



Preparing storage systems for SnapMirror and SnapVault replication

SnapManager Oracle

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapmanager-oracle/unix-installation-cmode/concept_understanding_the_differences_between_snapmirror_and_snapvault.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Preparing storage systems for SnapMirror and SnapVault replication. 1
 - Understanding the differences between SnapMirror and SnapVault 1
 - Preparing storage systems for SnapMirror replication 2
 - Preparing storage systems for SnapVault replication 3

Preparing storage systems for SnapMirror and SnapVault replication

You can use SnapManager with Data ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with Data ONTAP SnapVault technology to archive backups efficiently to disk. Before you can perform these tasks in SnapManager, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.



You cannot configure both a SnapMirror relationship and a SnapVault relationship on the same clustered Data ONTAP source volume. You must configure these relationships on different source volumes.

Related information

[Understanding the differences between SnapMirror and SnapVault](#)

[Preparing storage systems for SnapMirror replication](#)

[Preparing storage systems for SnapVault replication](#)

Understanding the differences between SnapMirror and SnapVault

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. SnapVault is archiving technology designed for disk-to-disk Snapshot copy replication for standards compliance and other governance-related purposes.

These objectives account for the different balance each technology strikes between the goals of backup currency and backup retention:

- SnapMirror stores *only* the Snapshot copies that reside in primary storage, because, in the event of a disaster, you need to be able to fail over to the most recent version of primary data you know to be good. Your organization, for example, might mirror hourly copies of production data over a ten-day span. As the failover use case implies, the equipment on the secondary system needs to be equivalent or nearly equivalent to the equipment on the primary system to serve data efficiently from mirrored storage.
- SnapVault, in contrast, stores Snapshot copies *whether or not* they currently reside in primary storage because, in the event of an audit, access to historical data is likely to be as important as access to current data. You might want to keep monthly Snapshot copies of your data over a 20-year span (to comply with government accounting regulations for your business, for example). Since there is no requirement to serve data from secondary storage, you can use slower, less expensive disks on the vault system.

Of course, the different weights SnapMirror and SnapVault give to backup currency and backup retention ultimately derive from the 255-Snapshot copy limit for each volume. Where SnapMirror retains the most recent copies, SnapVault retains the copies taken over the longest period of time.

Preparing storage systems for SnapMirror replication

Before you can use SnapManager's integrated SnapMirror technology to mirror Snapshot copies, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship. Upon initialization, SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other, less recent Snapshot copies on the source volume to the destination volume.

- You must create the source and destination volumes in peered clusters with peered Storage Virtual Machines (SVMs). For more information, see the *Clustered Data ONTAP Cluster Peering Express Guide*.
- You must be a cluster administrator.
- For Snapshot copy verification on the destination volume, the source and destination Storage Virtual Machines (SVMs) must have a management LIF as well as a data LIF. The management LIF must have the same DNS name as the SVM. Set the management LIF role to data, the protocol to none, and the firewall policy to mgmt.

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create a SnapMirror relationship. The following procedure assumes you are using the CLI. For information about how to create SnapMirror relationship by using OnCommand System Manager, see the *Clustered Data ONTAP Volume Disaster Recovery Preparation Express Guide*.

The following illustration shows the procedure for initializing a SnapMirror relationship:

1. Identify the destination cluster.
2. On the destination cluster, use the volume create command with the -typeDP option to create a SnapMirror destination volume that is either the same or greater in size than the source volume.



The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2-GB destination volume named dstvolB in SVM2 on the aggregate node01_aggr:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate  
node01_aggr -type DP  
-size 2GB
```

3. On the destination SVM, use the snapmirror create command with the -type DP parameter to create a SnapMirror relationship.

The DP type defines the relationship as a SnapMirror relationship.

The following command creates a SnapMirror relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2. By default, the command assigns the default SnapMirror policy DPDefault:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path  
SVM2:dstvolB  
-type DP
```



Do not define a mirror schedule for the SnapMirror relationship. SnapManager does that for you when you create a backup schedule.

If you do not want to use the default SnapMirror policy, you can invoke the `snapmirror policy create` command to define a SnapMirror policy.

4. Use the `snapmirror initialize` command to initialize the relationship.

The initialization process performs a baseline transfer to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume. It also transfers any other Snapshot copies on the source volume to the destination volume.

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Related information

[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)

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

Preparing storage systems for SnapVault replication

Before you can use SnapManager's integrated SnapVault technology to archive Snapshot copies to disk, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship. On initialization, SnapVault makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

- You must create the source and destination volumes in peered clusters with peered Storage Virtual Machines (SVMs). For more information, see the *Clustered Data ONTAP Cluster Peering Express Guide*.
- You must be a cluster administrator.

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create SnapVault relationships. The following procedure assumes you are using the CLI. For information about how to create SnapVault relationship by using OnCommand System Manager, see the *Clustered Data ONTAP Volume Backup Using SnapVault Express Guide*.

The following illustration shows the procedure for initializing a SnapVault relationship:

1. Identify the destination cluster.
2. On the destination cluster, use the volume create command with the -typeDP option to create a SnapVault destination volume that is the same or greater in size than the source volume.



The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2-GB destination volume named dstvolB in SVM2 on the aggregate node01_aggr:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate  
node01_aggr -type DP  
-size 2GB
```

3. On the destination SVM, use the snapmirror policy create command to create a SnapVault policy.

The following command creates the SVM-wide policy SVM1-vault:

```
SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-vault
```



Do not define a cron schedule or Snapshot copy policy for the SnapVault relationship. SnapManager does that for you when you create a backup schedule.

4. Use the snapmirror policy add-rule command to add a rule to the policy that defines the following Snapshot copy labels and the retention policy for each label:

- Daily
- Weekly
- Monthly
- Hourly
- Unlimited **Important:** The labels are case-sensitive.

These are fixed labels that SnapManager uses. You select one of these options when you archive a backup. You must execute this command once for each of the rules you are adding.

+ The following command adds a rule to the SVM1-vault policy that defines the “daily” label and specifies that thirty Snapshot copies matching the label should be kept in the vault:

+

```
SVM2::> snapmirror policy add-rule -vserver SVM2 -policy SVM1-vault  
-snapmirror-label Daily -keep 30
```

5. Use the snapmirror create command with the -type XDP parameter and the -policy parameter to create a SnapVault relationship and assign a vault policy.

The XDP type defines the relationship as a SnapVault relationship.

The following command creates a SnapVault relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2. It assigns the policy named SVM1-vault:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path  
SVM2:dstvolB  
-type XDP -policy SVM1-vault
```

6. Use the snapmirror initialize command to initialize the relationship.

The initialization process performs a baseline transfer to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Related information

[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)

[Clustered Data ONTAP 8.3 Volume Backup Using SnapVault Express Guide](#)

Copyright information

Copyright © 2024 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.