



Restoring and recovering SAP HANA databases

Snap Creator Framework

NetApp
August 28, 2025

This PDF was generated from https://docs.netapp.com/us-en/snap-creator-framework/sap-hana-ops/task_restoring_and_recovering_databases_from_primary_storage.html on August 28, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- Restoring and recovering SAP HANA databases. 1
 - Restoring and recovering databases from primary storage 2
 - Restoring and recovering databases from secondary storage 14
- Resuming a SnapVault relationship after a restore 24
 - Resuming a SnapVault relationship with Data ONTAP operating in 7-Mode 25
 - Resuming a SnapVault relationship with clustered Data ONTAP 26
- Restoring databases after primary storage failure 27
 - Restoring databases after a primary storage failure with Data ONTAP operating in 7-Mode 27
 - Restoring databases after a primary storage failure with clustered Data ONTAP 28

Restoring and recovering SAP HANA databases

You use SAP HANA Studio and Snap Creator to restore and recover SAP HANA databases.

1. Within SAP HANA Studio:
 - a. Select Recover for the SAP HANA system.
 - b. SAP HANA system is shut down.
 - c. Select the recovery type.
 - d. Provide log backup locations.
 - e. List of data backups is shown
 - f. Select backup to see the external backup ID.
2. For a storage system running clustered Data ONTAP only:
 - a. Only required if any other backup than the latest has been used for the restore.
 - b. Only required for "Volume SnapRestore" from primary storage.
 - c. Deactivate SnapVault relationships
3. Within Snap Creator:
 - a. Select "Restore" for the SAP HANA system.
 - b. Select restore from primary or secondary storage, depending on the availability of the backup at the primary storage.
 - c. Select storage controller, volume name, and Snapshot copy name. Snapshot copy name correlates with the backup ID earlier.
 - d. For multinode SAP HANA systems, multiple volumes need to be restored:
 - i. Choose **Add more restore items**.
 - ii. Select storage controller, volume name, and Snapshot copy name.
 - iii. Repeat this process for all required volumes.
 - e. For multitenant database containers (MDC) single tenant database systems, both the SYSTEM and TENANT databases are restored.
 - f. Restore process is started
 - g. Restore finished for all volumes.
4. At the database nodes, unmount and mount all data volumes to clean "Stale NFS Handles."
5. Within SAP HANA Studio:
 - a. Select **Refresh** on backup list.
 - b. Select available backup for recovery (green item).
 - c. Start recovery process.
 - d. For multitenant database containers (MDC) single tenant database systems, start the recovery process first for the SYSTEM database, and then for the TENANT database.
 - e. The SAP HANA system is started.
6. (Optional) Resume SnapVault relationships for all restored volumes.



At the storage systems, this step is only required if a backup other than the latest one has been used for the restore.

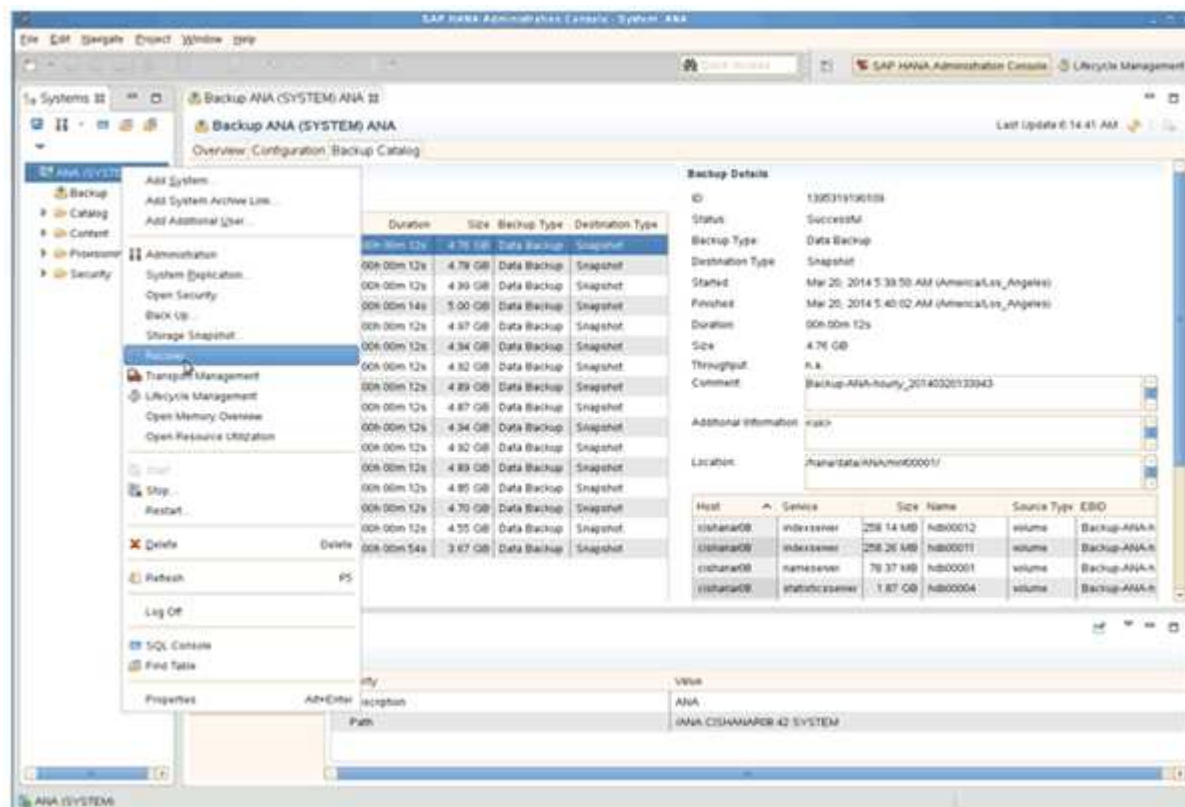
Restoring and recovering databases from primary storage

You can restore and recover the database from the primary storage.



You cannot restore file-based backup copies from Snap Creator.

1. Within SAP HANA Studio, select **Recover** for the SAP HANA system.



The SAP HANA system shuts down.

2. Select the recovery type and click **Next**.

Recovery of System ANA (on vshanar08)

Specify Recovery Type

Select a recovery type.

☒ Recover the database to its most recent state¹

☐ Recover the database to the following point in time¹

Date: 2014-03-20 Time: 03:28:17

Select Time Zone: (GMT-07:00) Pacific Daylight Time


System time used (GMT): 2014-03-20 10:28:17

☐ Recover Database to a Specific Data Backup¹

Advanced >>


? < Back Next > Cancel Finish

3. Provide the log backup locations and click **Next**.



The list of available backups you see is based on the content of the backup catalog.

4. Select the required backup and record the external backup ID.



5. Deactivate the SnapVault relationship.



This step is only required with clustered Data ONTAP.


If you need to restore a Snapshot copy that is older than the Snapshot copy currently used as the base Snapshot copy for SnapVault, you must first deactivate the SnapVault relationship in clustered Data ONTAP. To do that, execute the following commands on the backup cluster console:

```
hana::> snapmirror quiesce -destination-path hana2b:backup_hana_data
Operation succeeded: snapmirror quiesce for destination
hana2b:backup_hana_data.

hana::> snapmirror delete -destination-path hana2b:backup_hana_data
Operation succeeded: snapmirror delete the relationship with destination
hana2b:backup_hana_data.


hana::> snapmirror release -destination-path hana2b:backup_hana_data
[Job 6551] Job succeeded: SnapMirror Release Succeeded
```

6. In the Snap Creator GUI, select the SAP HANA system, then select **Actions > Restore**.

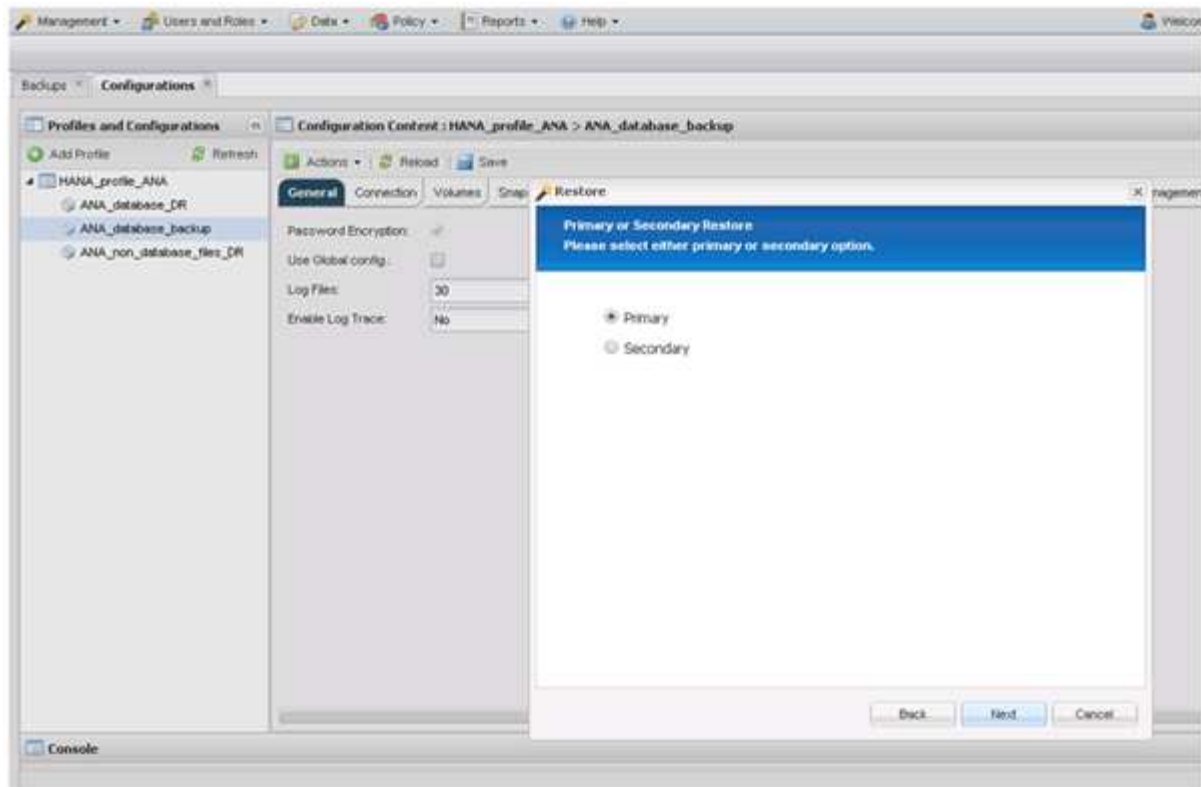


The Welcome to the Snap Creator Framework Restore Wizard screen appears.

7. Click **Next**.



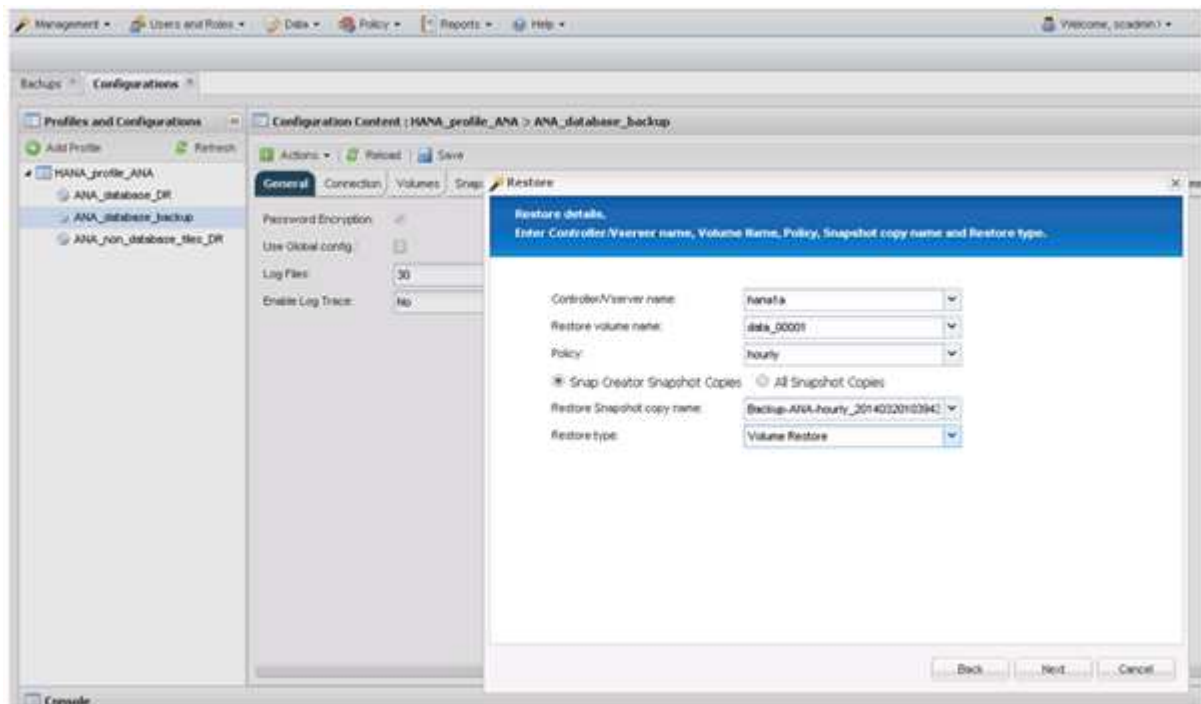
8. Select **Primary** and click **Next**.




9. Select restore from primary storage.

10. Select the storage controller, the volume name, and the Snapshot name.


The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio.



11. Click **Finish**.




12. Click **Yes** to add more restore items.




13. Select the storage controller, the additional volume name, and the Snapshot name.

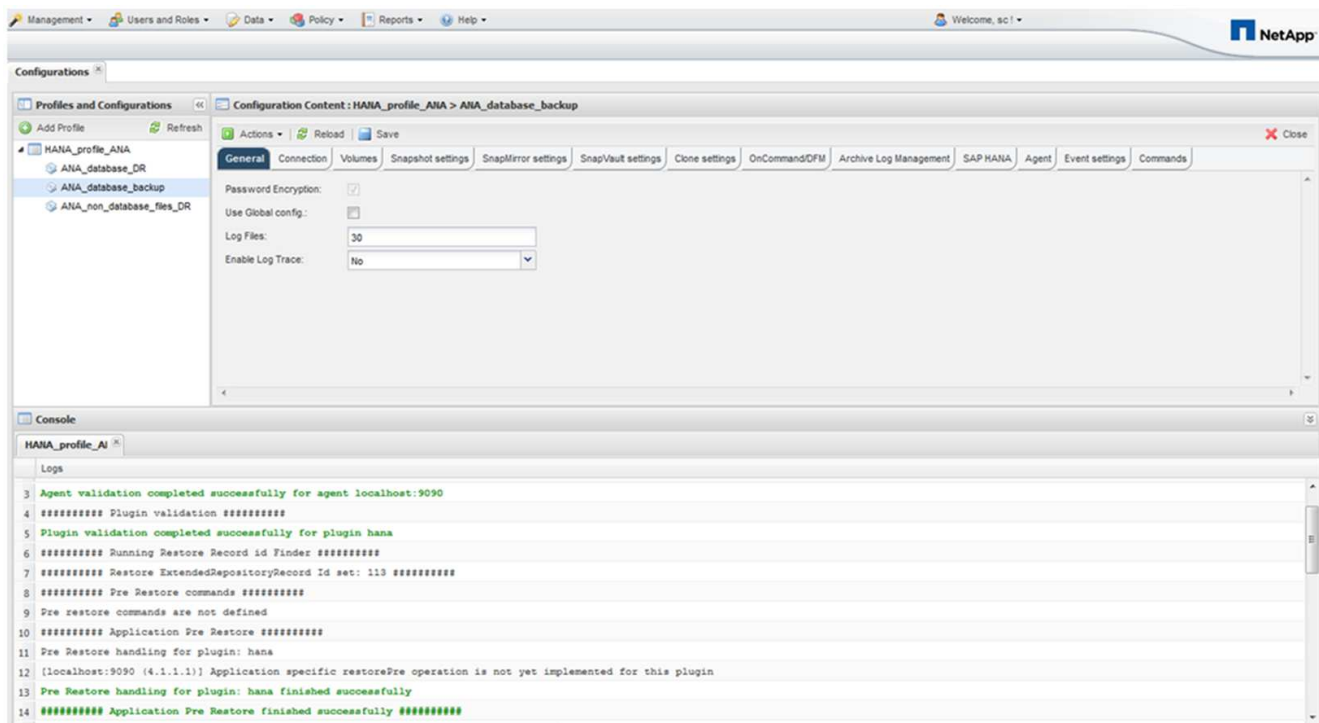
The Snapshot name correlates with the backup ID that has been selected within SAP HANA Studio.



14. Repeat steps 10 through 13 until all required volumes are added; in our example, data_00001, data_00002, and data_00003 need to be selected for the restore process.
15. When all volumes are selected, click **OK** to start the restore process.



The restore process is started.




Wait until the restore process is finished.

16. On each database node, remount all data volumes to clean Stale NFS Handles.

In the example, all three volumes need to be remounted at each database node.

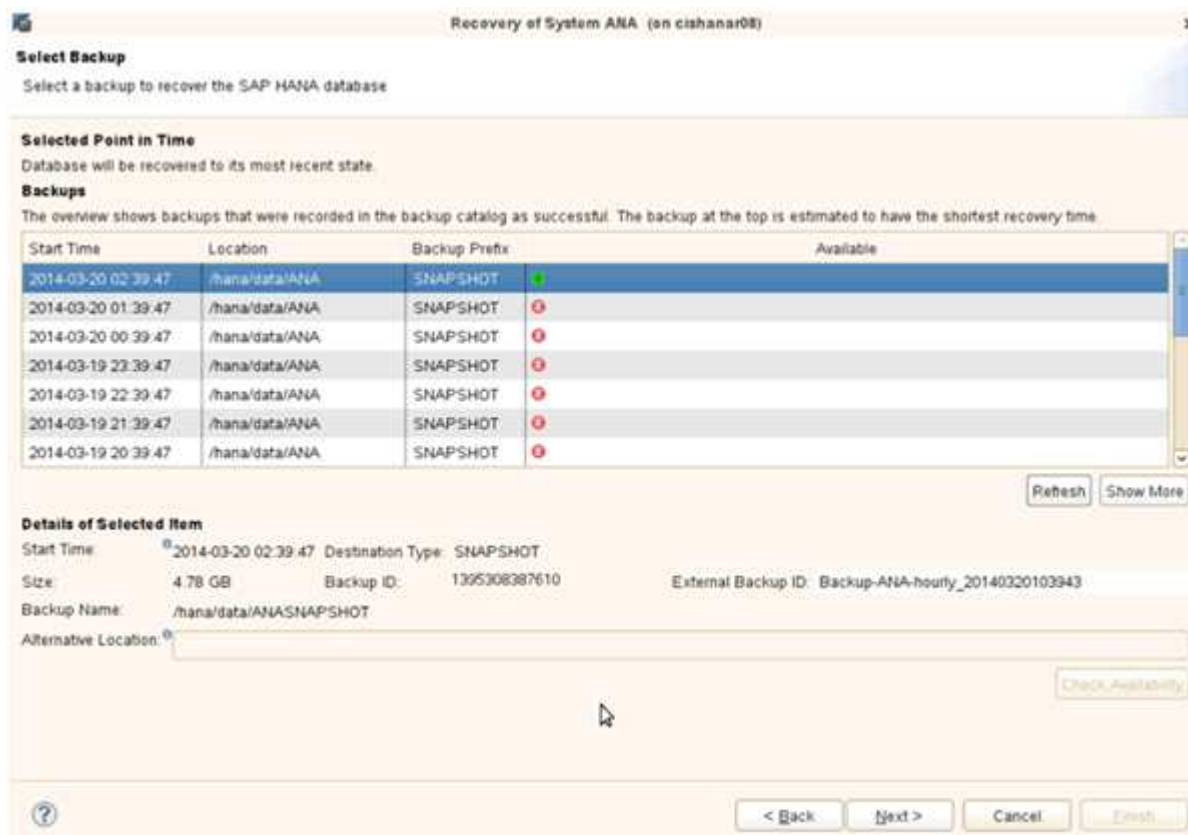
```
mount -o remount /hana/data/ANA/mnt00001
mount -o remount /hana/data/ANA/mnt00002
mount -o remount /hana/data/ANA/mnt00003
```

17. Go to SAP HANA Studio and click **Refresh** to update the list of available backups.



The backup that has been restored with Snap Creator is shown with a green icon in the list of backups.

18. Select the backup and click **Next**.




19. Select other settings as required and click **Next**.

The screenshot shows a window titled "Recovery of System ANA: (sn.cishanar08)". The "Other Settings" tab is active, with the instruction "Ensure that the snapshot is available in the SAP HANA system." Below this, there are three sections:


- Check Availability of Log Backups:** A text block explaining that the system can check for log backups at the start of recovery. Below this is a checkbox for "File System" (checked) and a checkbox for "Third-Party Backup Tool (Backupint)" (unchecked).
- Initialize Log Area:** A text block explaining that log entries will be deleted after recovery. Below this is a checkbox for "Initialize Log Area" (unchecked).
- Install New License Key:** A text block explaining that the old license key is no longer valid. Below this is a checkbox for "Install New License Key" (unchecked), followed by a text input field and a "Browse" button.

At the bottom of the window, there is a navigation bar with a help icon, "< Back", "Next >" (highlighted with a mouse cursor), "Cancel", and "Finish" buttons.


20. Click **Finish**.



The recovery process begins.




21. After recovery is finished, resume the SnapVault relationships, if required.



Restoring and recovering databases from secondary storage

You can restore and recover the database from the secondary storage.

1. Within SAP HANA Studio, select **Recover** for the SAP HANA system.



The SAP HANA system will be shut down.

2. Select the recovery type and click **Next**.

Recovery of System ANA (on cishanar08)

Specify Recovery Type

Select a recovery type.

☒ Recover the database to its most recent state[?]

☐ Recover the database to the following point in time[?]

Date: 2014-04-07 Time: 00:44:22

Select Time Zone: (GMT-07:00) Pacific Daylight Time


System time used (GMT): 2014-04-07 07:44:22

☐ Recover Database to a Specific Data Backup[?]

Advanced >>


? < Back Next > Cancel Finish

3. Provide log backup locations and click **Next**.




The list of available backups appear based on the content of the backup catalog.


4. Select the required backup and write down external backup ID.



5. Go to the Snap Creator GUI.
6. Select the SAP HANA system, and then click **Actions > Restore**.




The Welcome screen appears.




7. Click **Next**.


8. Select **Secondary** and click **Next**.




9. Enter the required information. The Snapshot name correlates with the backup ID that has been selected in




10. Select **Finish**.



11. Click **Yes** to add more items to restore.



12. Provide the required information for all volumes that need to be restored. In the setup data_00001, data_00002, and data_00003 need to be selected for the restore process.



13. When all volumes are selected, select **OK** to start the restore process.


Wait until the restore process is finished.

14. On each database node remount all data volumes to clean “Stale NFS Handles.”

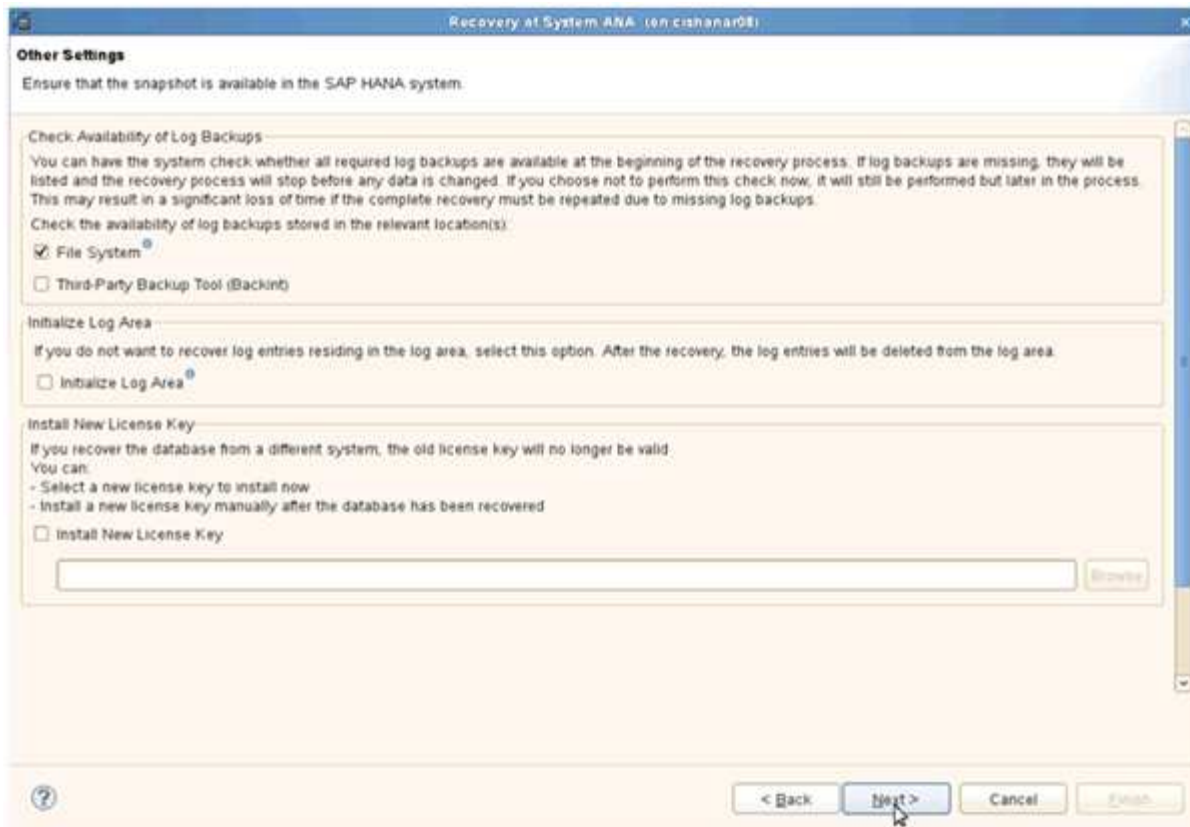
In the example, all three volumes need to be remounted at each database node.

```
mount -o remount /hana/data/ANA/mnt00001
mount -o remount /hana/data/ANA/mnt00002
mount -o remount /hana/data/ANA/mnt00003
```


15. Go to SAP HANA Studio and click **Refresh** to update the backup list.




16. The backup that has been restored with Snap Creator is shown with a green icon in the list of backups. Select the backup and click **Next**.
17. Select other settings as required and click **Next**.



18. Click **Finish**.



The recovery process begins.



19. After the recovery process is finished, resume the SnapVault Relationships, if required.



Resuming a SnapVault relationship after a restore

Any restore that is not done using the latest Snapshot backup will delete the SnapVault relationship at the primary storage systems.

After the restore and recovery process is finished, the SnapVault relationship has to be resumed so that backups can be executed again with Snap Creator. Otherwise, Snap Creator will issue an error message, because it can't find the SnapVault relationship anymore at the primary storage systems.

The data transfer that is required will be based on a delta transfer, if there is still a common Snapshot copy between the source volume and the destination volume.

Resuming a SnapVault relationship with Data ONTAP operating in 7-Mode

If you restore using a Snapshot backup other than the latest one, you need to resume the SnapVault relationship so that Snap Creator can continue to run backups.

1. Resume the SnapVault relationship with Data ONTAP operating in 7-Mode by entering the following command. `snapvault start -r -S source_controller:source_volumebackup_controller:backup_volume`

Perform this step for all volumes belonging to the SAP HANA database.

```
hana2b> snapvault start -r -S hana1a:/vol/data_00001/mnt00001
hana2b:/vol/backup_data_00001/mnt00001
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr 7 14:08:21 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00001/mnt00001 to hana1a:/vol/data_00001/mnt00001 was
successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana2b> snapvault start -r -S hana1b:/vol/data_00002/mnt00002
hana2b:/vol/backup_data_00002/mnt00002
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr 7 14:09:49 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00002/mnt00002 to hana1b:/vol/data_00002/mnt00002 was
successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```

hana2b> snapvault start -r -S hana1a:/vol/data_00003/mnt00003
hana2b:/vol/backup_data_00003/mnt00003
The resync base snapshot will be: Backup-ANA-SV_daily_20140406200000
Resync may alter the data in this qtree.
Are you sure you want to resync the qtree? y
Mon Apr  7 14:10:25 CEST [hana2b:replication.dst.resync.success:notice]:
SnapVault resync of
/vol/backup_data_00003/mnt00003 to hana1a:/vol/data_00003/mnt00003 was
successful.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.

```

When the data transfer is finished, you can again schedule backups by using Snap Creator.

Resuming a SnapVault relationship with clustered Data ONTAP

If you restore using a Snapshot backup other than the latest one, you need to resume the SnapVault relationship so that Snap Creator can continue to run backups.

1. Re-create and resynchronize the SnapVault relationship.

```

hana::> snapmirror create -source-path hana1a:hana_data -destination
-path
hana2b:backup_hana_data -type XDP
Operation succeeded: snapmirror create the relationship with destination
hana2b:backup_hana_data.

hana::> snapmirror resync -destination-path hana2b:backup_hana_data
-type XDP

Warning: All data newer than Snapshot copy sc-backup-
daily_20140430121000 on volume
hana2b:backup_hana_data will be deleted.
Do you want to continue? {y|n}: y
[Job 6554] Job is queued: initiate snapmirror resync to destination
"hana2b:backup_hana_data".
[Job 6554] Job succeeded: SnapMirror Resync Transfer Queued

```

2. To actually restart the SnapVault transfer, a manual Snapshot copy is required.

```
hana::> snapshot create -vserver hanala -volume hana_data -snapshot
sv_resync

hana::> snapshot modify -vserver hanala -volume hana_data -snapshot
sv_resync -snapmirror-label daily

hana::> snapmirror update -destination-path hana2b:backup_hana_data
Operation is queued: snapmirror update of destination
hana2b:backup_hana_data.
```

3. Verify that the SnapVault relationship appears in the destination list.

```
hana::> snapmirror list-destinations -source-path hanala:hana_data
Progress
Source          Destination      Transfer  Last
Relationship
Path           Type  Path           Status Progress  Updated      Id
-----
-----
hanala:hana_data
          XDP  hana2b:backup_hana_data
                        Transferring
                        38.46KB   04/30 18:15:54
                        9137fb83-
cba9-11e3-85d7-123478563412
```

Restoring databases after primary storage failure

After a primary storage failure, or when all Snapshot copies are deleted from the volumes at the primary storage, Snap Creator will not be able to handle the restore, because there will no longer be a SnapVault relationship on the primary storage systems.

Restoring databases after a primary storage failure with Data ONTAP operating in 7-Mode

You can restore an SAP HANA database after a primary storage system running Data ONTAP operating in 7-Mode fails.

1. In this case, the restore has to be executed directly on the secondary storage system by using the following command: `snapvault restore --s snapshot_name -S backup_controller:backup_volumesource_controller:source_volume`

Perform this step for all volumes belonging to the SAP HANA database.

```
hana1a> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00001/mnt00001 hana1a:/vol/data_00001/mnt00001
Restore will overwrite existing data in /vol/data_00001/mnt00001.
Are you sure you want to continue? y
Thu Apr 10 11:55:55 CEST [hana1a:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00001/mnt00001.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana1a> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00003/mnt00003 hana1a:/vol/data_00003/mnt00003
Restore will overwrite existing data in /vol/data_00003/mnt00003.
Are you sure you want to continue? y
Thu Apr 10 11:58:18 CEST [hana1a:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00003/mnt00003.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

```
hana1b> snapvault restore -s Backup-ANA-SV_hourly_20140410103943 -S
hana2b:/vol/backup_data_00002/mnt00002 hana1b:/vol/data_00002/mnt00002
Restore will overwrite existing data in /vol/data_00002/mnt00002.
Are you sure you want to continue? y
Thu Apr 10 12:01:29 CEST [hana1b:vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/data_00002/mnt00002.
Transfer started.
Monitor progress with 'snapvault status' or the snapmirror log.
```

When the restore process is finished, you use SAP HANA to perform the recovery.

Restoring databases after a primary storage failure with clustered Data ONTAP

You can restore an SAP HANA database after a primary storage system running clustered Data ONTAP fails.

Assuming the primary volume is lost completely, you need to create a new primary volume and then restore from the backup volume.

1. Create a primary volume with type data protection.

```
hana::> volume create -vserver hanala -volume hana_data -aggregate
aggr_sas_101 -size 300G -state online -type DP -policy default -autosize
-mode grow_shrink -space-guarantee none
-snapshot-policy none -foreground true
[Job 6744] Job is queued: Create hana_data.
[Job 6744] Job succeeded: Successful
```

2. Restore all data from the backup volume.

```
hana::> snapmirror restore -destination-path hanala:hana_data -source
-path hana2b:backup_hana_data -source-snapshot sc-backup-
daily_20140505121000
[Job 6746] Job is queued: snapmirror restore from source
"hana2b:backup_hana_data" for the
snapshot sc-backup-daily_20140505121000.

hana::> job show -id 6746
Owning
Job ID Name Vserver Node State
-----
6746 SnapMirror restore hana hana01 Running
Description: snapmirror restore from source
"hana2b:backup_hana_data" for the snapshot sc-backup-
daily_20140505121000
```

When the restore process is finished, you use SAP HANA to perform the recovery.

Copyright information

Copyright © 2025 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.