



Replicate SVM configurations

ONTAP 9

NetApp
February 20, 2026

Table of Contents

- Replicate SVM configurations 1
 - ONTAP SnapMirror SVM replication workflow 1
 - Criteria for placing volumes on ONTAP SnapMirror destination SVMs 1
 - Replicate an entire ONTAP SVM configuration 1
 - Exclude LIFs and related network settings from ONTAP SnapMirror SVM replication 5
 - Exclude network, name service, and other settings from SVM replication with ONTAP 8
 - Specify local tiers to use for ONTAP SnapMirror SVM DR relationships 11
 - Create an SMB server for an ONTAP SnapMirror destination SVM in a DR relationship 11
 - Exclude volumes from an ONTAP SnapMirror SVM DR relationship 13

Replicate SVM configurations

ONTAP SnapMirror SVM replication workflow

SnapMirror SVM replication involves creating the destination SVM, creating a replication job schedule, and creating and initializing a SnapMirror relationship.

You should determine which replication workflow best suits your needs:

- [Replicate an entire SVM configuration](#)
- [Exclude LIFs and related network settings from SVM replication](#)
- [Exclude network, name service, and other settings from SVM configuration](#)

Criteria for placing volumes on ONTAP SnapMirror destination SVMs

When replicating volumes from the source SVM to the destination SVM, it's important to know the criteria for selecting aggregates.

Aggregates are selected based on the following criteria:

- Volumes are always placed on non-root aggregates.
- Non-root aggregates are selected based on the available free space and the number of volumes already hosted on the aggregate.

Aggregates with more free space and fewer volumes are given priority. The aggregate with the highest priority is selected.

- Source volumes on FabricPool aggregates are placed on FabricPool aggregates on the destination with the same tiering-policy.
- If a volume on the source SVM is located on a Flash Pool aggregate, then the volume is placed on a Flash Pool aggregate on the destination SVM, if such an aggregate exists and has enough free space.
- If the `-space-guarantee` option of the volume that is replicated is set to `volume`, only aggregates with free space greater than the volume size are considered.
- The volume size grows automatically on the destination SVM during replication, based on the source volume size.

If you want to pre-reserve the size on the destination SVM, you must resize the volume. The volume size does not shrink automatically on the destination SVM based on the source SVM.

If you want to move a volume from one aggregate to another, you can use the `volume move` command on the destination SVM.

Replicate an entire ONTAP SVM configuration

You can create an SVM disaster recovery (SVM DR) relationship to replicate one SVM configuration to another. In the event of a disaster at the primary site, you can quickly

activate the destination SVM.

Before you begin

The source and destination clusters and SVMs must be peered. For more information, see [Create a cluster peer relationship](#) and [Create an SVM intercluster peer relationship](#).

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

About this task

This workflow assumes that you are already using a default policy or a custom replication policy.

Beginning with ONTAP 9.9.1, when you use the mirror-vault policy, you can create different snapshot policies on the source and destination SVM, and the snapshots on the destination are not overwritten by snapshots on the source. For more information, see [Understanding SnapMirror SVM replication](#).

Complete this procedure from the destination. If you need to create a new protection policy, for instance, when your source storage VM has SMB configured, you should create the policy and use the **Identity preserve** option. For details see [Create custom data protection policies](#).

Steps

You can perform this task from System Manager or the ONTAP CLI.

System Manager

1. On the destination cluster, click **Protection > Relationships**.
2. Under **Relationships**, click **Protect** and choose **Storage VMs (DR)**.
3. Select a protection policy. If you created a custom protection policy, select it, then choose the source cluster and storage VM you want to replicate. You can also create a new destination storage VM by entering a new storage VM name.
4. If desired, change the destination settings to override identity preserve and to include or exclude network interfaces and protocols.
5. Click **Save**.

CLI

1. Create a destination SVM:

```
vserver create -vserver <SVM_name> -subtype dp-destination
```

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named `svm_backup`:

```
cluster_dst:> vserver create -vserver svm_backup -subtype dp-destination
```

Learn more about `vserver create` in the [ONTAP command reference](#).

2. From the destination cluster, create an SVM peer relationship using the `vserver peer create` command.

For more information, see [Create an SVM intercluster peer relationship](#).

Learn more about `vserver peer create` in the [ONTAP command reference](#).

3. Create a replication job schedule:

```
job schedule cron create -name <job_name> -month <month> -dayofweek <day_of_week> -day <day_of_month> -hour <hour> -minute <minute>
```

For `-month`, `-dayofweek`, and `-hour`, you can specify `all` to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek
saturday -hour 3 -minute 0
```

Learn more about `job schedule cron create` in the [ONTAP command reference](#).

4. From the destination SVM or the destination cluster, create a replication relationship:

```
snapmirror create -source-path <SVM_name>: -destination-path
<SVM_name>: -type <DP|XDP> -schedule <schedule> -policy <policy>
-identity-preserve true
```



You must enter a colon (:) after the SVM name in the `-source-path` and `-destination-path` options.

The following example creates a SnapMirror DR relationship using the default `MirrorAllSnapshots` policy:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy
MirrorAllSnapshots -identity-preserve true
```

The following example creates a unified replication relationship using the default `MirrorAndVault` policy:

```
cluster_dst:> snapmirror create -source-path svm1: -destination-path
svm_backup: -type XDP -schedule my_daily -policy MirrorAndVault
-identity-preserve true
```

Assuming you have created a custom policy with the policy type `async-mirror`, the following example creates a SnapMirror DR relationship:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy my_mirrored
-identity-preserve true
```

Assuming you have created a custom policy with the policy type `mirror-vault`, the following example creates a unified replication relationship:

```
cluster_dst::> snapmirror create -source-path svm1: -destination
-path svm_backup: -type XDP -schedule my_daily -policy my_unified
-identity-preserve true
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

5. Stop the destination SVM:

```
vserver stop -vserver <SVM_name>
```

The following example stops a destination SVM named `svm_backup`:

```
cluster_dst::> vserver stop -vserver svm_backup
```

Learn more about `vserver stop` in the [ONTAP command reference](#).

6. From the destination SVM or the destination cluster, initialize the SVM replication relationship:

```
snapmirror initialize -source-path <SVM_name>: -destination-path  
<SVM_name>:
```



You must enter a colon (:) after the SVM name in the `-source-path` and `-destination-path` options.

The following example initializes the relationship between the source SVM, `svm1`, and the destination SVM, `svm_backup`:

```
cluster_dst::> snapmirror initialize -source-path svm1: -destination  
-path svm_backup:
```

Learn more about `snapmirror initialize` in the [ONTAP command reference](#).

Exclude LIFs and related network settings from ONTAP SnapMirror SVM replication

If the source and destination SVMs are in different subnets, you can use the `-discard -configs network` option of the `snapmirror policy create` command to exclude LIFs and related network settings from SVM replication.

Before you begin

The source and destination clusters and SVMs must be peered.

For more information, see [Create a cluster peer relationship](#) and [Create an SVM intercluster peer relationship](#).

About this task

The `-identity-preserve` option of the `snapmirror create` command must be set to `true` when you

create the SVM replication relationship.

Steps

1. Create a destination SVM:

```
vserver create -vserver SVM -subtype dp-destination
```

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named `svm_backup`:

```
cluster_dst:> vserver create -vserver svm_backup -subtype dp-destination
```

2. From the destination cluster, create an SVM peer relationship using the `vserver peer create` command.

For more information, see [Create an SVM intercluster peer relationship](#).

Learn more about `vserver peer create` in the [ONTAP command reference](#).

3. Create a job schedule:

```
job schedule cron create -name job_name -month month -dayofweek day_of_week  
-day day_of_month -hour hour -minute minute
```

For `-month`, `-dayofweek`, and `-hour`, you can specify `all` to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek  
"Saturday" -hour 3 -minute 0
```

4. Create a custom replication policy:

```
snapmirror policy create -vserver SVM -policy policy -type async-  
mirror|vault|mirror-vault -comment comment -tries transfer_tries -transfer  
-priority low|normal -is-network-compression-enabled true|false -discard  
-configs network
```

The following example creates a custom replication policy for SnapMirror DR that excludes LIFs:

```
cluster_dst::> snapmirror policy create -vserver svml -policy  
DR_exclude_LIFs -type async-mirror -discard-configs network
```

The following example creates a custom replication policy for unified replication that excludes LIFs:

```
cluster_dst::> snapmirror policy create -vserver svml -policy
unified_exclude_LIFs -type mirror-vault -discard-configs network
```



Consider creating the same custom SnapMirror policy on the source cluster for future failover and failback scenarios.

Learn more about `snapmirror policy create` in the [ONTAP command reference](#).

5. From the destination SVM or the destination cluster, run the following command to create a replication relationship:

```
snapmirror create -source-path SVM: -destination-path SVM: -type DP|XDP
-schedule schedule -policy policy -identity-preserve true|false -discard
-configs true|false
```



You must enter a colon (:) after the SVM name in the `-source-path` and `-destination-path` options. See the examples below.

The following example creates a SnapMirror DR relationship that excludes LIFs:

```
cluster_dst::> snapmirror create -source-path svml: -destination-path
svm_backup: -type XDP -schedule my_weekly -policy DR_exclude_LIFs
-identity-preserve true
```

The following example creates a SnapMirror unified replication relationship that excludes LIFs:

```
cluster_dst::> snapmirror create -source-path svml: -destination-path
svm_backup: -type XDP -schedule my_weekly -policy unified_exclude_LIFs
-identity-preserve true -discard-configs true
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

6. Stop the destination SVM:

```
vserver stop
```

SVM name

The following example stops the destination SVM named `svm_backup`:

```
cluster_dst::> vserver stop -vserver svm_backup
```

7. From the destination SVM or the destination cluster, initialize a replication relationship:

```
snapmirror initialize -source-path SVM: -destination-path SVM:
```

The following example initializes the relationship between the source, `svm1` and the destination, `svm_backup`:

```
cluster_dst::> snapmirror initialize -source-path svm1: -destination  
-path svm_backup:
```

Learn more about `snapmirror initialize` in the [ONTAP command reference](#).

After you finish

You must configure the network and protocols on the destination SVM for data access in the event a disaster occurs.

Related information

- [snapmirror create](#)
- [snapmirror initialize](#)
- [snapmirror policy create](#)

Exclude network, name service, and other settings from SVM replication with ONTAP

You might want to exclude network, name service, and other settings from an SVM replication relationship to avoid conflicts or configuration differences with the destination SVM.

You can use the `-identity-preserve false` option of the `snapmirror create` command to replicate only the volumes and security configurations of an SVM. Some protocol and name service settings are also preserved.

About this task

For a list of preserved protocol and name service settings, see [Configurations replicated in SVM DR relationships](#).

Before you begin

The source and destination clusters and SVMs must be peered.

For more information, see [Create a cluster peer relationship](#) and [Create an SVM intercluster peer relationship](#).

Steps

1. Create a destination SVM:

```
vserver create -vserver SVM -subtype dp-destination
```

The SVM name must be unique across the source and destination clusters.

The following example creates a destination SVM named `svm_backup`:

```
cluster_dst:> vserver create -vserver svm_backup -subtype dp-destination
```

2. From the destination cluster, create an SVM peer relationship using the `vserver peer create` command.

For more information, see [Create an SVM intercluster peer relationship](#).

Learn more about `vserver peer create` in the [ONTAP command reference](#).

3. Create a replication job schedule:

```
job schedule cron create -name job_name -month month -dayofweek day_of_week  
-day day_of_month -hour hour -minute minute
```

For `-month`, `-dayofweek`, and `-hour`, you can specify `all` to run the job every month, day of the week, and hour, respectively.



The minimum supported schedule (RPO) for FlexVol volumes in an SVM SnapMirror relationship is 15 minutes. The minimum supported schedule (RPO) for FlexGroup volumes in an SVM SnapMirror relationship is 30 minutes.

The following example creates a job schedule named `my_weekly` that runs on Saturdays at 3:00 a.m.:

```
cluster_dst::> job schedule cron create -name my_weekly -dayofweek  
"Saturday" -hour 3 -minute 0
```

4. Create a replication relationship that excludes network, name service, and other configuration settings:

```
snapmirror create -source-path SVM: -destination-path SVM: -type DP|XDP  
-schedule schedule -policy policy -identity-preserve false
```



You must enter a colon (:) after the SVM name in the `-source-path` and `-destination-path` options. See the examples below. You must run this command from the destination SVM or the destination cluster.

The following example creates a SnapMirror DR relationship using the default `MirrorAllSnapshots` policy. The relationship excludes network, name service, and other configuration settings from SVM replication:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path  
svm_backup: -type XDP -schedule my_daily -policy MirrorAllSnapshots  
-identity-preserve false
```

The following example creates a unified replication relationship using the default `MirrorAndVault` policy. The relationship excludes network, name service, and other configuration settings:

```
cluster_dst:> snapmirror create svm1: -destination-path svm_backup:
-type XDP -schedule my_daily -policy MirrorAndVault -identity-preserve
false
```

Assuming you have created a custom policy with the policy type `async-mirror`, the following example creates a SnapMirror DR relationship. The relationship excludes network, name service, and other configuration settings from SVM replication:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path
svm_backup: -type XDP -schedule my_daily -policy my_mirrored -identity
-preserve false
```

Assuming you have created a custom policy with the policy type `mirror-vault`, the following example creates a unified replication relationship. The relationship excludes network, name service, and other configuration settings from SVM replication:

```
cluster_dst::> snapmirror create -source-path svm1: -destination-path
svm_backup: -type XDP -schedule my_daily -policy my_unified -identity
-preserve false
```

Learn more about `snapmirror create` in the [ONTAP command reference](#).

5. Stop the destination SVM:

```
vserver stop
```

SVM name

The following example stops a destination SVM named `dvs1`:

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

6. If you are using SMB, you must also configure an SMB server.

See [SMB only: Creating an SMB server](#).

7. From the destination SVM or the destination cluster, initialize the SVM replication relationship:

```
snapmirror initialize -source-path SVM_name: -destination-path SVM_name:
```

Learn more about `snapmirror initialize` in the [ONTAP command reference](#).

After you finish

You must configure the network and protocols on the destination SVM for data access in the event a disaster occurs.

Specify local tiers to use for ONTAP SnapMirror SVM DR relationships

After a disaster recovery SVM is created, you can use the `aggr-list` option with `vserver modify` command to limit which local tiers are used to host SVM DR destination volumes.

Steps

1. Create a destination SVM:

```
vserver create -vserver SVM -subtype dp-destination
```

2. Modify the disaster recovery SVM's `aggr-list` to limit the local tiers that are used to host the disaster recovery SVM's volume:

```
cluster_dest::> vserver modify -vserver SVM -aggr-list <comma-separated-list>
```

Create an SMB server for an ONTAP SnapMirror destination SVM in a DR relationship

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

Steps

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

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

Learn more about `vserver start` in the [ONTAP command reference](#).

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

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

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

Learn more about `vserver show` in the [ONTAP command reference](#).

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

```
destination_cluster::>network interface create -vserver dvs1 -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
```

Learn more about `network interface create` in the [ONTAP command reference](#).

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

```
destination_cluster::>network route create -vserver dvs1 -destination
0.0.0.0/0
-gateway 192.0.2.1
```

Network management

Learn more about `network route create` in the [ONTAP command reference](#).

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

```
destination_cluster::>vserver services dns create -domains
mydomain.example.com -vserver
dvs1 -name-servers 192.0.2.128 -state enabled
```

Learn more about `vserver services dns create` in the [ONTAP command reference](#).

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

```
destination_cluster::>vserver cifs domain preferred-dc add -vserver dvs1
-preferred-dc
192.0.2.128 -domain mydomain.example.com
```

Learn more about `vserver cifs domain preferred-dc add` in the [ONTAP command reference](#).

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

```
destination_cluster::>vserver cifs create -vserver dvs1 -domain
mydomain.example.com
-cifs-server CIFS1
```

Learn more about `vserver cifs create` in the [ONTAP command reference](#).

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

```
destination_cluster::> vserver stop -vserver dvs1
[Job 46] Job succeeded: DONE
```

Learn more about `vserver stop` in the [ONTAP command reference](#).

Exclude volumes from an ONTAP SnapMirror SVM DR relationship

By default, all RW data volumes of the source SVM are replicated. If you do not want to protect all the volumes on the source SVM, you can use the `-vserver-dr -protection unprotected` option of the `volume modify` command to exclude volumes from SVM replication.

Steps

1. Exclude a volume from SVM replication:

```
volume modify -vserver SVM -volume volume -vserver-dr-protection unprotected
```

Learn more about `volume modify` in the [ONTAP command reference](#).

The following example excludes the volume `volA_src` from SVM replication:

```
cluster_src::> volume modify -vserver SVM1 -volume volA_src -vserver-dr
-protection unprotected
```

If you later want to include a volume in the SVM replication that you originally excluded, run the following command:

```
volume modify -vserver SVM -volume volume -vserver-dr-protection protected
```

The following example includes the volume `volA_src` in the SVM replication:

```
cluster_src::> volume modify -vserver SVM1 -volume volA_src -vserver-dr
-protection protected
```

2. Create and initialize the SVM replication relationship as described in [Replicating an entire SVM configuration](#).

Copyright information

Copyright © 2026 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.