



NetApp Cloud Volumes ONTAP 9

SVM Disaster Recovery Express Guide

June 2018 | 215-12840_BO
doccomments@netapp.com

Contents

Deciding whether to use the SVM Disaster Recovery Express Guide	3
SVM disaster recovery workflow	4
Quiescing SnapMirror transfers.....	4
Aborting any ongoing SnapMirror transfers.....	5
Breaking the SVM disaster recovery relationship	5
Stopping the source SVM.....	6
Starting the destination SVM	6
Source SVM reactivation workflow	8
Reactivating the source SVM	8
Creating the new source SVM.....	9
Setting up the existing source SVM	11
Stopping the destination SVM.....	14
Updating the SnapMirror relationship	14
Breaking the SVM disaster recovery relationship	15
Starting the source SVM.....	15
Resynchronizing the destination SVM from the source SVM.....	16
Where to find additional information	18
Copyright information	19
Trademark information.....	20
How to send comments about documentation and receive update notifications	21

Deciding whether to use the SVM Disaster Recovery Express Guide

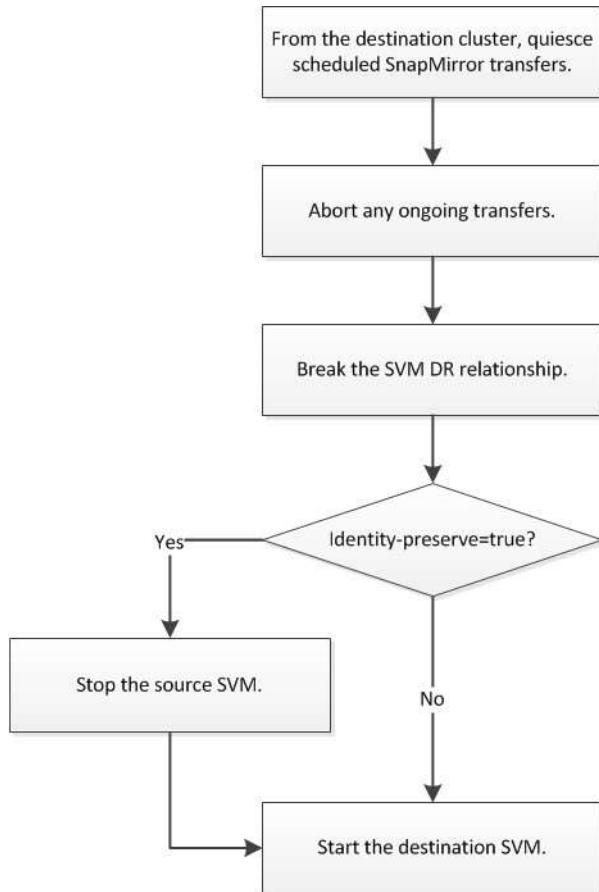
This guide describes how cluster administrators can quickly activate a destination Storage Virtual Machine (SVM) after a disaster, and then reactivate the source SVM. You can use the procedures listed in this guide to recover from an actual disaster.

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

- You are a cluster administrator.
- You have prepared the source SVM for disaster recovery by configuring the destination SVM.
Cloud Volumes ONTAP SVM disaster recovery express preparation
- You are using the 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 scheduled SnapMirror transfers, aborting any ongoing SnapMirror transfers, breaking the SVM disaster recovery (DR) relationship, stopping the source SVM, and starting the destination SVM.



Note: The ONTAP version of the destination SVM must be at or above the version of the source. This is not a requirement for volume async-mirror and XDP relationship.

During a disaster, any new data that is written on the source SVM after the last SnapMirror transfer is lost.

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

Progress
Source      Destination   Mirror      Relationship   Total
Path        Type        Path        State        Status    Progress  Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----
vs1:       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

Progress
Source      Destination   Mirror      Relationship   Total
Path        Type        Path        State        Status    Progress  Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----
vs1:       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      Destination   Mirror      Relationship  Total          Progress
Path        Type        Path        State        Status        Progress    Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----
vs1:      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` 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
Admin      Operational  Root
Vserver   Type      Subtype    State      State      Volume   Aggregate
-----  -----  -----  -----  -----  -----  -----
default  stopped  stopped    vol1      aggr1      vs1      data
```

Starting the destination SVM

In case of a disaster, 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. Start the destination SVM by using the `vserver start` command.

Example

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

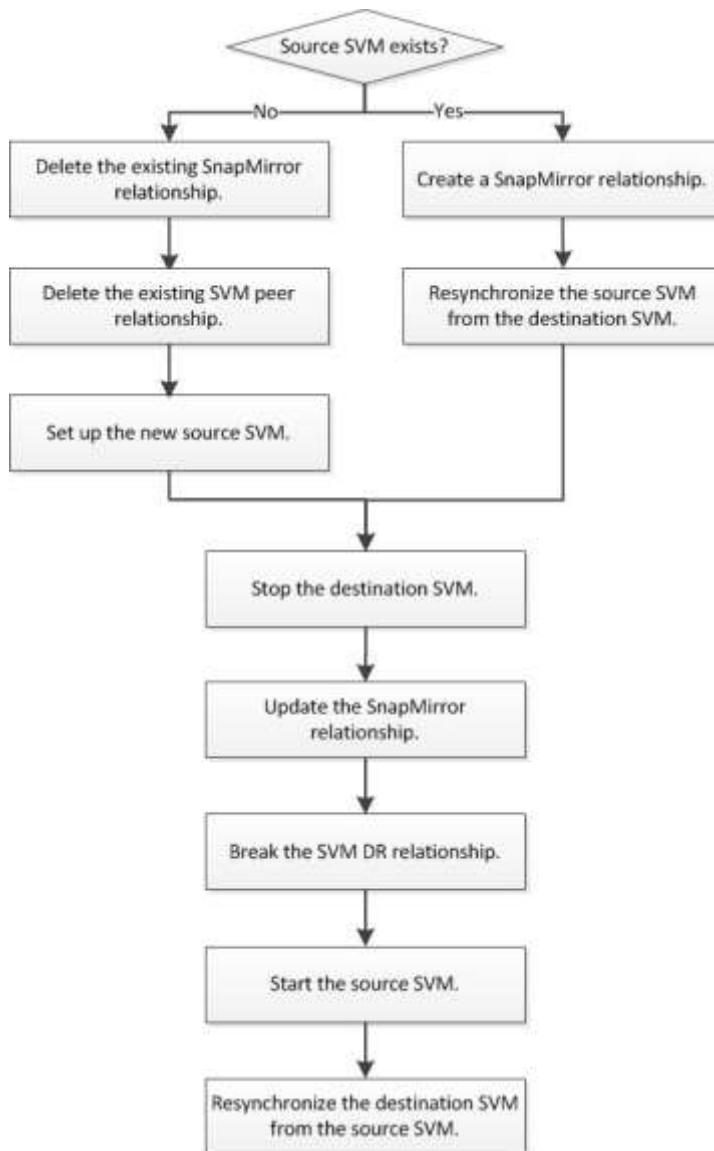
2. Verify that the destination SVM is in the `running` state and the subtype is `default` by using the `vserver show` command.

Example

```
destination_cluster::> vserver show
      Admin          Operational   Root
Vserver  Type   Subtype   State   State   Volume   Aggregate
-----+-----+-----+-----+-----+-----+-----+
  dvs1   data   default  running  running   rv      aggr1
```

Source SVM reactivation workflow

If the source SVM exists after a disaster, you can reactivate it and protect it by recreating 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 9
- [Setting up the existing source SVM](#) on page 11

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 9
2. [Deleting the SVM peer relationship](#) on page 9
3. [Setting up a new source SVM](#) on page 10

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

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
          Peer      Peer      Peering
Vserver  Vserver  State    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 dvs1 -peer-vserver vs1
Info: [Job 47] 'vserver peer delete' job queued
```

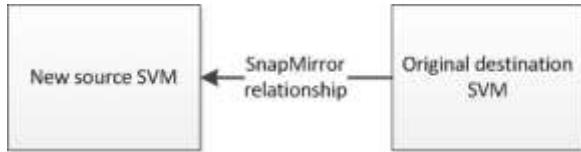
Setting up a new source SVM

After a disaster has occurred, you can set up a new source SVM by creating an SVM disaster recovery relationship to replicate the data and configuration from the original destination SVM.

About this task

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

You must perform this task from the new source cluster.



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

[Cloud Volumes ONTAP SVM disaster recovery express preparation](#)

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 a minimum free space of 10 GB exists.
2. Create a source SVM of subtype `dp-destination`, which is now the destination SVM.
3. Create an SVM peer relationship between the new source SVM and the original destination SVM.
4. Create a SnapMirror policy with the `-discard-configs network` option.
5. Create a SnapMirror relationship between the new source SVM and the original destination SVM.
6. For CIFS: If you chose to replicate data and a subset of the SVM configuration by setting the `identity-preserve` option to `false`, create a CIFS server.

You must use the SnapMirror policy that was created with the `-discard-configs network` option.

7. Initialize the new source SVM.
8. Configure NAS LIFs on the new source SVM.
9. If you chose to replicate data and a subset of the SVM configuration by setting the `-identity-preserve` option to `false`, 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 11
2. [Resynchronizing the source SVM from the destination SVM](#) on page 13

Creating a SnapMirror relationship

When recovering from a disaster, you must create a SnapMirror relationship between the existing source SVM and the destination SVM 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 replicating the configuration. The best practice is to have at least two non-root aggregates with a minimum free space of 10 GB each.
- The source cluster and destination cluster 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.

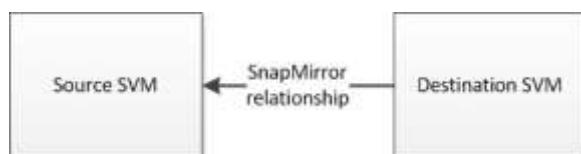
About this task

You must perform this task only if you are reactivating the source SVM for the first time.

If you reactivated the source SVM earlier, you do not have to perform this task because the SnapMirror relationship between the source and destination SVMs is in the `Broken-off` state.

You must set up the disaster recovery relationship by using the same method and configuration that 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 the configuration details when recreating the SnapMirror relationship between the new source and destination SVMs.

You must perform this task from the source cluster.



Steps

1. Create a SnapMirror policy to exclude LIFs from replication by using the `snapmirror policy create` command.

Example

```
source_cluster::> snapmirror policy create -vserver vs1 -policy
exclude_LIF -type async-mirror -discard-configs network
```

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

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

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

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

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

The following command creates the SVM SnapMirror relationship with the SVM names as the `-destination-path` and `-source-path` parameters and uses the SnapMirror policy `exclude_LIF` to exclude LIFs from replication:

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

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

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

3. 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      Destination   Mirror      Relationship   Total      Progress
Path        Type    Path        State        Status      Progress   Healthy  Updated
-----      -----  -----      -----  -----      -----  -----
dvs1:      DP      vs1:      Broken-off  Idle      -        true      -
```

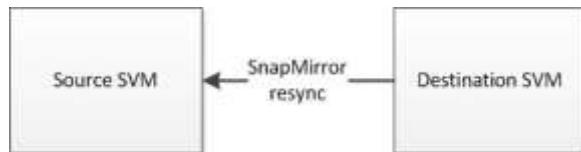
Resynchronizing the source SVM from the destination SVM

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

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.

To view the detailed status of the SnapMirror relationship, you can use the `-instance` option.

Example

```
source_cluster::> snapmirror show
Source          Destination   Mirror          Relationship  Total          Progress
Path           Type     Path      State       Status       Progress  Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----
dvs1:        DP       vs1:     Snapmirrored  Idle        -       true    -
source_cluster::> snapmirror show -instance
Source Path:  dvs1:
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 resynchronization, you cannot delete load-sharing mirrors from the source SVM; you can only promote 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
          Admin          Operational  Root
Vserver  Type   Subtype   State   State   Volume   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
Path        Type        Path        State        Status        Progress  Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----  -----
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      Destination  Mirror          Relationship  Total          Progress
Path        Type        Path        State        Status        Progress  Healthy  Last
-----  -----  -----  -----  -----  -----  -----  -----  -----
dvs1:      DP          vs1:      Broken-off  Idle        -          true    -
```

The source SVM continues to be in the Stopped state and the subtype changes from `dpdestination` 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.

About this task

You must perform this task from the source cluster.

Steps

- 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
```

The `-status-admin` option of the LIFs configured on the source SVM is set to **up**.

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

Example

```
source_cluster::> vserver show
      Admin      Operational  Root
Vserver   Type   Subtype   State   State   Volume   Aggregate
-----  -----  -----  -----
vs1      data   default  running  running  vol1     aggr1
```

Resynchronizing the destination SVM from the source SVM

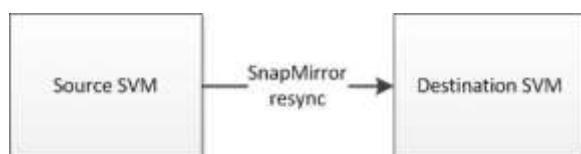
You can protect the reactivated source SVM by resynchronizing the data and configuration details from the source SVM to the destination SVM.

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, and 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 SVM root volume protection.
- The destination SVM must not contain any new protected volumes.
You must delete such volumes on the destination SVM to prevent resynchronization failure.

About this task

The following illustration shows the resynchronization of the destination volume.



Steps

1. Ensure that a SnapMirror relationship exists between the source and the destination SVMs:
 - a. Verify that a SnapMirror relationship exists by using the `snapmirror show` command.
 - b. If a SnapMirror relationship does not exist, then create a SnapMirror relationship by using the `snapmirror create` command.
2. From the destination cluster, resynchronize the destination SVM from the source SVM by using the `snapmirror resync` command.

Example

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

3. 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          Destination   Mirror          Relationship  Total          Progress
Path            Type        Path          State          Status        Progress      Last
-----          -----        -----          -----          -----        -----        -----
vs1:           DP          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.

- Man pages for ONTAP commands
 - *ONTAP 9 commands*
- Data protection conceptual information
 - *ONTAP concepts*

Copyright information

Copyright © 1994–2018 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

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, 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.

<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.

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.

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