



Preparing for IBM Domino backup and restore

Snap Creator Framework

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snap-creator-framework/domino-ops/reference_storage_layout_requirements.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Preparing for IBM Domino backup and restore 1
- Storage layout requirements 1
- SnapMirror and SnapVault setup 2

Preparing for IBM Domino backup and restore

Before you deploy the IBM Domino plug-in, make sure that your storage system and hosts meet minimum resource requirements. You also need to configure storage system layouts for databases, and optionally set up SnapMirror and SnapVault relationships.

For Snap Creator Server and Agent installation requirements, see the [Snap Creator Framework 4.1.2 Installation Guide](#). Pay particular attention to the IBM Domino preinstallation requirements for the Agent host:

- On UNIX hosts, you must create symbolic links to IBM Domino shared object files.
- On Windows hosts, you must add the IBM Domino installation path to the PATH environment variable.

Storage layout requirements

A typical IBM Domino environment has at least three Domino volumes, one each for Domino data, Domino transaction logs, and the plug-in changeinfo directory. Many sites also have volumes for Domino DAOS and for view rebuilds.

The IBM Domino plug-in uses the changeinfo directory for changes recorded during backup operations and for copies of transaction logs used in up-to-the-minute restore operations. It is a best practice to store the changeinfo directory on a separate volume, to avoid inadvertently overwriting the information and to make it easier to back up.

You may also find it useful to have separate volumes for Domino DAOS (if it is enabled) and for view rebuilds. When Domino rebuilds a view (for example, when a user opens a view whose index has been deleted or when `updll -R` is run), it may generate temporary files to sort the data for rapid view rebuilding.

By default, these temporary files are located in the system's temporary folder or in the Domino data folder. IBM recommends changing the location of the temporary files to a different drive to distribute disk I/O and to ensure adequate space to rebuild views. To change the temporary folder used for view rebuilds, add the `View_Rebuild_Dir` setting to the `notes.ini` file.

The following table shows the preferred volume layout:

Volume	Contents	Notes
Volume 1	Domino data	FC, SAS, or SSD drives preferred.
Volume 2	Domino transaction logs	FC, SAS, or SSD drives preferred.
Volume 3	changeinfo	Stores changes recorded during backup operations and copies of transaction logs for use in up-to-the-minute restore operations.
Volume 4	View rebuild	Optional. Stores temp files created during index updates. Can use RAM disk. Add <code>View_Rebuild_Dir</code> setting to <code>notes.ini</code> file.

Volume	Contents	Notes
Volume 5	DAOS repository	Optional. Contains .dlo files from DAOS. Low I/O requirements make this a good candidate for SATA drives.



In virtual environments, guest-mounted disks are preferred.

SnapMirror and SnapVault setup

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.

Before you can use Snap Creator with these products, you need to configure a data-protection relationship between the source and destination volumes, then initialize the relationship.



The procedures in this section describe how to set up replication relationships in clustered Data ONTAP. You can find information about setting up these relationships in Data ONTAP operating in 7-Mode in the .

Preparing storage systems for SnapMirror replication

Before you can use 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 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 documents CLI usage.



If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

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, and 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. 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
```

Preparing storage systems for SnapVault replication

Before you can use to perform disk-to-disk backup replication, 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 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 documents CLI usage.



If you are storing database files and transaction logs on different volumes, you must create relationships between the source and destination volumes for the database files and between the source and destination volumes for the transaction logs.

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 size as or larger 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. does that for you when you create a backup schedule.

4. 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, and assigns the policy SVM1-vault:

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

5. 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
```

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.