




NetApp[®]

New for 8.3.1

Clustered Data ONTAP 8.3

SVM 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: 215-10149_B0
August 2015

Contents

Deciding whether to use this guide	4
Deciding whether to replicate SVM network configuration	6
Configurations replicated in an SVM disaster recovery relationship	6
SVM disaster recovery preparation workflow	9
Preparing the destination cluster	9
Creating a destination SVM	11
Creating the SVM peer relationship	12
Creating a SnapMirror relationship	13
CIFS only: Creating a CIFS server	14
Different subnet: Configuring LIFs on the source SVM	16
Initializing the destination SVM	17
Configuring the network and protocols for data access on the destination SVM	18
Configuring the network and NAS protocols	19
Configuring the network and SAN protocols	19
Monitoring the SnapMirror relationship status	20
Where to find additional information	21
Copyright information	22
Trademark information	23
How to send comments about documentation and receive update notifications	24
Index	25

Deciding whether to use this guide

This guide describes how cluster administrators quickly prepare a data-serving Storage Virtual Machine (SVM) for disaster recovery. You can create and configure a destination SVM in the destination cluster, and replicate data and configuration from the source SVM to the destination SVM.

SVM disaster recovery is asynchronous mirroring of SVM data and configuration. You can choose to replicate all or a subset of the SVM configuration (excluding the network and protocol configuration).

You should use this guide if you want to create and configure a destination SVM for disaster recovery 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 cannot create destination SVMs for SVMs with Infinite Volume.
- The source or destination clusters are not in a MetroCluster configuration.
- Both the source and destination clusters must have SnapMirror license installed.
- The source and destination clusters are peered.
[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)
- The source SVM does not contain data protection volumes (DP) and transition volumes (TDP).
- 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.

If these assumptions are not correct for your situation, you should see the following documentation:

- Synchronous disaster recovery in a MetroCluster configuration
[Clustered Data ONTAP 8.3 MetroCluster Management and Disaster Recovery Guide](#)
- Volume-level disaster recovery by using SnapMirror technology
 - [Clustered Data ONTAP 8.3 Volume Disaster Recovery Preparation Express Guide](#)
 - [Clustered Data ONTAP 8.3 Volume Disaster Recovery Express Guide](#)
- Volume-level backup by using SnapVault
 - [Clustered Data ONTAP 8.3 Volume Backup Using SnapVault Express Guide](#)
 - [Clustered Data ONTAP 8.3 Volume Restore Using SnapVault Express Guide](#)
- Data protection by using tape technology
 - [Clustered Data ONTAP 8.3 NDMP Configuration Express Guide](#)
 - [Clustered Data ONTAP 8.3 Data Protection Tape Backup and Recovery Guide](#)
- Data protection by using SnapMirror and SnapVault technologies
 - [Clustered Data ONTAP 8.3 Data Protection 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.

Deciding whether to replicate SVM network configuration

Depending on the data and network configuration of the source SVM, you can choose to replicate all or a subset of SVM configuration (excluding the network and protocol configuration), when setting up the destination SVM for disaster recovery.

Choices

- Replicate data and all of the SVM configuration in the following scenarios:
 - The destination SVM must have the same NAS configuration as the source SVM.
 - The source SVM is not configured for SAN data access.
 - The destination SVM is not required to provide read-only access.

If the source and destination clusters are in the same network subnet, the destination SVM will use the same IP addresses as the source SVM so that clients can access data by using the same IP addresses and network configuration when the source SVM is not available.

If the source and destination clusters are in different network subnets, on the source SVM, you must create destination LIFs belonging to the subnet of the destination SVM.

LIFs that belong to the subnet of the destination SVM will not be reachable until the destination SVM is started. After the destination SVM is started, the LIFs that are replicated and that belong to the subnet of the source SVM will not be reachable.

- Replicate data and a subset of the SVM configuration in the following scenarios:
 - The destination SVM must have a different NAS configuration from the source SVM.
 - The source SVM is configured only for SAN data access.
 - The destination SVM is required to provide read-only access.

Note: Network, protocol, and name services configurations are not replicated.

You can set up the destination SVM to provide read-only access for NFS clients and SAN hosts by configuring the network and protocols on the destination SVM.

Configurations replicated in an SVM disaster recovery relationship

When you set up the SVM disaster recovery relationship, the value you select for the `-identity-preserve` option of the `snapmirror create` command determines the configurations that are replicated in the destination SVM.

If you set the `-identity-preserve` option to **true**, all the configuration details except the SAN configuration are replicated. If you set the option to **false**, only a subset of the configuration details that are not associated with the network configuration are replicated.

The following table lists the configuration details that are replicated when the `-identity-preserve` option is set to **true** or **false**:

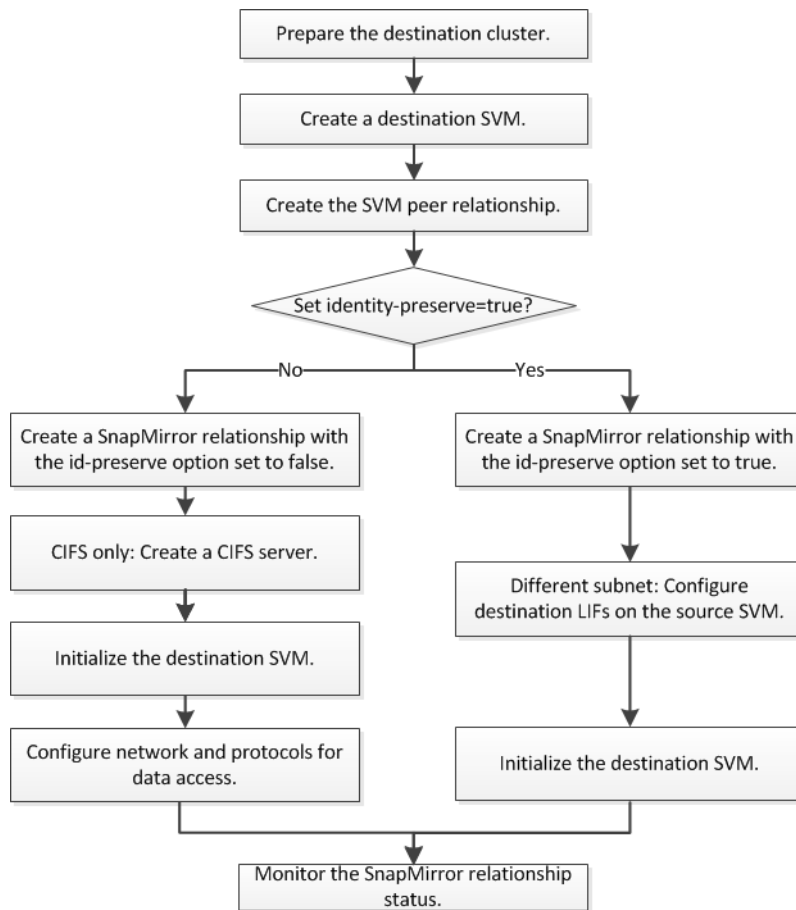
Configuration		Replicated if the <code>-identity-preserve</code> is set to true	Replicated if the <code>-identity-preserve</code> is set to false
CIFS	CIFS server	Yes	No
CIFS policy	Local groups and local user	Yes	Yes
	Privilege	Yes	Yes
	Shadow copy	Yes	Yes
	BranchCache	Yes	Yes
	Server options	Yes	Yes
	Server security	Yes	No
	Home directory, share	Yes	Yes
	Symlink	Yes	Yes
	Fpolicy policy, Fsecurity policy, and Fsecurity NTFS	Yes	Yes
	Name mapping and group mapping	Yes	Yes
Audit information	Yes	Yes	
NFS	Export policies	Yes	No
	Export policy rules	Yes	No
	NFS server and Kerberos configuration	Yes	No
Network	NAS LIFs	Yes	No
	SAN LIFs	No	No
	Firewall policies	Yes	No
	Routes	Yes	No
	Broadcast domain	No	No
	Subnet	No	No
	IPspace	No	No
RBAC	Security certificates	Yes	No
	Login user, public key, role, and role configuration	Yes	Yes
	SSL	Yes	No

Configuration		Replicated if the <code>-identity-preserve</code> is set to true	Replicated if the <code>-identity-preserve</code> is set to false
Name services	DNS and DNS hosts	Yes	No
	UNIX user and UNIX group	Yes	Yes
	Kerberos realm and Kerberos keyblocks	Yes	No
	LDAP and LDAP client	Yes	No
	Netgroup	Yes	No
	NIS	Yes	No
	Web and web access	Yes	No
Volume	Object	Yes	Yes
	Snapshot copies, Snapshot policy, and autodelete policy	Yes	Yes
	Efficiency policy	Yes	Yes
	Quota policy and quota policy rule	Yes	Yes
	Recovery queue	Yes	Yes
Root volume	Namespace	Yes	Yes
	User data	No	No
	Qtrees	No	No
	Quotas	No	No
	File-level QoS	No	No
	Attributes: state of the root volume, space guarantee, size, export policies, autosize, and total number of files	No	No
Storage QoS	QoS policy group	Yes	Yes
Fibre Channel (FC)		No	No
iSCSI		No	No
LUNs	Object	Yes	Yes
	igroups	No	No
	portsets	No	No
SNMP	v3 users	Yes	No

Note: Cluster-level objects such as aggregates are not replicated.

SVM disaster recovery preparation workflow

Preparing the SVM for disaster recovery involves preparing the destination cluster, creating the destination SVM, creating the SVM peer relationship, creating a SnapMirror relationship, initializing the destination SVM, configuring the destination SVM for data access, and monitoring the SnapMirror relationship status.



Preparing the destination cluster

Before you create and configure the destination SVM, you must verify that the cluster peer relationship between the source and destination clusters is healthy, and prepare the destination cluster with required licenses, custom schedules, and sufficient free space.

Steps

1. Verify that the source and destination clusters are peered and the peer cluster is available by using the `cluster peer show` command.

Example

```
destination_cluster::> cluster peer show
Peer Cluster Name   Cluster Serial Number Availability Authentication
-----
source_cluster      1-80-000011         Available         absent
```

2. Install the licenses of the features and protocols used by the source SVM on the destination cluster:
 - a. Identify the licensed features and protocols on the source cluster by using the `license show` command.

Example

```
source_cluster::> license show
(system license show)

Serial Number: 1-80-000011
Owner: source_cluster
Package      Type      Description      Expiration
-----
Base         site     Cluster Base License -
NFS          site     NFS License      -
CIFS         site     CIFS License     -
iSCSI        site     iSCSI License    -
FCP          site     FCP License      -
SnapMirror   site     SnapMirror License -
FlexClone    site     FlexClone License -
7 entries were displayed.
```

- b. Verify that the licenses of the required features and protocols are installed on the destination cluster by using the `system license show` command.

Example

```
destination_cluster::> system license show

Serial Number: 1-80-000011
Owner: source_cluster
Package      Type      Description      Expiration
-----
Base         site     Cluster Base License -
NFS          site     NFS License      -
CIFS         site     CIFS License     -
SnapMirror   site     SnapMirror License -
FlexClone    site     FlexClone License -
5 entries were displayed.
```

- c. If any required feature or protocol is not licensed on the destination cluster, then add the license by using the `system license add` command.

Example

The following example adds FCP license on the destination cluster:

```
destination_cluster::> system license add -license-code
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

License for package "FCP" installed successfully.
```

3. On the destination cluster, create the same custom schedules as on the source cluster:
 - a. Identify the custom schedules on the source cluster by using the `job schedule cron show` command.

Example

```
source_cluster::> job schedule cron show
```

Name	Description
5min	@:00,:05,:10,:15,:20,:25,:30,:35,:40,:45,:50,:55
8hour	@2:15,10:15,18:15
daily	@0:10
hourly	@:05
weekly	Sun@0:15

5 entries were displayed.

- b. On the destination cluster, create the same custom schedules by using the `job schedule cron create` command.

You must specify the exact job names and schedule as that on the source cluster because job schedules are case-sensitive.

Example

```
destination_cluster::> job schedule cron create -name weekly -dayofweek
"Sunday" -hour 0 -minute 15
```

4. Ensure that the destination cluster has 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.

- a. Verify that the non-root aggregate has a minimum free space of 10 GB by using the `storage aggregate show` command.

Example

```
destination_cluster::> storage aggregate show
```

Aggregate	Size Available	Used%	State	#Vols	Nodes	RAID Status
aggr0	6.04GB 1.15GB	81%	online	2	destination_cluster-01	raid_dp, normal
aggr1	5.14GB 2.47GB	52%	online	15	destination_cluster-02	raid_dp, normal

- b. If there is no non-root aggregate with a minimum free space of 10 GB, create a non-root aggregate of size 10 GB by using the `storage aggregate create` command.

Example

```
destination_cluster::> storage aggregate create -aggregate aggr3 -nodes
destination_cluster_01 -diskcount 20 -disksize 10
```

Creating a destination SVM

For protecting the data and configuration information on the source SVM, you must create a destination SVM on the destination cluster.

Before you begin

- The cluster must have at least one non-root aggregate with sufficient space.
- If you want to assign IPspace, you must have created the IPspace.

- If you want to replicate all of the SVM configuration, the IPspaces of source and destination SVMs must have ports belonging to the same subnet.

About this task

A destination SVM can be used for only one source SVM and the destination SVM cannot be a source of any other SVM disaster recovery relationship.

Steps

1. Create the destination SVM by using the `vserver create` command with the subtype `dp-destination`.

You must use the fully qualified domain name (FQDN) of the SVM or another convention that ensures unique SVM names across clusters.

If you do not assign an IPspace, the SVM is created in the default IPspace.

Example

The following command creates a destination SVM in the default IPspace:

```
destination_cluster::> vserver create -vserver dvs1 -subtype dp-destination
[Job 383] Job succeeded:
Vserver creation completed
```

2. Verify the status of the newly created SVM by using the `vserver show` command.

Example

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

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
dvs1	data	dp-destination	running	stopped	-	-

Result

The destination SVM is created without a root volume and is in the `stopped` state.

Creating the SVM peer relationship

You must create an intercluster SVM peer relationship between the source and the destination SVMs to provide an infrastructure for SVM disaster recovery.

About this task

The cluster administrator of the peered cluster must authenticate the SVM peer relationship.

Steps

1. On the destination cluster, create the SVM peer relationship by using the `vserver peer create` command.

Example

```
destination_cluster::> vserver peer create -vserver dvs1 -peer-vserver vs1 -applications
snapmirror -peer-cluster source_cluster
```

The SVM peer relationship is in the `initiated` state.

2. On the source cluster, accept the SVM peer relationship by using the `vserver peer accept` command.

Example

```
source_cluster::> vserver peer accept -vserver vs1 -peer-vserver dvs1
Info: [Job 371] 'vserver peer accept' job queued
```

3. Verify that the SVM peer relationship is in the peered state.

Example

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

Vserver	Peer Vserver	Peer State	Peering Applications
dvs1	vs1	peered	snapmirror

Creating a SnapMirror relationship

You must create a SnapMirror relationship between the source and the destination SVMs for disaster recovery. You can choose to replicate data and all or a subset of the SVM configuration when creating the SnapMirror relationship.

Before you begin

- The CIFS audit consolidation path must be on a non-root volume.
- The SVM root volume must not have any qtrees.

About this task

You must perform this task from the destination cluster.

Steps

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

You can specify the SnapMirror policy and schedule when creating the SnapMirror relationship. The SnapMirror schedule is applicable to all the volumes and configuration of the source SVM.

You can specify the source and the destination SVMs as either 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 SVM disaster recovery relationship with SVM names as the `-destination-path` and `-source-path` parameters:

```
destination_cluster::> snapmirror create -source-path vs1: -destination-path dvs1: -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 vs1 -destination-vserver dvs1 -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 to **false**.

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

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

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 vs2 -destination-vserver dvs2
-type DP -throttle unlimited -policy DPDefault -schedule hourly -identity-preserve
false
```

2. Verify that the SnapMirror relationship is established and is in the **Uninitialized** 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:	DP dvs1:	Uninitialized	Idle	-	true	-

```
destination_cluster::> snapmirror show -instance
```

```

Source Path: vs1:
Destination Path: dvs1:
Relationship Type: DP
Relationship Group Type: vsserver
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
.....
.....

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

CIFS only: Creating a CIFS server

If the source SVM has CIFS configuration, and you chose to set `identity-preserve` to **false**, you must create a CIFS server for the destination SVM. CIFS server is required for some CIFS configurations, such as shares during initialization of the SnapMirror relationship.

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 subtype is `dp-destination` by using the `vserver show` command.

Example

```
destination_cluster::> vserver show
Vserver  Type      Subtype      Admin      Operational  Root
-----  -
dvsl     data      dp-destination  running    running      Volume      Aggregate
-----  -

```

3. Create a LIF by using the `network interface create` command.

Example

```
destination_cluster::>network interface create -vserver dvsl -lif NAS1 -role data -data-
protocol cifs -home-node destination_cluster-01 -home-port a0a-101 -address 192.0.2.128 -
netmask 255.255.255.128
```

4. Create a route by using the `network route create` command.

Example

```
destination_cluster::>network route create -vserver dvsl -destination 0.0.0.0/0
-gateway 192.0.2.1
```

[Clustered Data ONTAP 8.3 Network Management Guide](#)

5. Configure DNS by using the `vserver services dns create` command.

Example

```
destination_cluster::>vserver services dns create -domains mydomain.example.com -vserver
dvsl -name-servers 192.0.2.128 -state enabled
```

6. Add the preferred domain controller by using the `vserver cifs domain preferred-dc add` command.

Example

```
destination_cluster::>vserver cifs domain preferred-dc add -vserver dvsl -preferred-dc
192.0.2.128 -domain mydomain.example.com
```

7. Create the CIFS server by using the `vserver cifs create` command.

Example

```
destination_cluster::>vserver cifs create -vserver dvsl -domain mydomain.example.com
-cifs-server CIFS1
```

[Clustered Data ONTAP 8.3 File Access Management Guide for CIFS](#)

8. Stop the destination SVM by using the `vserver stop` command.

Example

```
destination_cluster::> vserver stop -vserver dvsl
[Job 46] Job succeeded: DONE
```

Different subnet: Configuring LIFs on the source SVM

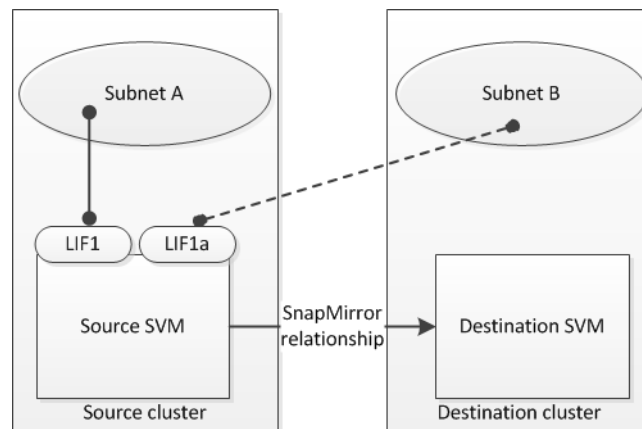
If the source and destination SVMs are in different network subnets and if you chose to set the `-identity-preserve` option to `true`, then you must configure LIFs belonging to the destination SVM network subnet on the source SVM.

Before you begin

You must be aware of the network subnet configuration of the destination SVM.

About this task

In the following illustration, the source SVM is associated with subnet A and the destination SVM is associated with subnet B.



For every LIF on the source SVM in subnet A, you must create a LIF for the destination SVM in subnet B.

In this example, both LIF1 and LIF1a are replicated from the source SVM to the destination SVM. When the source SVM is available, only LIF1 must be active and functional. When the destination SVM is activated during a disaster, only LIF1a must be active and functional to serve data.

You must create a subnet object on the destination cluster and create LIFs on the source cluster.

Steps

1. On the destination cluster, create a subnet object:
 - a. Create a broadcast domain and assign a network port by using the `broadcast-domain create` command.

Example

```
source_cluster::> broadcast-domain create -ipSpace Default
-broadcast-domain bd2 -mtu 1500 -ports localhost:e0c
```

- b. Create a subnet object in the broadcast domain by using the `network subnet create` command.

Example

```
destination_cluster::>network subnet create -subnet-name 20subnet
-broadcast-domain bd2 -subnet 192.0.2.128/25 -ipspace Default
```

2. On the source cluster, create a LIF in the subnet of the destination SVM by using the `network interface create` command.

Example

```
source_cluster::> network interface create -vserver vs1 -lif LIF1a -role data -
data-protocol cifs,nfs,
-home-node source_cluster_01 -home-port e0c -address 192.0.2.129 -netmask-
length 25 -status-admin down
```

3. Create a route for the new LIF by using the `network route create` command.

Example

```
source_cluster::> network route create -vserver vs1 -destination 0.0.0.0/0 -
gateway 192.0.2.1 -metric 20
```

4. For the newly created LIF, verify the status by using the `network interface show`, and verify the route information by using the `network route show`.

Example

```
source_cluster::> network interface show -vserver vs1
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	LIF1 LIF1a	up/up down/down	192.0.2.66/24 192.0.2.129/25	source_cluster_01 source_cluster_01	e0d e0c	true true

```
source_cluster::> network route show -vserver vs1
```

Vserver	Destination	Gateway	Metric
vs1	0.0.0.0/0	192.0.2.1	20

Initializing the destination SVM

You must initialize the destination SVM for the baseline transfer of data and configuration details from the source SVM.

Before you begin

- The source 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 be in the `stopped` state.

About this task

You cannot use tape seeding to initialize the SnapMirror relationship between the source and destination SVMs.

Steps

1. Use the `snapmirror initialize` command to perform a baseline transfer from the source to the destination SVM.

Example

```
destination_cluster::> snapmirror initialize dvs1:
```

2. Use the `snapmirror show` command to verify that the status of the SnapMirror relationship is in the Snapmirrored state.

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

Example

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

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

```
destination_cluster::> snapmirror show -instance
```

```

Source Path: vsl:
Destination Path: dvs1:
Relationship Type: DP
Relationship Group Type: vserver
SnapMirror Schedule: -
SnapMirror Policy Type: async-mirror
SnapMirror Policy: DPDefault
.....
.....
Total Transfer Bytes: -
Total Transfer Time in Seconds: -
```

After you finish

You must not associate the destination SVM with any other source SVM.

Configuring the network and protocols for data access on the destination SVM

If you chose to set the `identity-preserve` option to `false` or if the source SVM has SAN configuration, you must configure the network and protocols on the destination SVM for data access when a disaster occurs.

Before you begin

The destination SVM must be started and in the **running** state.

About this task

You must configure the SAN network and protocols if the source SVM is configured for SAN protocols because SAN network configuration is not replicated.

Choices

- [Configuring the network and NAS protocols](#) on page 19
- [Configuring the network and SAN protocols](#) on page 19

Configuring the network and NAS protocols

If the source SVM has NAS configuration, you must configure the network and NAS protocols on the destination SVM for data access in the event of a disaster.

About this task

This procedure provides the high-level steps that are required to complete the network and NAS protocols configuration on the destination SVM. Detailed information about these steps are available in other clustered Data ONTAP documentation.

- [Clustered Data ONTAP 8.3 Network Management Guide](#)
- [Clustered Data ONTAP 8.3 File Access Management Guide for NFS](#)
- [Clustered Data ONTAP 8.3 File Access Management Guide for CIFS](#)

Steps

1. Start the destination SVM by using the `vserver start` command.
2. Create data LIFs by using the `network interface create` command.
3. Create routes for the data LIFs by using the `network route create` command.
4. Configure name services such as LDAP, NIS, and DNS by using the `vserver services` command.
5. Configure CIFS, NFS, or both the protocols by using the `vserver cifs create` and `vserver nfs create` commands.
6. If the source SVM has CIFS configuration, stop the destination SVM by using the `vserver stop` command.

After you finish

You can set up read-only access for NFS clients from the destination SVM.

Configuring the network and SAN protocols

If the source SVM has SAN configuration, you must configure the network and SAN protocols on the destination SVM for data access in the event of a disaster.

About this task

This procedure provides the high-level steps that are required to complete the network and SAN protocols configuration on the destination SVM. Detailed information about these steps are available in other clustered Data ONTAP documentation.

- [Clustered Data ONTAP 8.3 Network Management Guide](#)
- [Clustered Data ONTAP 8.3 SAN Administration Guide](#)

Steps

1. Start the destination SVM by using the `vserver start` command.
2. Create data LIFs by using the `network interface create` command.
3. Create igroups for the LUNs by using the `lun igroup create` command.

4. Map the LUNs to the igroups by using the `lun mapping create` command.
5. Configure iSCSI, FC, or both the protocols by using the `vserver iscsi create` and `vserver fcp create` commands.

After you finish

You can set up read-only access for SAN hosts from the destination SVM.

Monitoring the SnapMirror relationship status

You can monitor the status of the SnapMirror relationship between the source and the destination SVMs to ensure that the updates are occurring per the schedule.

About this task

SNMP is not supported for monitoring the SnapMirror relationships between the source and destination SVMs.

Step

1. Use the `snapmirror show -instance` command to view the details of the SnapMirror relationship status.

Example

```
destination_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
        Relationship Status: Idle
        ..
        ..
        Snapshot Checkpoint: -
        Newest Snapshot: vserverdr.4eb1flaa-
e4ba-11e3-9b97-005056af93d7.2014-05-26_095857
        Newest Snapshot Timestamp: 05/26 09:58:57
        Exported Snapshot: vserverdr.4eb1flaa-
e4ba-11e3-9b97-005056af93d7.2014-05-26_095857
        Exported Snapshot Timestamp: 05/26 09:58:57
        Healthy: true
        Unhealthy Reason: -
        ...
        ...
        Last Transfer Type: update
        Last Transfer Error: -
        Last Transfer From: vs1:
        Last Transfer End Timestamp: 05/26 10:05:24
        ...
        ...
Lag Time: 2:0:15
Identity Preserve Vserver DR: true
```

Where to find additional information

Additional documentation is available to help you activate the destination Storage Virtual Machine (SVM) to test the disaster recovery setup or when a disaster occurs. You can also learn more about how to re-create and reactivate the source SVM after the disaster.

Reference guides

You can activate the destination SVM, and re-create and reactivate the source SVM by using the following documentation:

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

You can use the `snapmirror` commands to manage the SVM disaster recovery relationships.

- [Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#)
- [Clustered Data ONTAP 8.3 Data Protection 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

C

- CIFS
 - configuring the network and protocols on destination SVM [19](#)
 - creating a CIFS server for SVMs [14](#)
- clusters
 - verifying the peer relationship for the source and destination [9](#)
- comments
 - how to send feedback about documentation [24](#)
- configuration details
 - deciding to replicate for disaster recovery [6](#)
- configuring
 - network and protocols on the destination SVM [18](#)
- creating
 - CIFS server for SVMs [14](#)
 - destination SVM for disaster recovery [11](#)
 - intercluster SVM peer relationships [12](#)
 - LIFs on the source SVM [16](#)

D

- destination cluster
 - preparing for SVM disaster recovery [9](#)
- destination SVMs
 - creating for disaster recovery [11](#)
 - creating LIFs [16](#)
 - initializing [17](#)
 - requirements for using Express Guide to configure for disaster recovery [4](#)
- disaster recovery
 - additional information about SVM [21](#)
 - creating a destination SVM for [11](#)
 - creating a SnapMirror relationship for SVMs [13](#)
 - creating LIFs on the source [16](#)
 - deciding to replicate all or some configuration details [6](#)
 - preparation workflow for SVM [9](#)
 - requirements for using Express Guide to prepare the SVM [4](#)
 - SVM configuration details that are replicated [6](#)
- documentation
 - additional information about SVM disaster recovery [21](#)
 - how to receive automatic notification of changes to [24](#)
 - how to send feedback about [24](#)

E

- express guides
 - additional documentation on SVM disaster recovery [21](#)
 - requirements for using SVM disaster recovery preparation guide [4](#)
 - SVM disaster recovery preparation workflow [9](#)

F

- FC
 - configuring the network and protocols on the destination SVM [19](#)
- feedback
 - how to send comments about documentation [24](#)
- flowcharts
 - SVM disaster recovery preparation [9](#)

I

- information
 - how to send feedback about improving documentation [24](#)
- initializing
 - the destination SVM [17](#)
- iSCSI
 - configuring the network and protocols on the destination SVM [19](#)

L

- licenses
 - preparing destination cluster for disaster recovery [9](#)
- LIFs
 - creating on the source SVM [16](#)

M

- monitoring
 - SnapMirror relationship status [20](#)

N

- NAS
 - configuring on destination SVM [19](#)
 - SVM configuration details replicated for disaster recovery [6](#)
- network configuration
 - on the destination SVM [18](#)
- NFS
 - configuring the network and protocols on destination SVM [19](#)

P

- peer relationships
 - creating between intercluster SVMs [12](#)
 - verifying for the source and destination clusters [9](#)
- preparation
 - workflow for SVM disaster recovery [9](#)
- preparing
 - destination cluster for disaster recovery [9](#)
- protocols
 - configuring on the destination SVM [18](#)
 - NAS, configuring the network on destination SVM [19](#)

SAN, configuring on destination SVM [19](#)

R

relationships

- creating intercluster SVM peer [12](#)
- creating SnapMirror [13](#)

S

SAN

- configuring on destination SVM [19](#)
- SVM configuration details replicated for disaster recovery [6](#)

schedules

- preparing destination cluster for disaster recovery [9](#)

SnapMirror relationships

- creating for SVM disaster recovery [13](#)
- monitoring the status of [20](#)
- requirements for using Express Guide to configure SVM disaster recovery [4](#)

source SVMs

- creating LIFs for the destination SVM [16](#)

suggestions

- how to send feedback about documentation [24](#)

SVMs

- configuration details replicated in a disaster recovery relationship [6](#)
- configuring the network and NAS protocols on destination [19](#)
- configuring the network and SAN protocols on destination [19](#)

creating a CIFS server [14](#)

creating a SnapMirror relationship for disaster recovery [13](#)

creating destination for disaster recovery [11](#)

creating intercluster peer relationships [12](#)

creating LIFs on the source SVM [16](#)

deciding to replicate all or some configuration details for disaster recovery [6](#)

initializing the destination [17](#)

monitoring the SnapMirror status of [20](#)

preparation workflow for disaster recovery [9](#)

requirements for using the Express Guide to prepare for disaster recovery [4](#)

T

twitter

- how to receive automatic notification of documentation changes [24](#)

V

verifying

- the cluster peer relationship [9](#)

W

workflows

- SVM disaster recovery preparation [9](#)