrors apart from the load-sharing mirror created for the SVM root volume protection.
- The destination SVM must not contain any new volumes that are not replicated on the source SVM.
You must delete such volumes on the destination SVM to prevent resynchronization failure.

About this task



Steps

- Ensure that the SnapMirror relationship exists between the source and the destination SVMs:
 - Verify that the SnapMirror relationship exists by using the `snapmirror show` command.
 - If the SnapMirror relationship does not exist, then create the SnapMirror relationship by using the `snapmirror-create` command.

Creating a SnapMirror relationship

- From the destination cluster, resynchronize the destination SVM from the source SVM by using the `snapmirror resync` command.

Example

```
destination_cluster::> snapmirror resync dvs1:
```

- Verify that the resynchronization operation is complete and the SnapMirror relationship is in the Snapmirrored state by using the `snapmirror show` command.

For viewing the detailed status of the relationship, you can use the `-instance` option.

Example

```
destination_cluster::> snapmirror show
```

| Source Path | Destination Path | Mirror State | Relationship Status | Total Progress | Healthy | Progress Last Updated |
|-------------|------------------|--------------|---------------------|----------------|---------|-----------------------|
| vs1: | dvs1: | Snapmirrored | Idle | - | true | - |

```
source_cluster::> snapmirror show -instance
```

```

Source Path: vs1:
Destination Path: dvs1:
Relationship Type: DP
Relationship Group Type: vserver
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
Mirror State: Snapmirrored

.....
.....

Total Transfer Bytes: -
Total Transfer Time in Seconds: -
```

After the resynchronization, you can only promote load-sharing mirrors and cannot delete them from the destination SVM.

Where to find additional information

Additional information is available to help you to manage the Storage Virtual Machine (SVM) disaster recovery relationships and set up other data protection solutions.

Reference guides

You can use the following documentation for details about the `snapmirror` commands:

- Man pages for the clustered Data ONTAP commands
[Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)

You can use the following documentation for other data protection solutions:

- Volume-level disaster recovery by using SnapMirror technology between peered clusters
 - [Clustered Data ONTAP 8.3 Volume Disaster Recovery Preparation Express Guide](#)
 - [Clustered Data ONTAP 8.3 Volume Disaster Recovery Express Guide](#)
- Data protection by using tape technology
 - [Clustered Data ONTAP 8.3 Data Protection Tape Backup and Recovery Guide](#)
 - [Clustered Data ONTAP 8.3 NDMP Configuration Express Guide](#)
- Data protection by using SnapMirror and SnapVault technologies
 - [Clustered Data ONTAP 8.3 Data Protection Guide](#)
 - [Clustered Data ONTAP 8.3 Volume Backup Using SnapVault Express Guide](#)
 - [Clustered Data ONTAP 8.3 Volume Restore Using SnapVault Express Guide](#)

Copyright information

Copyright © 1994–2015 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, AltaVault, 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, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to docomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

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

A

- aborting
 - SnapMirror transfers for the SVM disaster recovery relationship [6](#)
- activating
 - destination SVM, workflow for [5](#)

B

- breaking
 - SVM disaster recovery relationship from the destination SVM [6](#)
 - SVM disaster recovery relationship from the source SVM [15](#)

C

- choosing
 - new or existing SVM for reactivation [9](#)
- comments
 - how to send feedback about documentation [22](#)
- creating
 - SVM SnapMirror disaster recovery relationship [12](#)

D

- deleting
 - SnapMirror relationship [10](#)
 - SVM peer relationship [10](#)
- destination
 - requirements for using Express Guide to activate the SVM after a disaster [4](#)
 - resynchronizing the SVM for disaster recovery [17](#)
 - starting the SVM during disaster recovery [7](#)
 - stopping the SVM during disaster recovery [14](#)
 - workflow for activating the SVM [5](#)
- disaster recovery
 - creating a SnapMirror relationship for SVMs [12](#)
 - performing a SnapMirror update [15](#)
 - reactivation workflow for the source SVM [9](#)
 - requirements for using Express Guide to activate the destination SVM [4](#)
 - resynchronizing the destination SVM [17](#)
 - resynchronizing the source SVM [13](#)
 - setting up new source SVM [11](#)
 - starting the destination SVM [7](#)
 - starting the source SVM [16](#)
 - stopping SnapMirror transfers [5](#)
 - stopping the destination SVM [14](#)
 - stopping the source SVM [7](#)
- disaster recovery relationships
 - aborting SnapMirror transfers for the SVM [6](#)
 - breaking from the destination SVM [6](#)
 - breaking from the source SVM [15](#)
- documentation
 - additional information about SVM disaster recovery [19](#)

- how to receive automatic notification of changes to [22](#)
- how to send feedback about [22](#)

E

- express guides
 - additional documentation [19](#)
 - requirements for using SVM disaster recovery guide [4](#)

F

- feedback
 - how to send comments about documentation [22](#)
- flowcharts
 - workflow for activating the destination SVM [5](#)

I

- information
 - how to send feedback about improving documentation [22](#)

P

- peer relationships
 - deleting for SVM [10](#)

Q

- quiescing
 - SnapMirror transfers [5](#)

R

- reactivating
 - existing source SVM [12](#)
 - new source SVM [10](#)
 - workflow for the source SVM [9](#)
- recovering
 - workflow for the source SVM [9](#)
- relationships
 - deleting SnapMirror [10](#)
- resynchronizing
 - destination SVM for disaster recovery [17](#)
 - source the SVM for disaster recovery [13](#)

S

- setting up
 - a new source SVM for disaster recovery [11](#)
- SnapMirror relationships
 - breaking, for SVM disaster recovery [6](#)
 - creating for SVM disaster recovery [12](#)
 - deleting [10](#)

- for disaster recovery, breaking between source and destination SVMs [15](#)
- stopping transfers for SVM disaster recovery relationship [5](#)
- SnapMirror transfers
 - aborting for SVM disaster recovery relationship [6](#)
- Snapmirror updates
 - performing for disaster recovery [15](#)
- source
 - reactivating the existing SVM [12](#)
 - reactivating the new SVM [10](#)
 - requirements for using Express Guide to reactivate the SVMs after a disaster [4](#)
 - resynchronizing SVM for disaster recovery [13](#)
 - setting up the SVM for disaster recovery [11](#)
 - starting SVMs for disaster recovery [16](#)
 - stopping SVMs during disaster recovery [7](#)
- starting
 - destination SVM during disaster recovery [7](#)
 - source SVM after a disaster [16](#)
- stopping
 - destination SVM during disaster recovery [14](#)
 - SnapMirror transfers for the SVM disaster recovery relationship [6](#)
 - source SVM during disaster recovery [7](#)
- suggestions
 - how to send feedback about documentation [22](#)
- SVMs
 - aborting the disaster recovery relationship for [6](#)
 - breaking the disaster recovery relationship for [6](#), [15](#)
 - creating a SnapMirror relationship for disaster recovery [12](#)
 - deleting the peer relationship [10](#)
 - quiescing the SnapMirror relationship [5](#)
 - reactivating a new source [10](#)
 - reactivating the existing source [12](#)
 - reactivation workflow for source [9](#)
 - requirements for using Express Guide to activate the destination after a disaster [4](#)

- resynchronizing the destination for disaster recovery [17](#)
- resynchronizing the source for disaster recovery [13](#)
- setting up a new source [11](#)
- SnapMirror updates between destination and source for disaster recovery [15](#)
- starting the destination during disaster recovery [7](#)
- starting the source after a disaster [16](#)
- stopping the destination during disaster recovery [14](#)
- stopping the source during disaster recovery [7](#)
- workflow for activating destination [5](#)

T

- technical reports
 - additional information about SVM disaster recovery [19](#)
- testing
 - starting the destination SVM in a disaster recovery setup [7](#)
 - stopping the destination in a disaster recovery setup [14](#)
 - stopping the source SVM in a disaster recovery setup [7](#)
- twitter
 - how to receive automatic notification of documentation changes [22](#)

U

- updating
 - SnapMirror relationship for disaster recovery [15](#)

W

- workflows
 - activating the destination SVM [5](#)
 - reactivating the source SVM [9](#)