



## **Disaster recovery testing**

NetApp solutions for SAP

NetApp  
December 09, 2025

# Table of Contents

Disaster recovery testing .....	1
Disaster Recovery Testing .....	1
Prepare the target host .....	2
Target server host name and IP address .....	2
Install required software .....	2
Configure users, ports, and SAP services .....	2
Prepare HANA log volume .....	3
Prepare log backup volume .....	3
Prepare file system mounts .....	3
Create new volumes based on snapshot backups at the disaster recovery site .....	4
Mount the new volumes at the target host .....	8
HANA database recovery .....	9
Recovery to latest HANA data volume backup savepoint .....	10
Recovery with forward recovery using log/catalog backups .....	12
Check consistency of latest log backups .....	14

# Disaster recovery testing

## Disaster Recovery Testing

To implement an effective disaster recovery strategy, you must test the required workflow. Testing demonstrates whether the strategy works and whether the internal documentation is sufficient, and it also allows administrators to train on the required procedures.

ANF Cross-Region Replication enables disaster recovery testing without putting RTO and RPO at risk. Disaster recovery testing can be done without interrupting data replication.

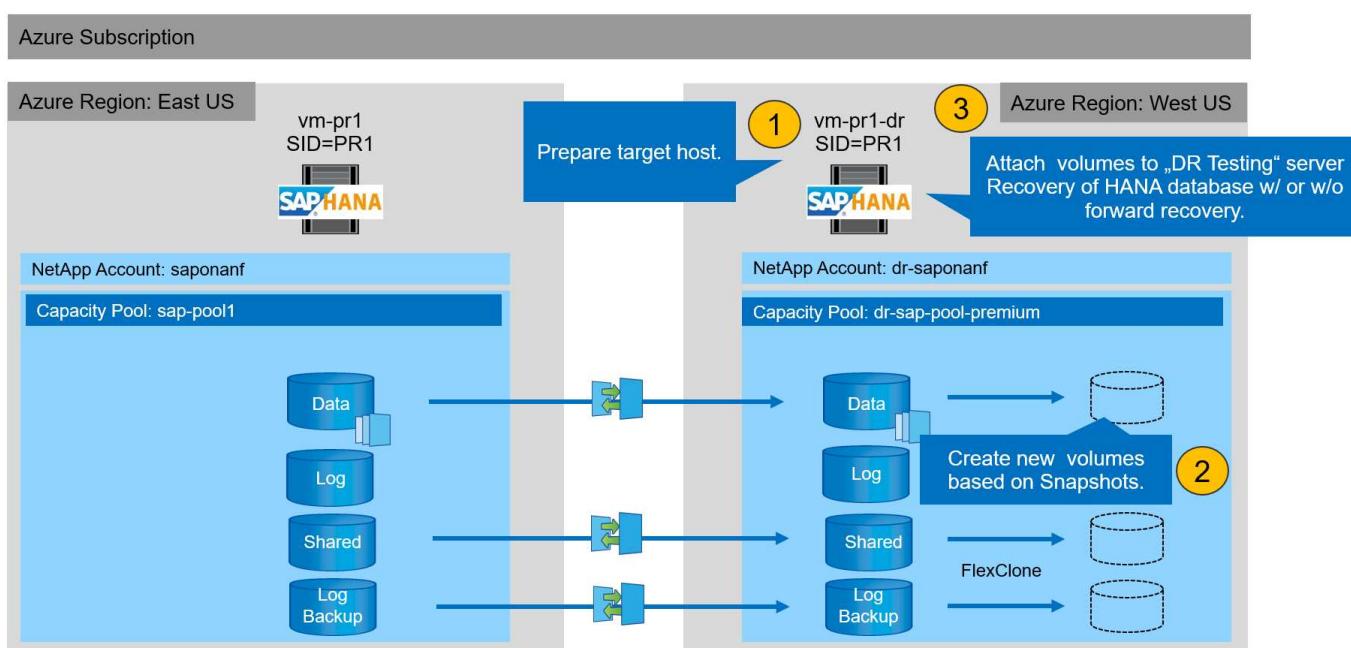
The disaster recovery testing workflow leverages the ANF feature set to create new volumes based on existing Snapshot backups at the disaster recovery target. See [How Azure NetApp Files snapshots work | Microsoft Docs](#).

Depending on whether log backup replication is part of the disaster recovery setup or not, the steps for disaster recovery are slightly different. This section describes the disaster recovery testing for data-backup-only replication as well as for data volume replication combined with log backup volume replication.

To perform disaster recovery testing, complete the following steps:

1. Prepare the target host.
2. Create new volumes based on Snapshot backups at the disaster recovery site.
3. Mount the new volumes at the target host.
4. Recover the HANA database.
  - Data volume recovery only.
  - Forward recovery using replicated log backups.

The following subsections describe these steps in detail.



# Prepare the target host

This section describes the preparation steps required at the server that is used for the disaster recovery failover.

During normal operation, the target host is typically used for other purposes, for example, as a HANA QA or test system. Therefore, most of the described steps must be executed when disaster failover testing is executed. On the other hand, the relevant configuration files, like `/etc/fstab` and `/usr/sap/sapservices`, can be prepared and then put in production by simply copying the configuration file. The disaster recovery failover procedure ensures that the relevant prepared configuration files are configured correctly.

The target host preparation also includes shutting down the HANA QA or test system as well as stopping all services using `systemctl stop sapinit`.

## Target server host name and IP address

The host name of the target server must be identical to the host name of the source system. The IP address can be different.

 Proper fencing of the target server must be established so that it cannot communicate with other systems. If proper fencing is not in place, then the cloned production system might exchange data with other production systems, resulting in logically corrupted data.

## Install required software

The SAP host agent software must be installed at the target server. For full information, see the [SAP Host Agent](#) at the SAP help portal.

 If the host is used as a HANA QA or test system, the SAP host agent software is already installed.

## Configure users, ports, and SAP services

The required users and groups for the SAP HANA database must be available at the target server. Typically, central user management is used; therefore, no configuration steps are necessary at the target server. The required ports for the HANA database must be configured at the target hosts. The configuration can be copied from the source system by copying the `/etc/services` file to the target server.

The required SAP services entries must be available at the target host. The configuration can be copied from the source system by copying the `/usr/sap/sapservices` file to the target server. The following output shows the required entries for the SAP HANA database used in the lab setup.

```
vm-pr1:~ # cat /usr/sap/sapservices
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/PR1/HDB01/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
limit.descriptors=1048576
```

## Prepare HANA log volume

Because the HANA log volume is not part of the replication, an empty log volume must exist at the target host. The log volume must include the same subdirectories as the source HANA system.

```
vm-pr1:~ # ls -al /hana/log/PR1/mnt00001/
total 16
drwxrwxrwx 5 root      root      4096 Feb 19 16:20 .
drwxr-xr-x 3 root      root      22 Feb 18 13:38 ..
drwxr-xr-- 2 pr1adm    sapsys   4096 Feb 22 10:25 hdb00001
drwxr-xr-- 2 pr1adm    sapsys   4096 Feb 22 10:25 hdb00002.00003
drwxr-xr-- 2 pr1adm    sapsys   4096 Feb 22 10:25 hdb00003.00003
vm-pr1:~ #
```

## Prepare log backup volume

Because the source system is configured with a separate volume for the HANA log backups, a log backup volume must also be available at the target host. A volume for the log backups must be configured and mounted at the target host.

If log backup volume replication is part of the disaster recovery setup, the replicated log backup volume is mounted at the target host, and it is not necessary to prepare an additional log backup volume.

## Prepare file system mounts

The following table shows the naming conventions used in the lab setup. The volume names at the disaster recovery site are included in `/etc/fstab`.

HANA PR1 volumes	Volume and subdirectories at disaster recovery site	Mount point at target host
Data volume	PR1-data-mnt00001-sm-dest	/hana/data/PR1/mnt00001
Shared volume	PR1-shared-sm-dest/shared PR1-shared-sm-dest/usr-sap-PR1	/hana/shared /usr/sap/PR1
Log backup volume	hanabackup-sm-dest	/hanabackup



The mount points from this table must be created at the target host.

Here are the required `/etc/fstab` entries.

```

vm-pr1:~ # cat /etc/fstab
# HANA ANF DB Mounts
10.0.2.4:/PR1-data-mnt0001-sm-dest /hana/data/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-log-mnt0001-dr /hana/log/PR1/mnt0001 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA ANF Shared Mounts
10.0.2.4:/PR1-shared-sm-dest/hana-shared /hana/shared nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
10.0.2.4:/PR1-shared-sm-dest/usr-sap-PR1 /usr/sap/PR1 nfs
rw,vers=4,minorversion=1,hard,timeo=600,rsize=262144,wsize=262144,intr,noa
time,lock,_netdev,sec=sys 0 0
# HANA file and log backup destination
10.0.2.4:/hanabackup-sm-dest /hanabackup nfs
rw,vers=3,hard,timeo=600,rsize=262144,wsize=262144,nconnect=8,bg,noatime,n
oclock 0 0

```

## Create new volumes based on snapshot backups at the disaster recovery site

Depending on the disaster recovery setup (with or without log backup replication), two or three new volumes based on snapshot backups must be created. In both cases, a new volume of the data and the HANA shared volume must be created.

A new volume of the log backup volume must be created if the log backup data is also replicated. In our example, data and the log backup volume have been replicated to the disaster recovery site. The following steps use the Azure Portal.

1. One of the application-consistent snapshot backups is selected as a source for the new volume of the HANA data volume. Restore to New Volume is selected to create a new volume based on the snapshot backup.

PR1-data-mnt00001-sm-dest (dr-saponanf/dr-sap-pool1/PR1-data-mnt00001-sm-dest)

PR1-data-mnt00001-sm-dest (dr-saponanf/dr-sap-pool1/PR1-data-mnt00001-sm-dest) | Snapshots

Volume

Search (Ctrl+ /) < + Add snapshot Refresh

Overview

Activity log

Access control (IAM)

Tags

Settings

Properties

Locks

Storage service

Mount instructions

Export policy

Snapshots

Replication

Monitoring

Metrics

Automation

Tasks (preview)

Export template

Support + troubleshooting

New support request

Search snapshots

Name	Location	Created	...
azacsnap_2021-02-16T134021-9431230Z	West US	02/16/2021, 02:40:27 PM	...
azacsnap_2021-02-16T134917-6284160Z	West US	02/16/2021, 02:49:20 PM	...
azacsnap_2021-02-16T135737-3778546Z	West US	02/16/2021, 02:57:41 PM	...
azacsnap_2021-02-16T160002-1354654Z	West US	02/16/2021, 05:00:05 PM	...
azacsnap_2021-02-16T200002-0790339Z	West US	02/16/2021, 09:00:08 PM	...
azacsnap_2021-02-17T000002-1753859Z	West US	02/17/2021, 01:00:06 AM	...
azacsnap_2021-02-17T040001-5454808Z	West US	02/17/2021, 05:00:05 AM	...
azacsnap_2021-02-17T080002-2933611Z	West US	02/17/2021, 09:00:18 AM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/17/2021, 12:46:22 PM	...
azacsnap_2021-02-17T120001-9196266Z	West US	02/17/2021, 01:00:08 PM	...
azacsnap_2021-02-17T160002-2801612Z	West US	02/17/2021, 05:00:06 PM	...
azacsnap_2021-02-17T200001-9149055Z	West US	02/17/2021, 09:00:05 PM	...
azacsnap_2021-02-18T000001-7955243Z	West US	02/18/2021, 01:00:07 PM	...
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 01:10:00 PM	...

Restore to new volume

Revert volume

Delete

2. The new volume name and quota must be provided in the user interface.

Home > Azure NetApp Files > dr-saponanf > dr-sap-pool1 (dr-saponanf/dr-sap-pool1) > PR1-data-mnt00001-sm-dest (d

## Create a volume

Basics Protocol Tags Review + create

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#)

### Volume details

Volume name \*

PR1-data-mnt00001-sm-dest-clone



Restoring from snapshot ⓘ

azacsnap\_2021-02-18T000001-7955243Z

Available quota (GiB) ⓘ

2096

2.05 TiB

Quota (GiB) \* ⓘ

500



500 GiB

Virtual network ⓘ

dr-vnet (10.2.0.0/16,10.0.2.0/24)



Delegated subnet ⓘ

default (10.0.2.0/28)



Show advanced section



3. Within the protocol tab, the file path and export policy are configured.

Home > Azure NetApp Files > dr-saponanf > dr-sap-pool1 (dr-saponanf/dr-sap-pool1) > PR1-data-mnt00001-sm-dest (d

## Create a volume

Basics    **Protocol**    Tags    Review + create

Configure access to your volume.

### Access

Protocol type

NFS  SMB  Dual-protocol (NFSv3 and SMB)

### Configuration

File path \* [\(i\)](#)

PR1-data-mnt00001-sm-dest-clone

Versions

NFSv4.1

Kerberos

Enabled  Disabled

### Export policy

Configure the volume's export policy. This can be edited later. [Learn more](#)

↑ Move up    ↓ Move down    ⌈ Move to top    ⌉ Move to bottom     Delete

<input checked="" type="checkbox"/> Index	Allowed clients	Access	Root Access
---	-----------------	--------	-------------

<input checked="" type="checkbox"/> 1	0.0.0.0/0	Read & Write	On

4. The Create and Review screen summarizes the configuration.

## Create a volume

Validation passed

Basics    Protocol    Tags    Review + create

### Basics

Subscription	Pay-As-You-Go
Resource group	dr-rg-sap
Region	West US
Volume name	PR1-data-mnt00001-sm-dest-clone
Capacity pool	dr-sap-pool1
Service level	Standard
Quota	500 GiB

### Networking

Virtual network	dr-vnet (10.2.0.0/16,10.0.2.0/24)
Delegated subnet	default (10.0.2.0/28)

### Protocol

Protocol	NFSv4.1
File path	PR1-data-mnt00001-sm-dest-clone

5. A new volume has now been created based on the HANA snapshot backup.

dr-saponanf | Volumes

NetApp account

Search (Ctrl+ /) < Add volume Add data replication Refresh

Overview    Activity log    Access control (IAM)    Tags    Settings    Properties    Locks

Azure NetApp Files    Active Directory connections

Storage service    Capacity pools    Volumes

Data protection    Snapshot policies

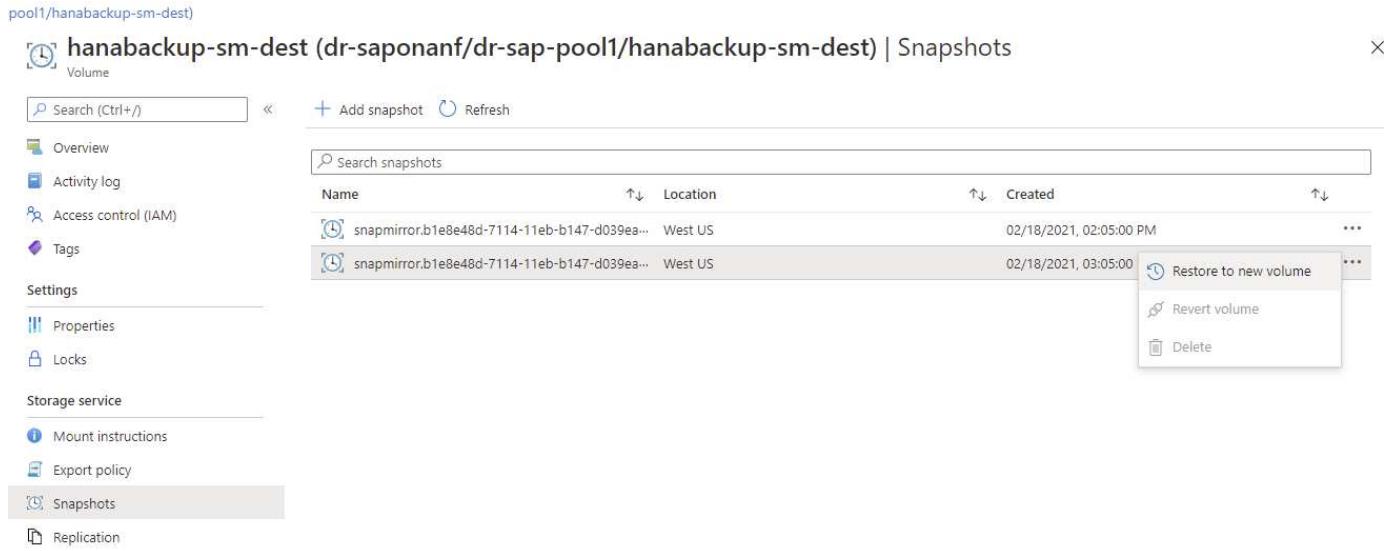
Automation    Tasks (preview)    Export template

Support + troubleshooting    New support request

Search volumes

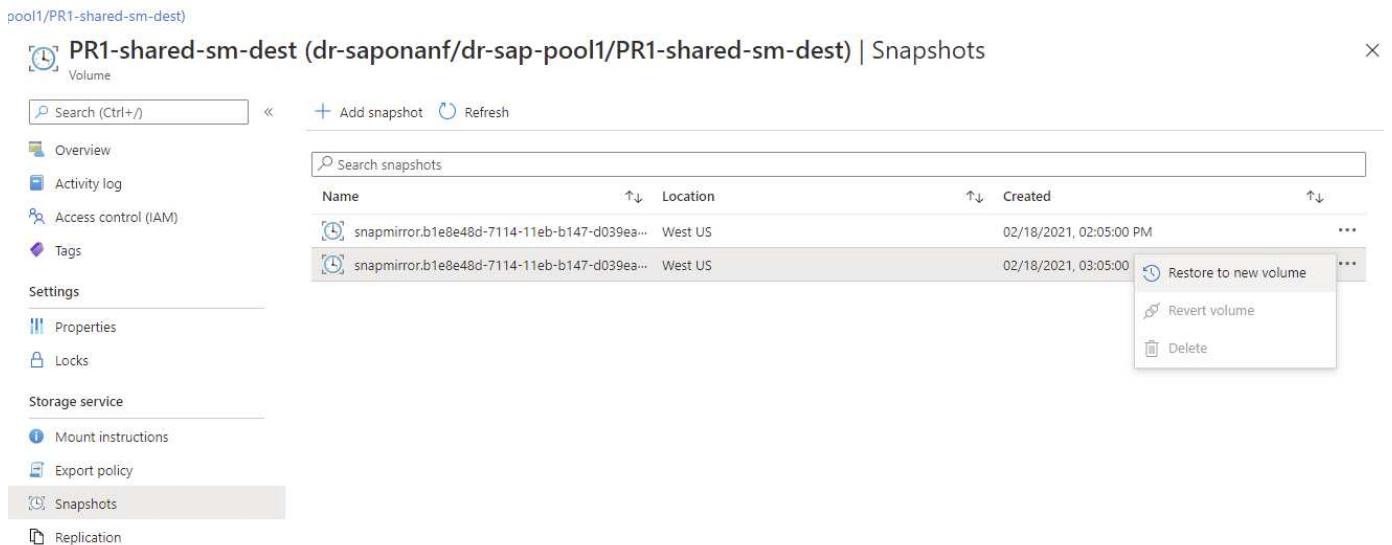
Name	Quota	Protocol type	Mount path	Service level	Capacity pool
hanabackup-sm-dest	1000 GiB	NFSv3	10.0.2.4:/hanabackup-sm-dest	Standard	dr-sap-pool1
PR1-data-mnt00001-sm-dest	500 GiB	NFSv4.1	10.0.2.4:/PR1-data-mnt00001-sm-dest	Standard	dr-sap-pool1
PR1-data-mnt00001-sm-dest-clone	500 GiB	NFSv4.1	10.0.2.4:/PR1-data-mnt00001-sm-dest-clone	Standard	dr-sap-pool1
PR1-log-mnt00001-dr	250 GiB	NFSv4.1	10.0.2.4:/PR1-log-mnt00001-dr	Standard	dr-sap-pool1
PR1-shared-sm-dest	250 GiB	NFSv4.1	10.0.2.4:/PR1-shared-sm-dest	Standard	dr-sap-pool1

The same steps must now be performed for the HANA shared and the log backup volume as shown in the following two screenshots. Since no additional snapshots have been created for the HANA shared and log backup volume, the newest SnapMirror Snapshot copy must be selected as the source for the new volume. This is unstructured data, and the SnapMirror Snapshot copy can be used for this use case.



Name	Location	Created
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 02:05:00 PM
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 03:05:00

The following screenshot shows the HANA shared volume restored to new volume.



Name	Location	Created
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 02:05:00 PM
snapmirror.b1e8e48d-7114-11eb-b147-d039ea...	West US	02/18/2021, 03:05:00



If a capacity pool with a low performance tier has been used, the volumes must now be moved to a capacity pool that provides the required performance.

All three new volumes are now available and can be mounted at the target host.

## Mount the new volumes at the target host

The new volumes can now be mounted at the target host, based on the `/etc/fstab` file created before.

```
vm-pr1:~ # mount -a
```

The following output shows the required file systems.

```
vm-pr1:/hana/data/PR1/mnt00001/hdb00001 # df
Filesystem                                1K-blocks      Used
Available  Use% Mounted on
devtmpfs                                     8190344        8
8190336  1% /dev
tmpfs                                         12313116      0
12313116  0% /dev/shm
tmpfs                                         8208744      17292
8191452  1% /run
tmpfs                                         8208744      0
8208744  0% /sys/fs/cgroup
/dev/sda4                                     29866736  2438052
27428684  9% /
/dev/sda3                                     1038336      101520
936816  10% /boot
/dev/sda2                                     524008      1072
522936  1% /boot/efi
/dev/sdb1                                     32894736      49176
31151560  1% /mnt
tmpfs                                         1641748      0
1641748  0% /run/user/0
10.0.2.4:/PR1-log-mnt00001-dr           107374182400      256
107374182144  1% /hana/log/PR1/mnt00001
10.0.2.4:/PR1-data-mnt00001-sm-dest-clone 107377026560  6672640
107370353920  1% /hana/data/PR1/mnt00001
10.0.2.4:/PR1-shared-sm-dest-clone/hana-shared 107377048320 11204096
107365844224  1% /hana/shared
10.0.2.4:/PR1-shared-sm-dest-clone/usr-sap-PR1 107377048320 11204096
107365844224  1% /usr/sap/PR1
10.0.2.4:/hanabackup-sm-dest-clone          107379429120  35293440
107344135680  1% /hanabackup
```

## HANA database recovery

The following shows the steps for HANA database recovery

Start the required SAP services.

```
vm-pr1:~ # systemctl start sapinit
```

The following output shows the required processes.

```
vm-pr1:/ # ps -ef | grep sap
root      23101      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saphostexec pf=/usr/sap/hostctrl/exe/host_profile
pr1adm    23191      1  3 11:29 ?          00:00:00
/usr/sap/PR1/HDB01/exe/sapstartsrv
pf=/usr/sap/PR1/SYS/profile/PR1_HDB01_vm-pr1 -D -u pr1adm
sapadm    23202      1  5 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile -D
root      23292      1  0 11:29 ?          00:00:00
/usr/sap/hostctrl/exe/saposcol -l -w60
pf=/usr/sap/hostctrl/exe/host_profile
root      23359  2597  0 11:29 pts/1    00:00:00 grep --color=auto sap
```

The following subsections describe the recovery process with and without forward recovery using the replicated log backups. The recovery is executed using the HANA recovery script for the system database and hdbsql commands for the tenant database.

## Recovery to latest HANA data volume backup savepoint

The recovery to the latest backup savepoint is executed with the following commands as user pr1adm:

- System database

```
recoverSys.py --command "RECOVER DATA USING SNAPSHOT CLEAR LOG"
```

- Tenant database

```
Within hdbsql: RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
```

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

### System database recovery

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py
--command="RECOVER DATA USING SNAPSHOT CLEAR LOG"
[139702869464896, 0.008] >> starting recoverSys (at Fri Feb 19 14:32:16
2021)
[139702869464896, 0.008] args: ()
[139702869464896, 0.009] keys: {'command': 'RECOVER DATA USING SNAPSHOT
CLEAR LOG'}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 14:32:16 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 14:32:16
stopped system: 2021-02-19 14:32:16
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 14:32:21
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T14:32:56+00:00  P0027646      177bab4d610  INFO      RECOVERY
RECOVER DATA finished successfully
recoverSys finished successfully: 2021-02-19 14:32:58
[139702869464896, 42.017] 0
[139702869464896, 42.017] << ending recoverSys, rc = 0 (RC_TEST_OK), after
42.009 secs
pr1adm@vm-pr1:/usr/sap/PR1/HDB01>

```

## Tenant database recovery

If a user store key has not been created for the pr1adm user at the source system, a key must be created at the target system. The database user configured in the key must have privileges to execute tenant recovery operations.

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbuserstore set PR1KEY vm-pr1:30113
<backup-user> <password>

```

The tenant recovery is now executed with hdbsql.

```
pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type:  \h for help with commands
      \q to quit
hdbsql SYSTEMDB=> RECOVER DATA FOR PR1 USING SNAPSHOT CLEAR LOG
0 rows affected (overall time 66.973089 sec; server time 66.970736 sec)
hdbsql SYSTEMDB=>
```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

## Recovery with forward recovery using log/catalog backups

Log backups and the HANA backup catalog are being replicated from the source system.

The recovery using all available log backups is executed with the following commands as user pr1adm:

- System database

```
recoverSys.py --command "RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT"
```

- Tenant database

```
Within hdbsql: RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
```



To recover using all available logs, you can just use any time in the future as the timestamp in the recovery statement.

You can also use HANA Studio or Cockpit to execute the recovery of the system and the tenant database.

The following command output show the recovery execution.

### System database recovery

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> HDBSettings.sh recoverSys.py --command
"RECOVER DATABASE UNTIL TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING
SNAPSHOT"
[140404915394368, 0.008] >> starting recoverSys (at Fri Feb 19 16:06:40
2021)
[140404915394368, 0.008] args: ()
[140404915394368, 0.008] keys: {'command': "RECOVER DATABASE UNTIL
TIMESTAMP '2021-02-20 00:00:00' CLEAR LOG USING SNAPSHOT"}
using logfile /usr/sap/PR1/HDB01/vm-pr1/trace/backup.log
recoverSys started: =====2021-02-19 16:06:40 =====
testing master: vm-pr1
vm-pr1 is master
shutdown database, timeout is 120
stop system
stop system on: vm-pr1
stopping system: 2021-02-19 16:06:40
stopped system: 2021-02-19 16:06:41
creating file recoverInstance.sql
restart database
restart master nameserver: 2021-02-19 16:06:46
start system: vm-pr1
sapcontrol parameter: ['-function', 'Start']
sapcontrol returned successfully:
2021-02-19T16:07:19+00:00  P0009897      177bb0b4416 INFO      RECOVERY
RECOVER DATA finished successfully, reached timestamp 2021-02-
19T15:17:33+00:00, reached log position 38272960
recoverSys finished successfully: 2021-02-19 16:07:20
[140404915394368, 39.757] 0
[140404915394368, 39.758] << ending recoverSys, rc = 0 (RC_TEST_OK), after
39.749 secs

```

## Tenant database recovery

```

pr1adm@vm-pr1:/usr/sap/PR1/HDB01> hdbsql -U PR1KEY
Welcome to the SAP HANA Database interactive terminal.
Type: \h for help with commands
      \q to quit

hdbsql SYSTEMDB=> RECOVER DATABASE FOR PR1 UNTIL TIMESTAMP '2021-02-20
00:00:00' CLEAR LOG USING SNAPSHOT
0 rows affected (overall time 63.791121 sec; server time 63.788754 sec)

hdbsql SYSTEMDB=>

```

The HANA database is now up and running, and the disaster recovery workflow for the HANA database has been tested.

## Check consistency of latest log backups

Because log backup volume replication is performed independently of the log backup process executed by the SAP HANA database, there might be open, inconsistent log backup files at the disaster recovery site. Only the latest log backup files might be inconsistent, and those files should be checked before a forward recovery is performed at the disaster recovery site using the `hdbbackupcheck` tool.

If the `hdbbackupcheck` tool reports an error for the latest log backups, the latest set of log backups must be removed or deleted.

```
pr1adm@hana-10: > hdbbackupcheck
/hanabackup/PR1/log/SYSTEMDB/log_backup_0_0_0_0.1589289811148
Loaded library 'libhdbcaccessor'
Loaded library 'libhdblivecach'
Backup '/mnt/log-backup/SYSTEMDB/log_backup_0_0_0_0.1589289811148'
successfully checked.
```

The check must be executed for the latest log backup files of the system and the tenant database.

If the `hdbbackupcheck` tool reports an error for the latest log backups, the latest set of log backups must be removed or deleted.

## 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.