

---

# SAP HANA DBA's Guide to Actifio GO

The logo for Actifio GO is located in the bottom right corner of the page. It consists of a blue rectangular background. On the left side of this rectangle, there is a pattern of overlapping hexagons in a lighter shade of blue. To the right of this pattern, the text "Actifio GO" is written in a white, sans-serif font.

**Actifio GO**

**Copyright, Trademarks, and other Legal Matter**

Copyright © 2021 Google LLC. All rights reserved.

Actifio™, OnVault™, and VDP™ are trademarks of Google LLC.

All other brands, product names, goods and/or services mentioned herein are trademarks or property of their respective owners.

# Contents

Preface .....	v
The ActifioNOW Customer Portal.....	v
Actifio Support Centers.....	v
Chapter 1 – SAP HANA DBA’s Introduction to Actifio Copy Data Management .....	1
Actifio Data Virtualization.....	1
Capturing Data.....	2
Replicating Data.....	2
Accessing Data.....	3
Introduction to Actifio SAP HANA Administration.....	4
SAP HANA Backup Methods.....	5
References.....	6
Chapter 2 – Preparing the SAP HANA 1.0 Database .....	7
Creating the Database User Account.....	7
Get the SQL Port ID.....	9
Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 (single container system).....	9
Chapter 3 – Preparing a HANA 2.0 Database .....	11
Creating the System Database and Tenant Database Users.....	11
Creating the System Database User Account from HANA STUDIO SPS01.....	11
Creating the System Database User Account from HANA STUDIO, SPS02, 03, & 04.....	13
Creating the System Database User Account from HANA STUDIO, SPS05.....	14
Creating the User under the Tenant DB.....	16
Getting the Instance and SQL Port Numbers.....	17
Creating the SAP HANA Hdbuserstore Key.....	17
Creating the SAP HANA Hdbuserstore Key for the System Database and Each Tenant Database in a Single Node System.....	18
Creating the SAP HANA Hdbuserstore Key for the System Database and each Tenant Database in a Scale-Out Multi-Node SAP HANA System.....	19
Using the Tenant DB User Store Key Prefix.....	20
Chapter 4 – Adding a SAP HANA Database Host and Discovering Databases .....	21
Adding the Host to AGM.....	21

Onboarding the HANA Databases from the App Manager.....	22
Chapter 5 - Defining Actifio GO Policy Templates and Resource Profiles .....	23
Chapter 6 - Protecting the HANA Database and Its Logs .....	25
Protecting the HANA Database.....	25
Protecting HANA Database Logs.....	27
Setting up the Log Mode and Log Backup in HANA Studio.....	27
Setting up the Log Backup in Actifio AGM.....	29
Chapter 7 - Restoring or Recovering an SAP HANA Database .....	31
Instant Recovery of a HANA Database from a Volume-Based Snapshot to the Source .....	31
Clean-up Procedure After a Failed Migrate Job.....	33
Instant Recovery of a HANA Database from a Block-Based LVM Snapshot to a New Target.....	34
Recovering a Single Tenant Database from a Volume-Based Snapshot .....	36
Restoring a HANA Database from a Volume-Based Snapshot Back to the Source .....	37
Restoring a HANA Database from a Full+Incremental Snapshot Back to the Source .....	38
Recovering from a Full+Incremental Snapshot to a New Target.....	39
Chapter 8 - Accessing an SAP HANA Database .....	41
Mount a Virtual SAP HANA Database from a Volume-Based Snapshot to a Target SAP HANA Host.....	41
Workflow to Automate Mount and Refresh of a Virtual Database.....	43
Chapter 9 - HANA Database Management Using actHANADBМ .....	45
Installing and Configuring actHANADBМ.pl.....	46
actHANADBМ Commands .....	47
agmconfig.....	47
createTemplate.....	47
hostDiscovery.....	48
protectApp.....	49
backup.....	49
listImageDetails.....	50
mount.....	51
unmountdelete.....	52
restore.....	53
runwf.....	54

---

# Preface

---

The information presented in this guide is intended for users who are familiar with basic Actifio processes and procedures as described in **Getting Started with Actifio Copy Data Management** and who are qualified to administer SAP HANA databases.

## The ActifioNOW Customer Portal

During the configuration and initialization of your Actifio appliance your Actifio representative provided you with a user name and password for the ActifioNOW customer portal.

From the customer portal you can obtain detailed reports about your Actifio appliance as well as search the portal's knowledge base for answers to specific questions.

To log into the ActifioNOW customer portal:

1. Go to: <https://now.actifio.com>
2. When prompted, enter the user name and password provided by your Actifio representative.

## Actifio Support Centers

To contact an Actifio support representative, you can:

- Send email to: [support@actifio.com](mailto:support@actifio.com)
- Call:

**From anywhere:** +1.315.261.7501

**US Toll-Free:** +1.855.392.6810

**Australia:** 0011 800-16165656

**Germany:** 00 800-16165656

**New Zealand:** 00 800-16165656

**UK:** 0 800-0155019



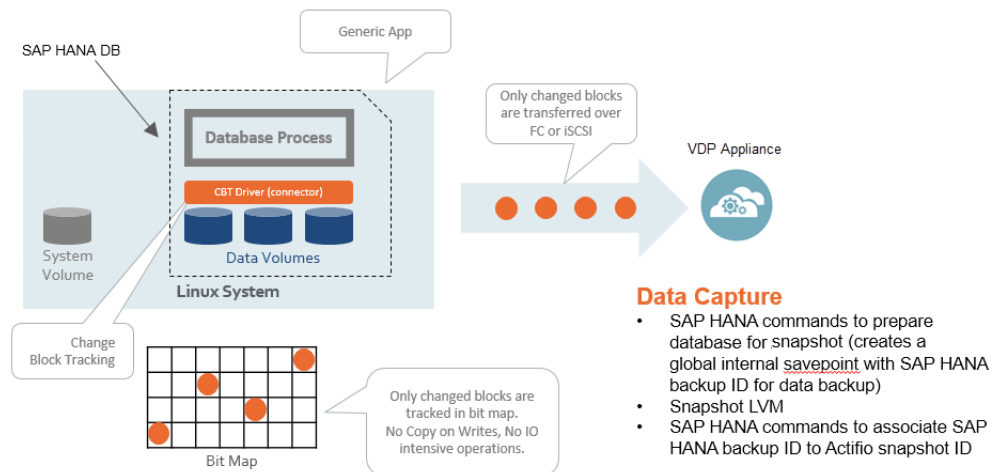
# 1 SAP HANA DBA's Introduction to Actifio Copy Data Management

This chapter introduces Actifio concepts and the procedures used to capture and access databases. It includes:

- [Actifio Data Virtualization](#) on page 1
- [Capturing Data](#) on page 2
- [Replicating Data](#) on page 2
- [Accessing Data](#) on page 3
- [Introduction to Actifio SAP HANA Administration](#) on page 4
- [SAP HANA Backup Methods](#) on page 5
- [References](#) on page 6

## Actifio Data Virtualization

An Actifio Appliance is a highly scalable copy data management platform that virtualizes application data to improve the resiliency, agility, and cloud mobility of your business. It works by virtualizing data in much the same way other technologies have virtualized servers and networks. This enables you to capture data from production systems, manage it in the most efficient way possible, and use virtual copies of the data however they are needed.



### SAP HANA for LVM with Linux Change Block Tracking

Application data is captured at the block level, in application native format, according to an SLA. A Golden copy of that data is created, and is then updated incrementally with only the changed blocks of data in an "incremental forever" model. Unlimited virtual copies of the data can be made available instantly for use, without proliferating physical copies and taking up additional storage infrastructure.

## Capturing Data

Capturing data consists of four steps:

1. Add servers that host databases.
2. Discover the database.
3. Define Actifio Policy Templates and Resource Profiles according to your RPOs and RTOs.
4. Assign Actifio Policy Templates and Resource Profiles to discovered databases.

## The Actifio Connector

The Actifio Connector is used to capture selected databases. The Actifio Connector is a small-footprint, lightweight service that can be installed on either virtual or physical servers.

Specifically, the Actifio Connector:

- Discovers the application to which data and log volumes will be added.
- Uses Linux changed block tracking to capture data at block level in incremental forever fashion.
- Identifies changes to database data for Actifio's incremental forever capture strategy.

## Replicating Data

Data can be replicated to a second Actifio Appliance or to the cloud for recovery, disaster recovery, or test/development purposes.

Data replication has traditionally been an inhibitor to efficient data management in a geographically distributed environment. Actifio replication addresses these issues with a global deduplication and compression approach that:

- Drives down overall network usage.
- Eliminates the need for a dedicated WAN accelerator/optimizer.
- Does not require storage array vendor licenses as data is sent from one Actifio Appliance to another.
- Is heterogeneous from any supported array to any supported array: Tier 1 to Tier 2 and/or Vendor A to Vendor B.
- Preserves write-order, even across multiple LUNs.
- Is fully integrated with VMware Site Recovery Manager (SRM) and Actifio Resiliency Director.

Actifio Replication is controlled by Actifio Policy Template policies:

- Production to Mirror policies have several options to replicate data to a second Actifio Appliance.
- Dedup Backup to Dedup DR policies use a fixed, Actifio proprietary replication engine to replicate data to a second Actifio Appliance. In addition, Dedup Backup to Dedup DR policies allow you to replicate data to two locations.
- Production to Vault policies use a fixed, Actifio proprietary replication engine to replicate data to the cloud.



# Accessing Data

Actifio VDP can instantly present a copy of the database rolled forward to a specific point of time. Access options include:

- Mounts
- LiveClones
- Restores
- Workflows

## Mounts

The Actifio mount function provides instant access to data without moving data. Captured copies of databases can be rolled forward via the Actifio user interface and mounted on any database server. Mounts are described in [Chapter 7, Restoring or Recovering an SAP HANA Database](#).

## LiveClones

The LiveClone is an independent copy of data that can be refreshed when the source data changes. The advantage of LiveClones is that they are independent copies of data that can be incrementally refreshed and masked before being made available to users. This allows teams such as development and test to ensure they are working on the latest set of data without having to manually manage the data and not access or interfere with the production environment.

## Restores

The restore function reverts the production data to a specified point in time. Restore operations actually move data. Typically restore operations are performed to restore a database to a valid state after a massive data corruption or storage array failure. The amount of time required to complete a restore operation depends on the amount of data involved. Restores are described in [Chapter 7, Restoring or Recovering an SAP HANA Database](#).

## Workflows

Workflows automate access to the captured database. Workflows are built with captured data. Workflows can present data as either a direct mount or as a LiveClone:

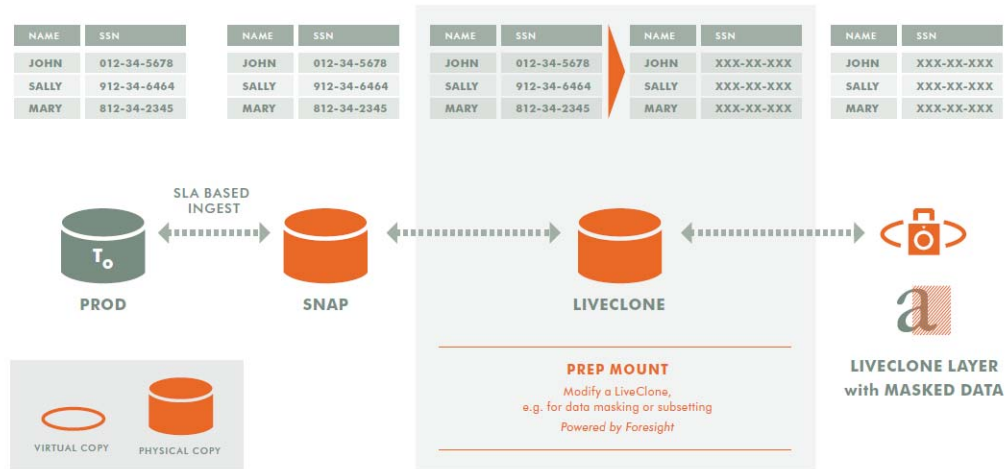
- Direct mounts (standard or virtual applications) work well for data that does not need to be masked prior to being presented. A mounted copy of data can be refreshed manually or on automatically on a schedule. Direct mounts allow you to instantly access captured data without actually moving the data.
- A LiveClone is a copy of your production data that can be updated manually or on a scheduled basis. You can mask sensitive data in a LiveClone prior to making it available to users.

Combining Actifio automated data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Users can provision their own environments almost instantly.

For example, you can create an SLA Template Policy that captures data according to a schedule. You can mark the captured data as sensitive and only accessible by users with the proper access rights. After access rights have been defined and data has been captured, you can create a Workflow that:

- Makes the captured data available as a LiveClone or as a direct mount
- Updates the LiveClone or mountable data on a scheduled or on-demand basis
- (Optional) Automatically applies scripts to the LiveClone's data after each update. This is useful for masking sensitive data.

Once the Workflow completes, users can provision environments with LiveClone or mounted data.



**Workflow With Masked Social Security Data**

## Introduction to Actifio SAP HANA Administration

Actifio can virtualize and protect:

- **Single Container system (HANA 1.0 or 2.0) Dedicated:** In single-container system the system database and tenant database are seen as a single unit and are administered as one.
- **MDC: Multiple-Container Systems (HANA 2.0):** Multiple isolated databases in a single SAP HANA system. These are referred to as multi-tenant database containers. A multiple-container system always has exactly one system database used for central system administration, and any number of multi-tenant databases (including zero), also called tenant databases.

**Table 1: Actifio Support for SAP HANA Configurations**

Configurations	SAP Storage Snapshot API	SAP File-Based API (hdbsql): Actifio Block Mapping	SAP File-Based API (hdbsql): Actifio NFS Mapping
Single Container System (HANA 1.0)	Yes (preferred)	Yes	Yes
MDC: Multiple-Container Systems (HANA 2.0) with one tenant database	Yes (preferred)	Yes	Yes
MDC: Multiple-Container Systems (HANA 2.0) with more than one tenant database	SAP HANA platform 2.0 SPS 04	Yes	Yes
Scale-Out MDC: Multiple-Container Systems (HANA 2.0) with one or more tenant databases			Yes
HANA 2.0 1+1 HA system replication	SAP HANA platform 2.0 SPS 04	Yes	Yes

### Notes

- SAP storage snapshot API - leverages Actifio CBT with incremental-forever and instant mount.
- SAP file-based API - traditional backup with weekly full, daily incremental & copy-based restore.
- NFS mapping is always to all HANA nodes.
- HANA log backup is automatic in all options and integrated with database backup policies.

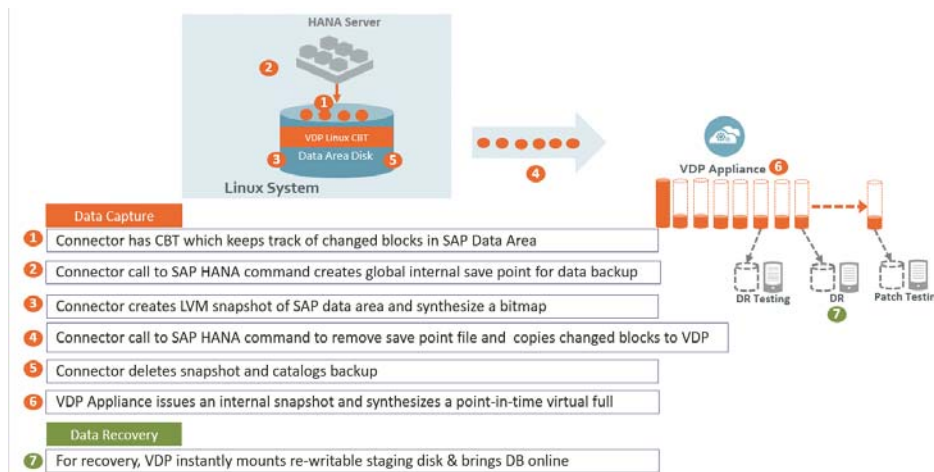
# SAP HANA Backup Methods

Actifio offers these methods of protecting SAP HANA databases:

- Block-Based LVM Snapshot with CBT Integrated with SAP HANA Database Storage Snapshot API
- File-Based Backup Integrated with HANA Traditional Backup API
- SAP HANA Log Backup

## Block-Based LVM Snapshot with CBT Integrated with SAP HANA Database Storage Snapshot API

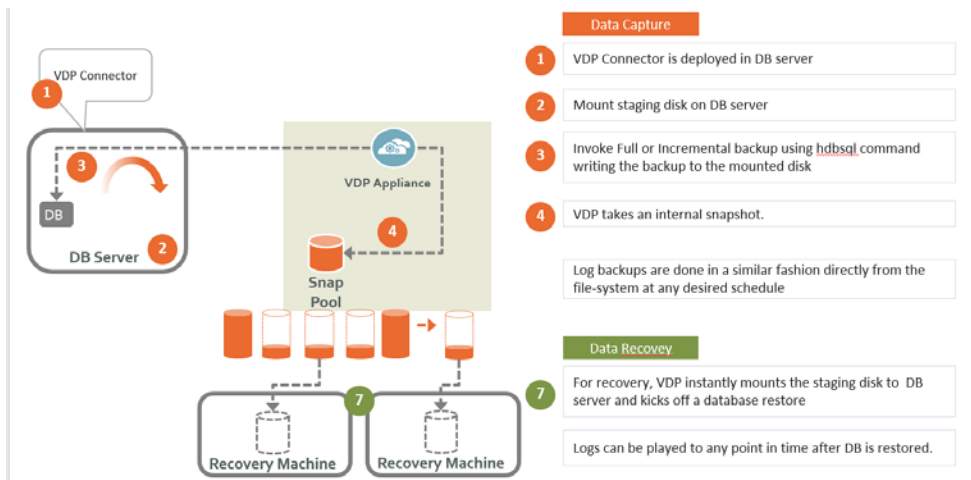
The SAP HANA database creates a database internal snapshot based on a system wide save point executed during the PREPARE step. The database internal snapshot is stored in the data volumes area.



### How it Works Using HANA Storage Snapshot API with Linux CBT and LVM Snapshot

## File-Based Backup Integrated with HANA Traditional Backup API

This provides the full and incremental backups of the data area, which is in backup format. The recovery API recovers the data area by overwriting the data area. When the data area is backed up, the entire payload data from all server nodes of the SAP HANA database instance is backed up. This applies in both single-host and multi-host environments.



### How it Works Using HANA File-Based (hdbsql API) Traditional Backup

## SAP HANA Log Backup

Log backups start automatically if the parameters `enable_auto_log_backup` and `log_mode = normal` have been configured. During a log backup, the payload of the log segments is copied from the log area to the location specified by the parameter `basepath_logbackup`.

## References

1. Category > Administration Guide: [http://help.sap.com/hana\\_platform](http://help.sap.com/hana_platform)
2. Storage Snapshots: [https://help.sap.com/saphelp\\_hanaplatform/helpdata/en/ac/114d4b34d542b99bc390b34f8ef375/content.htm](https://help.sap.com/saphelp_hanaplatform/helpdata/en/ac/114d4b34d542b99bc390b34f8ef375/content.htm)
3. 1642148 - FAQ: SAP HANA Database Backup & Recovery:  
<https://launchpad.support.sap.com/#/notes/1642148/E>
4. Create a homogeneous copy of an SAP HANA database by recovering an existing database to a different database:  
[https://help.sap.com/saphelp\\_hanaplatform/helpdata/en/ea/70213a0e114ec29724e4a10b6bb176/content.htm?frameset=/en/ca/c903c28b0e4301b39814ef41dbf568/frameset.htm&current\\_toc=/en/00/0ca1e3486640ef8b884cdf1a050fbb/plain.htm&node\\_id=773&show\\_children=false](https://help.sap.com/saphelp_hanaplatform/helpdata/en/ea/70213a0e114ec29724e4a10b6bb176/content.htm?frameset=/en/ca/c903c28b0e4301b39814ef41dbf568/frameset.htm&current_toc=/en/00/0ca1e3486640ef8b884cdf1a050fbb/plain.htm&node_id=773&show_children=false)

# 2 Preparing the SAP HANA 1.0 Database

---

**Note:** To prepare a HANA 2.0 database, see [Chapter 3, Preparing a HANA 2.0 Database](#).

---

## Prerequisites

- All the configured services (see SAP Note 1697613 and SAP Note 1649519) such as nameserver, indexserver, etc must be running. You can check this in the Overview of SAP HANA studio -> Operational State: All Services are started.
- Make sure log\_mode for database is set to normal (check under HANA Studio configuration tab).
- Use a SAP HANA hdbuserstore key to execute Backup and Recovery instead of a user name and password to communicate with HANA database using the SAP HANA Secure User Store. For HANA 1.0, a userstore key must be created for a single container under the database.

Preparing the HANA 1.0 database requires:

[Creating the Database User Account](#) on page 7

[Get the SQL Port ID](#) on page 9

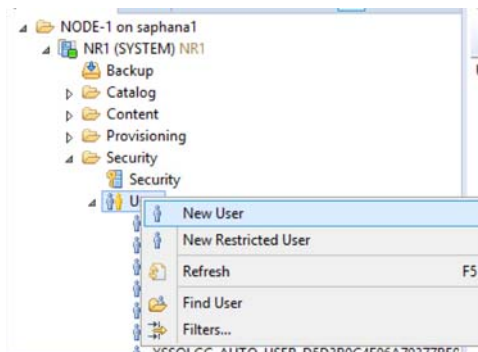
[Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 \(single container system\)](#) on page 9

## Creating the Database User Account

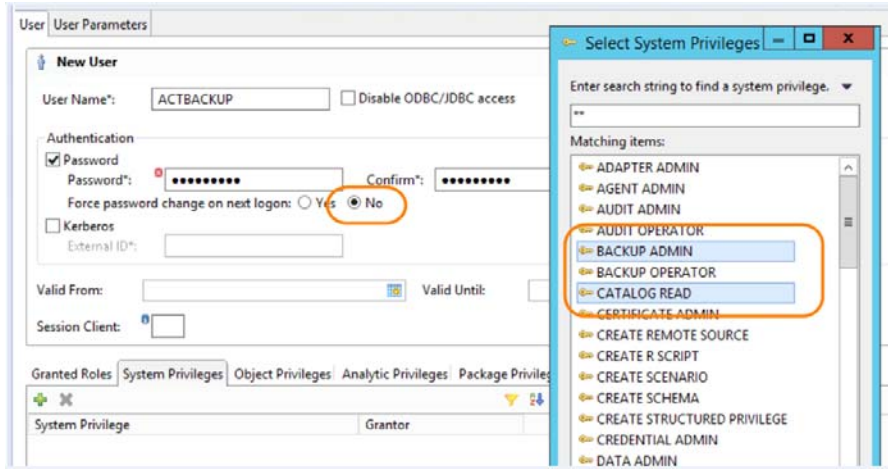
Make sure to create this user account under a single container database. Make sure to provide BACKUP ADMIN and CATALOG READ to back up the user created under database. Choose a database user name based on company's standard.

To create the user:

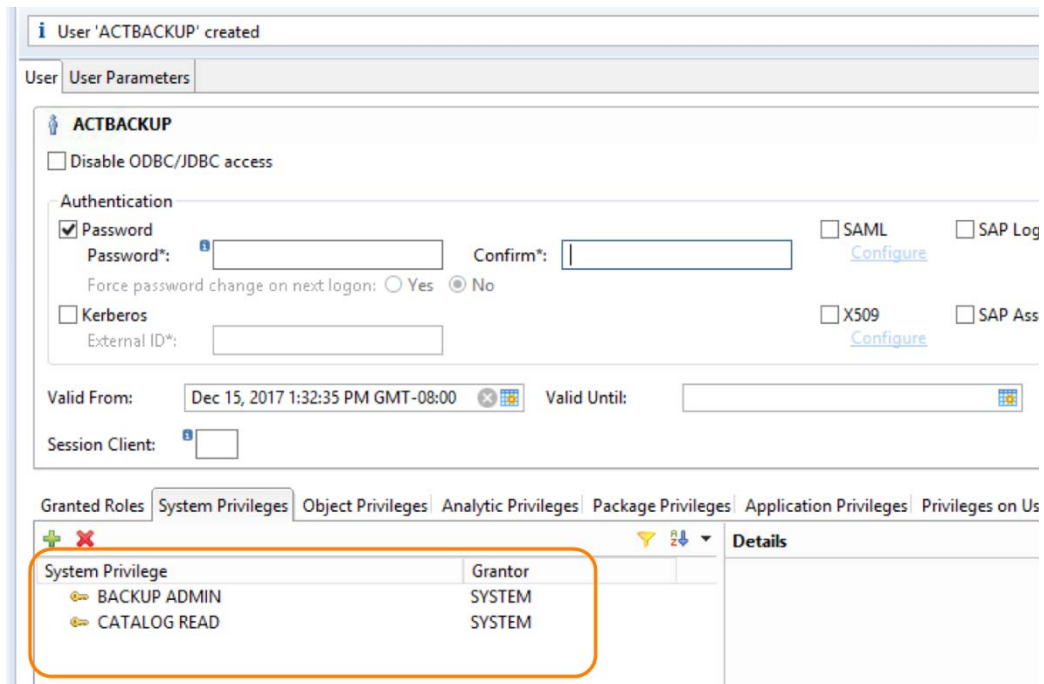
1. From SAP HANA Studio go to System > Security > Users > New User.



2. Assign a user name and a password.
3. Set Force password change on next logon to **No**.
4. Click on the System Privilege tab and assign privilege by selecting **BACKUP ADMIN** and **CATALOG READ**.

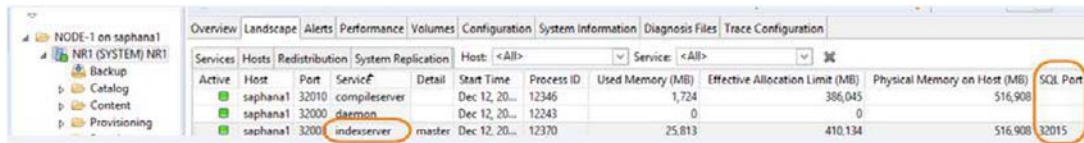


You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN and CATALOG READ privileges.



## Get the SQL Port ID

For a HANA 1.0 single container system, get the SQL PORT from HANA Studio. At System > Landscape, get the value of SQL Port for indexserver. In the example below, 32015 is the SQL port, and the instance number here is 20.



Active	Host	Port	Service	Detail	Start Time	Process ID	Used Memory (MB)	Effective Allocation Limit (MB)	Physical Memory on Host (MB)	SQL Port
Active	saphana1	32010	compileserver		Dec 12, 20...	12345	1,724	386,045	516,908	
Active	saphana1	32000	daemon		Dec 12, 20...	12243	0	0		
Active	saphana1	32000	indexserver	master	Dec 12, 20...	12370	25,813	410,134	516,908	32015

## Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 (single container system)

To communicate with HANA database, use a SAP HANA hdbuserstore key instead of a user name and password. Create the hdbuserstore key using the SAP HANA Secure User Store.

### Hdbuserstore Key Naming Convention

Set the key name = DATABASE BACKUP USERNAME.

For example:

DATABASE BACKUP USERNAME = ACTBACKUP

Set SYSTEMDB key name = ACTBACKUP

### Procedure

To create the SAP HANA hdbuserstore key:

1. Open the putty window to the HANA database server and login to <sid>adm by su to <sid>adm.
2. `cd exe`
3. Create entries in the hdbuserstore by calling:

```
# ./hdbuserstore SET <key_name> <server>:<port> <DB_user_name> <DB_user_password>
```

The <port> is the SQL port of the systemdb or tenant database, see above.

For example:

- DATABASE Backup username from above: ACTBACKUP
- KEY NAME: ACTBACKUP (same as database backup username)
- SQL Port from above: 32013
- Hostname : saphana3

```
./hdbuserstore SET ACTBACKUP saphana3:32013 ACTBACKUP <database backup user password>  
*****>
```

4. Check the keystore: `./hdbuserstore list`.





---

# 3 Preparing a HANA 2.0 Database

---

**Note:** To prepare a HANA 1.0 database, see [Chapter 2, Preparing the SAP HANA 1.0 Database](#).

---

## Prerequisites

- All the configured services (see SAP Note 1697613 and SAP Note 1649519) such as nameserver, indexserver, etc. must be running. You can check this in the Overview of SAP HANA studio -> Operational State: All Services are started.
- Make sure log\_mode for database is set to normal (check under the HANA Studio configuration tab).
- Use a SAP HANA hdbuserstore key to execute Backup and Recovery instead of a user name and password to communicate with HANA database using the SAP HANA Secure User Store. For HANA 2.0 userstore key needs to be created for SYSTEMDB and all tenant db.
- Create the database user account and hdbuserstore key names in accordance with the company's naming convention. Make sure to create this user account under SYSTEMDB and all tenant databases.

This includes:

[Creating the System Database and Tenant Database Users](#) on page 11

[Getting the Instance and SQL Port Numbers](#) on page 17

[Creating the SAP HANA Hdbuserstore Key](#) on page 17

## Creating the System Database and Tenant Database Users

[Creating the System Database User Account from HANA STUDIO SPS01](#) on page 11

[Creating the System Database User Account from HANA STUDIO, SPS02, 03, & 04](#) on page 13

[Creating the System Database User Account from HANA STUDIO, SPS05](#) on page 14

[Creating the User under the Tenant DB](#) on page 16

## Creating the System Database User Account from HANA STUDIO SPS01

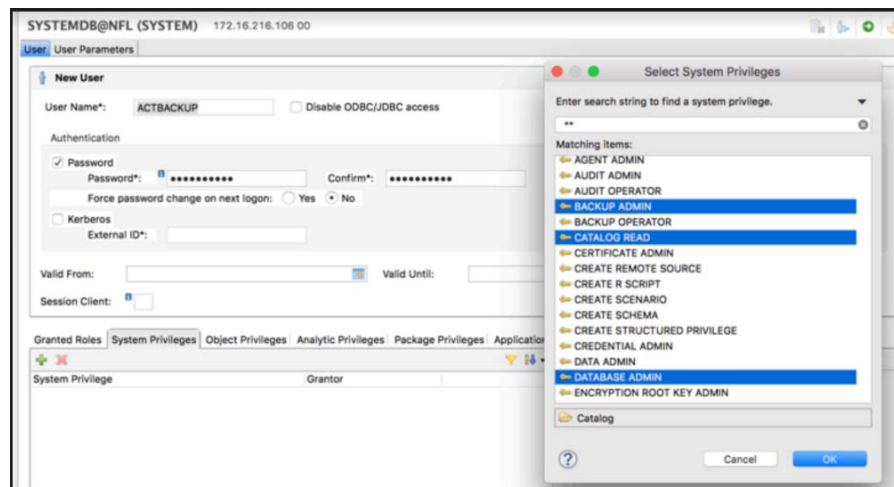
### Naming convention for database user account

Choose a database user name based on company's standard. Make sure to create this user account under SYSTEMDB. Make sure to provide BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN to the backup user created under database.


## Procedure

To create the system database user account:

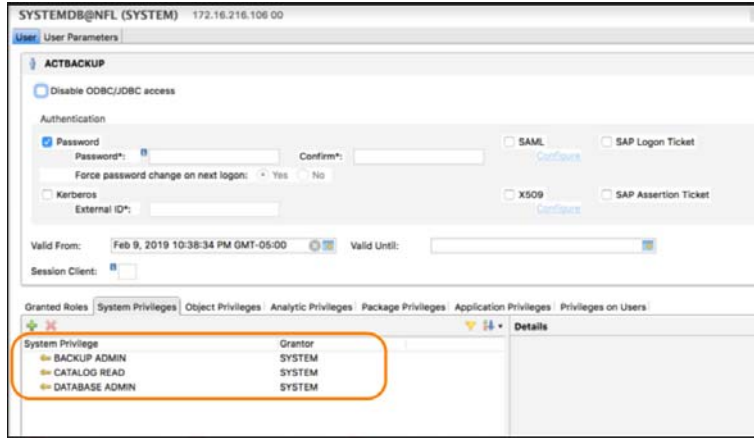
1. Create the USER under SYSTEMDB:
  - o Assign a User Name and a Password.
  - o Set Force password change on next logon to **No**.
  - o Click on the System Privilege tab and assign privileges by selecting **BACKUP ADMIN**, **CATALOG READ**, and **DATABASE ADMIN**.
  - o From SAP HANA Studio SYSTEMDB, go to System > Security > Users > New User.



2. Deploy the newly created user by clicking the green arrow in the top right corner.



You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN privileges.



## Creating the System Database User Account from HANA STUDIO, SPS02, 03, & 04

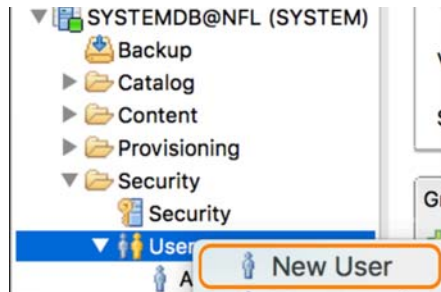
### Naming convention for database user account

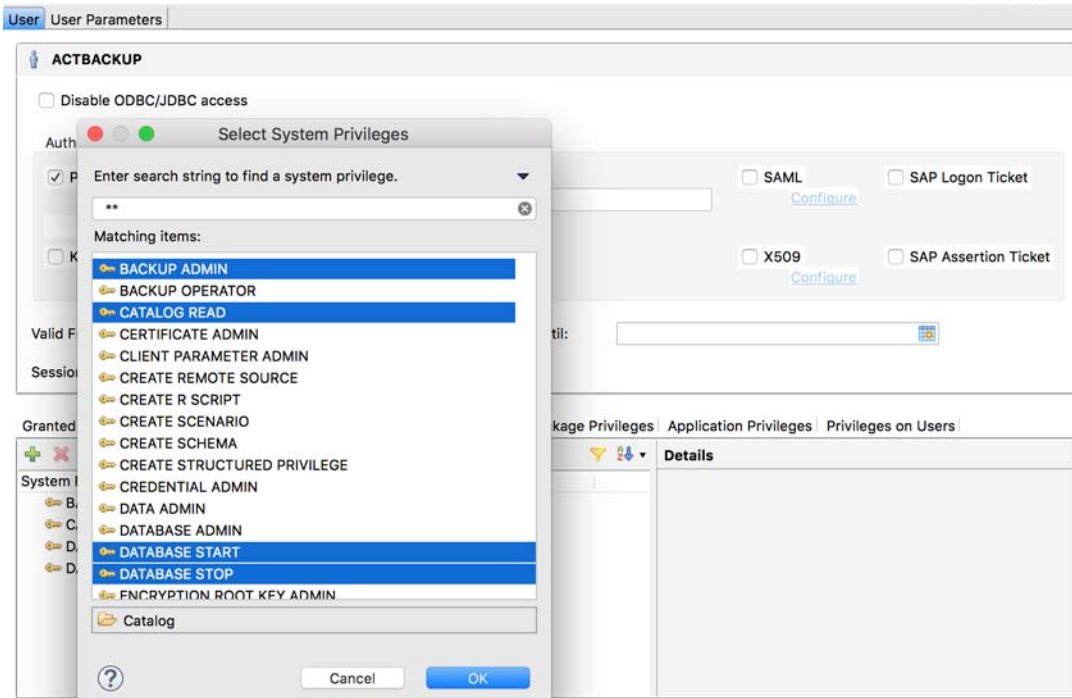
Choose a database user name based on company's standard. Make sure to create this user account under SYSTEMDB. Make sure to provide BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN to the backup user created under database.

### Procedure


To create the system database user account:

1. Create the USER under SYSTEMDB:
  - o Assign a user name and a password.
  - o Set Force password change on next logon to **No**.
  - o Click on the System Privilege tab and assign privileges by selecting **BACKUP ADMIN**, **CATALOG READ**, **DATABASE START**, and **DATABASE STOP**.
  - o From SAP HANA Studio SYSTEMDB, go to System > Security > Users > New User.

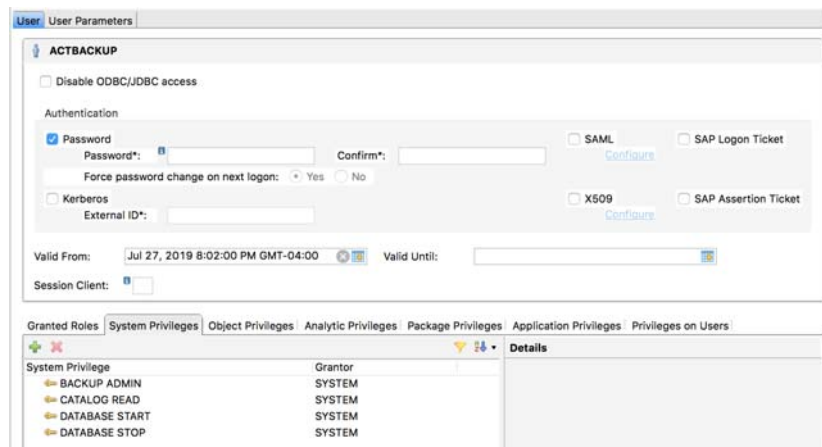




2. Deploy the newly created user by clicking the green arrow in the top right corner.



You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN, CATALOG READ, DATABASE START, and DATABASE STOP privileges.



## Creating the System Database User Account from HANA STUDIO, SPS05

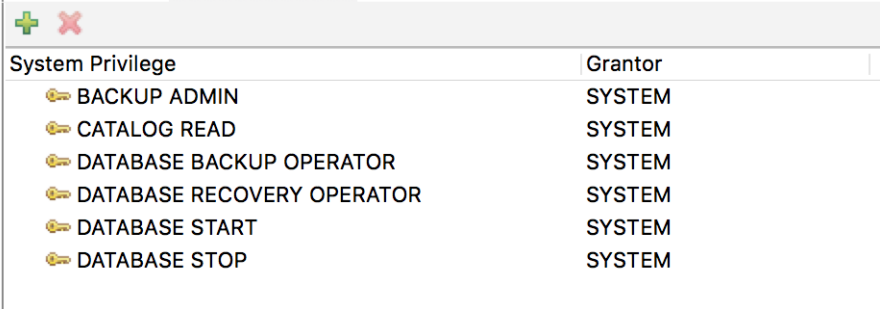
### Naming convention for database user account

Choose a database user name based on company's standard. Make sure to create this user account under SYSTEMDB. Make sure to provide BACKUP ADMIN, CATALOG READ, DATABASE BACKUP OPERATOR, DATABASE RECOVERY OPERATOR, DATABASE START, and DATABASE STOP to the backup user created under database.

## Procedure

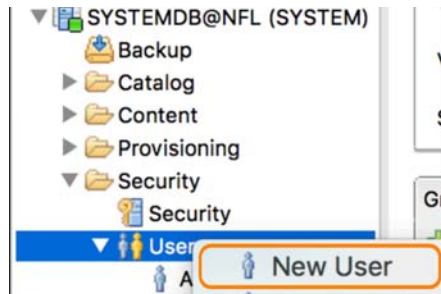
To create the system database user account:

1. Create the USER under SYSTEMDB:
  - o Assign a user name and a password.
  - o Set Force password change on next logon to **No**.
  - o Click on the System Privilege tab and assign privileges by selecting **BACKUP ADMIN**, **CATALOG READ**, **DATABASE BACKUP OPERATOR**, **DATABASE RECOVERY OPERATOR**, **DATABASE START**, and **DATABASE STOP**.




System Privilege	Grantor
BACKUP ADMIN	SYSTEM
CATALOG READ	SYSTEM
DATABASE BACKUP OPERATOR	SYSTEM
DATABASE RECOVERY OPERATOR	SYSTEM
DATABASE START	SYSTEM
DATABASE STOP	SYSTEM

- o From SAP HANA Studio SYSTEMDB, go to System > Security > Users > New User.



2. Deploy the newly created user by clicking the green arrow in the top right corner.

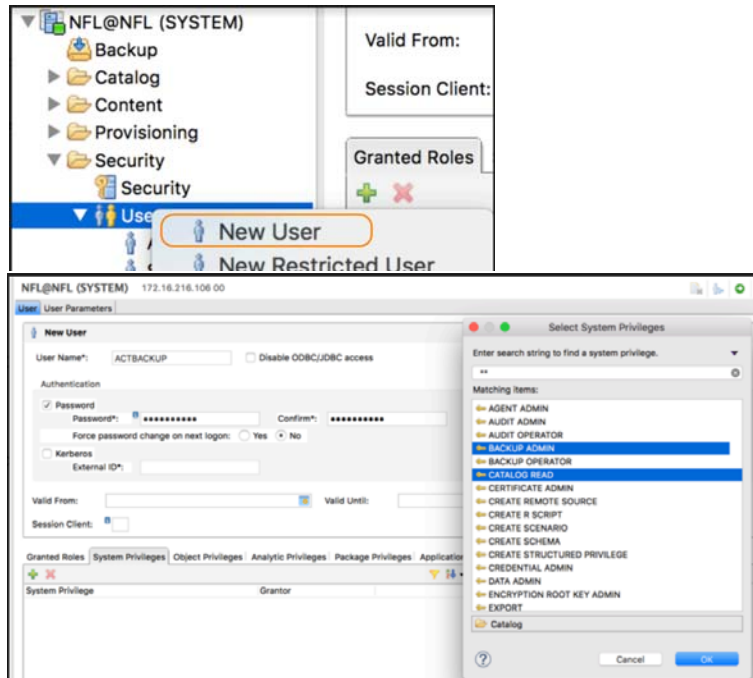


You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN, CATALOG READ, DATABASE BACKUP OPERATOR, DATABASE RECOVERY OPERATOR, DATABASE START, and DATABASE STOP privileges.


## Creating the User under the Tenant DB

To create the tenant database user account:

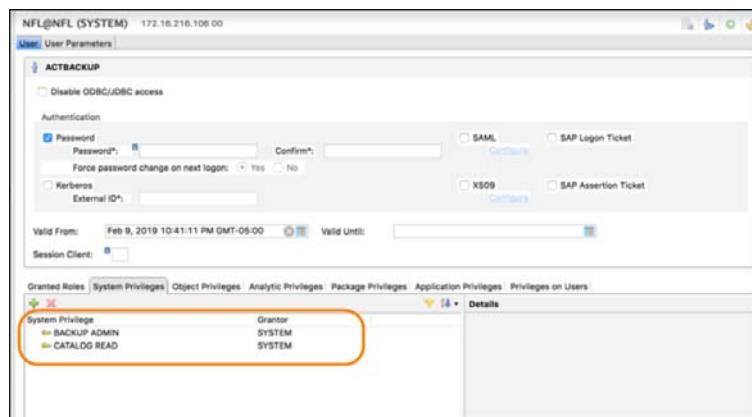
1. Create the USER under TENANTDB:
  - o Assign a user name and a password.
  - o Set Force password change on next logon to **No**.
  - o Click on the System Privilege tab and assign privileges by selecting **BACKUP ADMIN** and **CATALOG READ**.
  - o From SAP HANA Studio SYSTEMDB, go to TENANTDB > Security > Users > New User.



2. Deploy the newly created user by clicking the green arrow in the top right corner.



You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN and CATALOG READ privileges.



## Getting the Instance and SQL Port Numbers

**SYSTEMDB:** From SYSTEMDB go to System > Landscape and get the value of SQL port for the nameserver. In the example below, 30013 is the SQL port, and the instance number is 00.

Active Host	Port	Service	Detail	Start Time	Process ID	Us	Pe	Effective Allocation Limit (MB)	Physical Memory on Host (MB)	SQL Port
saphana-autovm6	30006	webdispatcher		Mar 13, 2019 11:00:42 PM	15286			16,850	40,074	
saphana-autovm6		sapstartsv								
saphana-autovm6	30000	daemon		Mar 7, 2019 4:45:08 PM		0		0		
saphana-autovm6	30001	nameserver	master	Mar 13, 2019 10:59:46 PM	14989			20,340	40,074	30013
saphana-autovm6	30010	compilesrv		Mar 13, 2019 11:00:39 PM	15243			16,619	40,074	
saphana-autovm6	30002	preprocessor		Mar 13, 2019 11:00:39 PM	15245			16,876	40,074	

**TENANT DB:** From HANA Studio. At tenantdb-System > Landscape, get the value of SQL Port for indexserver. The <port> is the SQL port of the specific tenant database, i.e. 3<instance>15

In the example below, 30015 is the SQL port, and the instance number here is 00.

Active Host	Port	Service	Detail	Start Time	Process ID	C	Memory	Used	Mr	Peak	U	Eff	Physic	SQL Port
saphana-autovm6	30006	webdispatcher		Mar 13,...	15286			1,564	1,564	1...	40,...			
saphana-autovm6	30007	xsengine		Mar 13,...	15485			2,746	3,260	1...	40,...			
saphana-autovm6	30000	daemon		Mar 7, 2...				0	0	0				
saphana-autovm6	30001	nameserver	master	Mar 13,...	14989			4,925	4,925	2...	40,...			
saphana-autovm6	30010	compilesrv		Mar 13,...	15243			1,333	1,333	1...	40,...			
saphana-autovm6	30002	preprocessor		Mar 13,...	15245			1,590	1,590	1...	40,...			
saphana-autovm6	30003	indexserver	master	Mar 13,...	15432			8,792	9,047	2...	40,...			30015

## Creating the SAP HANA Hdbuserstore Key

Use a SAP HANA hdbuserstore key to execute Backup and Recovery instead of a user name and password to communicate with HANA database using the SAP HANA Secure User Store. For HANA 2.0, the userstore key needs to be created for SYSTEMDB and all tenant db.

This includes:

[Creating the SAP HANA Hdbuserstore Key for the System Database and Each Tenant Database in a Single Node System on page 18](#)

[Creating the SAP HANA Hdbuserstore Key for the System Database and each Tenant Database in a Scale-Out Multi-Node SAP HANA System on page 19](#)

## Hdbuserstore Key Naming Convention

For SYSTEMDB set the key name = DATABASE BACKUP USERNAME.

For TENANTDB set the key name = <SYSTEMDB Key Name><TENANT DB Name>.

For example:

DATABASE BACKUP USERNAME = ACTBACKUP across SYSTEMDB and all TENANT DB

Set SYSTEMDB key name = ACTBACKUP

For tenant TDB, set TENANTDB key name = ACTBACKUPTDB

For tenant SDB, set TENANTDB key name = ACTBACKUPSDB

## Creating the SAP HANA Hdbuserstore Key for the System Database and Each Tenant Database in a Single Node System

1. Open the PuTTY window to the HANA database server and login to <sid>adm by su to <sid>adm.
2. `cd exe`
3. Create entries in hdbuserstore by calling:  
# `./hdbuserstore SET <key_name> <server>:<port> <DB_user_name> <DB_user_password>`  
The <port> is the SQL port of the systemdb or tenant database.
4. Check the keystore: `./hdbuserstore list`

### Example

Creating a SYSTEMDB hdbuserstore key:

```
./hdbuserstore SET ACTBACKUP saphana3:30013 ACTBACKUP <database backup user password  
*****>
```

Where:

- SYSTEM DB DATABASE (Backup username from above): ACTBACKUP
- KEY NAME (same as DATABASE backup username): ACTBACKUP
- SQL Port (for systemdb from above): 30013
- Hostname: saphana3

### Example

Creating a TENANTDB hdbuserstore key:

```
./hdbuserstore SET ACTBACKUPTBD saphana3:30015 ACTBACKUP <database backup user password  
*****>
```

Where:

- TENANT DB DATABASE Backup username from above: ACTBACKUP
- KEY NAME (systemdb key name postfix tenant db name): ACTBACKUPTDB
- SQL Port (for tenant db from above): 30015
- Hostname: saphana3



## Creating the SAP HANA Hdbuserstore Key for the System Database and each Tenant Database in a Scale-Out Multi-Node SAP HANA System

For a three node scale-out system with server 1, server 2, and server 3:

1. Open the PuTTY window to each HANA database server and login to <sid>adm by su to <sid>adm.
2. `cd exe`
3. On each of the HANA scale-out nodes, create entries in Hdbuserstore by running the command below:

```
# ./hdbuserstore SET <key_name> "<server 1>:<port>;<server 2>:<port>;<server 3>:<port>"  
<DB_user_name> <DB_user_password>
```

Where the <port> is the SQL port of the systemdb or tenant database.

4. Check the keystore: `./Hdbuserstore list`

### Example, SYSTEMDB hdbuserstore key

Where:

- SYSTEM DB DATABASE Backup username from above: ACTBACKUP
- KEY NAME: ACTBACKUP (same as DATABASE backup username)
- SQL Port for systemdb from above: 30013
- Hostname : saphana1, saphana 2, saphana 3

```
./hdbuserstore SET ACTBACKUP "saphana1:30013; saphana2:30013; saphana3:30013" ACTBACKUP  
<database backup user password *****>
```

### Example, TENANTDB (TDB) hdbuserstore key

TENANT DB DATABASE Backup username from above: ACTBACKUP

KEY NAME: ACTBACKUPTDB (systemdb key name, and append tenant db name)

SQL Port for tenant db from above: 30015

Hostname : saphana1, saphana2, saphana3

```
./hdbuserstore SET ACTBACKUPTDB "saphana1:30015; saphana2:30015; saphana3:30015" ACTBACKUP  
<database backup user password *****>
```

## Using the Tenant DB User Store Key Prefix

The default value for this field is <SYSTEMDB user store key><tenant db name>.

If the Tenant DB user store key uses the SYSTEMDB user store key as prefix, then you do not need a prefix value.

If the Tenant DB user store key does **not** use the SYSTEMDB user store key as prefix, then you must provide the prefix value in Application Details & Settings (see [Apply Policy Templates and Resource Profiles to Discovered Databases](#) on page 11).

### Use case 1

You have created a user store key and you have a tenant database TDI:

SYSTEMDB user store key = ACTBACKUP

TENANT DB user store key = ACTBACKUPTDI

Under Application Details & Settings:

1. At HANA DB USER STORE KEY, provide the user store key of SYSTEMDB:

HANA DB USER STORE KEY \*

2. Leave TENANT DB USER STORE KEY PREFIX value empty.

TENANT DB USER STORE KEY PREFIX

### Use case 2

You have created a user store key and you have a tenant database TDBACKUPTDI (tenant db name is different from system db name):

SYSTEMDB user store key = ACTBACKUP

TENANT DB user store key = TDBACKUPTDI

Under Application Details & Settings:

1. At HANA DB USER STORE KEY, provide the user store key of SYSTEMDB:

HANA DB USER STORE KEY \*

2. At TENANT DB USER STORE KEY PREFIX (at the bottom of the screen), enter the "TDBACKUP" part of the name as a prefix:

TENANT DB USER STORE KEY PREFIX

# 4 Adding a SAP HANA Database Host and Discovering Databases

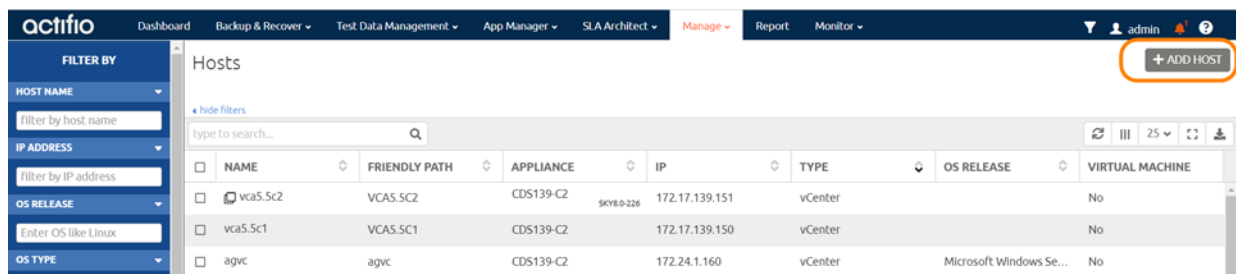
Before you can protect an SAP HANA database, you must add the host and discover the database. This requires:

1. [Adding the Host to AGM on page 21](#)
2. [Onboarding the HANA Databases from the App Manager on page 22](#)

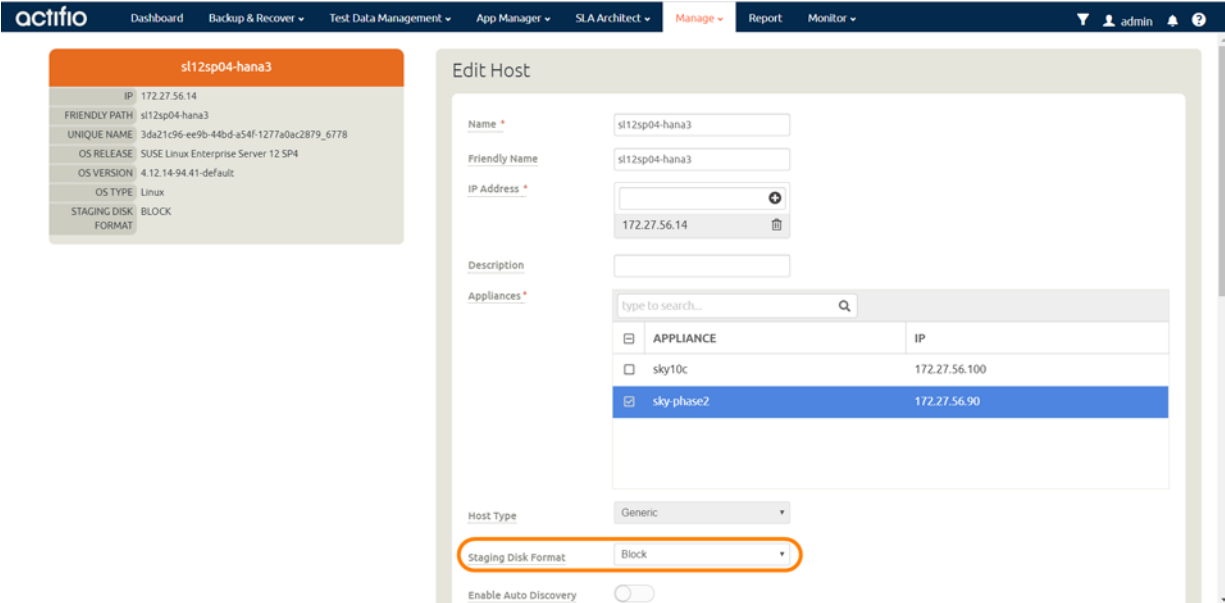
## Adding the Host to AGM

Add the host to AGM. If the host is already added then edit the host and make sure to set the Staging Disk Format correctly.

1. From the AGM Manage, Hosts page, click **+Add Host**.



2. On the Add Host page:
  - o **Name:** Provide the database server name.
  - o **IP Address:** Provide the database server IP and click the **+** in the right corner.
  - o **Appliances:** Select the check box for the Actifio Appliance that will manage the data.
  - o **Host Type:** Make sure this is Generic.
3. Click **Add** at bottom right to add the host. The Host is added.
4. Right-click the host and select **Edit**. On the Edit Host page, select the staging disk format:
  - o **Block**-based staging disks are the most useful for both backup/recovery and TDM usage. Actifio changed-block tracking (CBT) is only available on block-based staging disks, and virtual databases can only be mounted to block-based staging disks.
  - o **NFS** staging disks permit only traditional file-based backup with Full+Incremental file system backup. Select NFS only if Block is not an option in your network.




5. Select **Save** at the bottom of Edit Host page.

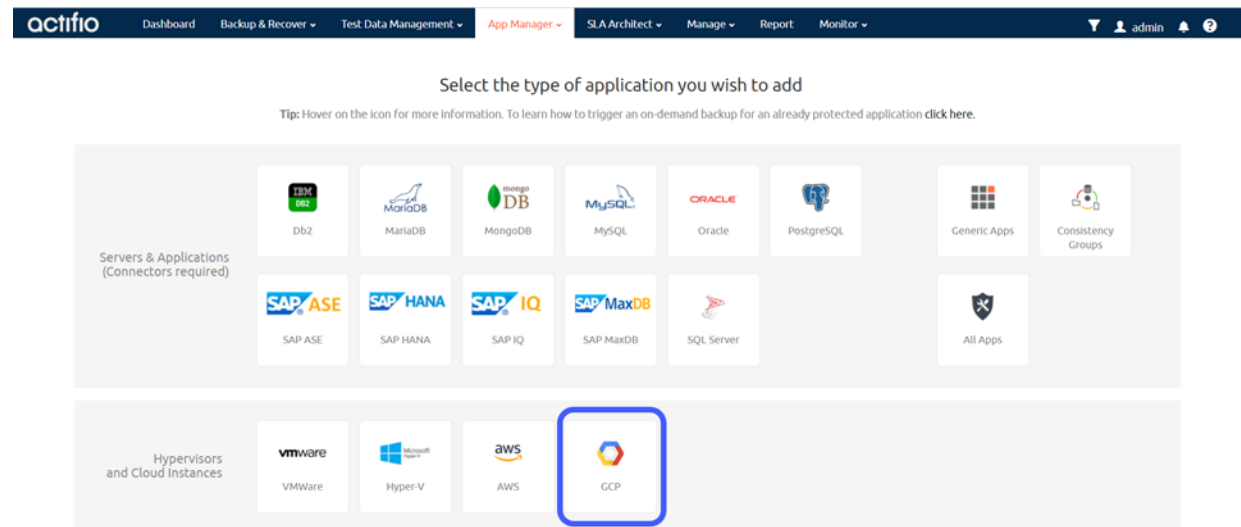
## Onboarding the HANA Databases from the App Manager

To discover and protect the HANA database applications:

1. From the AGM App Manager, Applications page, select **Add Application** in the upper right corner.



2. In the Add Application wizard, select **SAP HANA**. Then follow the wizard.



---

# 3 Defining Actifio GO Policy Templates and Resource Profiles

---

The policy template and the resource profile that make up the SLA define the type of data capture to perform and where to store the captured image. Based on the required RPOs and desired backup storage location we will need to create Templates and Profiles. The policy part of the template provides additional backup configuration like enabling logs backups.

A resource profile defines the Sky appliances that are identified as the primary (local) appliance hosting the disk pool to use for snapshots and a remote appliance for remote backup to be used for disaster recovery operations. The profile also defines the OnVault pools to be used to send backup data to an object storage device or cloud offering, such as Google Cloud Storage.

You can use the existing the template and profile to protect the application or refer to [Creating a Policy Template](#) and [Creating a Resource Profile](#) to create new template and profile respectively.

## Apply Policy Templates and Resource Profiles to Discovered Databases

Once you have discovered one or more applications, applying a policy template and resource profile to the application enables you to capture data operations.

You choose between two very different backup methods in the Application Details & Settings:

- **Use volume level backup (SAP HANA Savepoint API):** Use volume level LVM snapshots with CBT on Linux to a block-based staging disk. This option enables you to create application-aware virtual database copies from the snapshot images.
- **Use full+incremental backup (SAP Backint API):** This is the traditional file-based backup and recovery. This "file dump" method does not support the creation of virtual databases. You can select this for both Block and NFS staging disks. This method only supports traditional backup and physical recovery.

When creating a snapshot policy as part of creating a template you have the option of also capturing its log files at a specified frequency. Details are in [Creating a Policy Template, Step 9](#).

You can replicate database logs to a remote Sky appliance or to an OnVault. You can use the logs at the remote site for any database image within the retention range of the replicated logs. Log replication uses StreamSnap technology to perform the replication between the local and remote Sky appliances; log replication goes directly from the local snapshot pool to the snapshot pool on the remote appliance. For a log replication job to run, there must be a replication policy included in the template, and at least one successful replication of the database must first be completed.



# 6 Protecting the HANA Database and Its Logs

This includes:

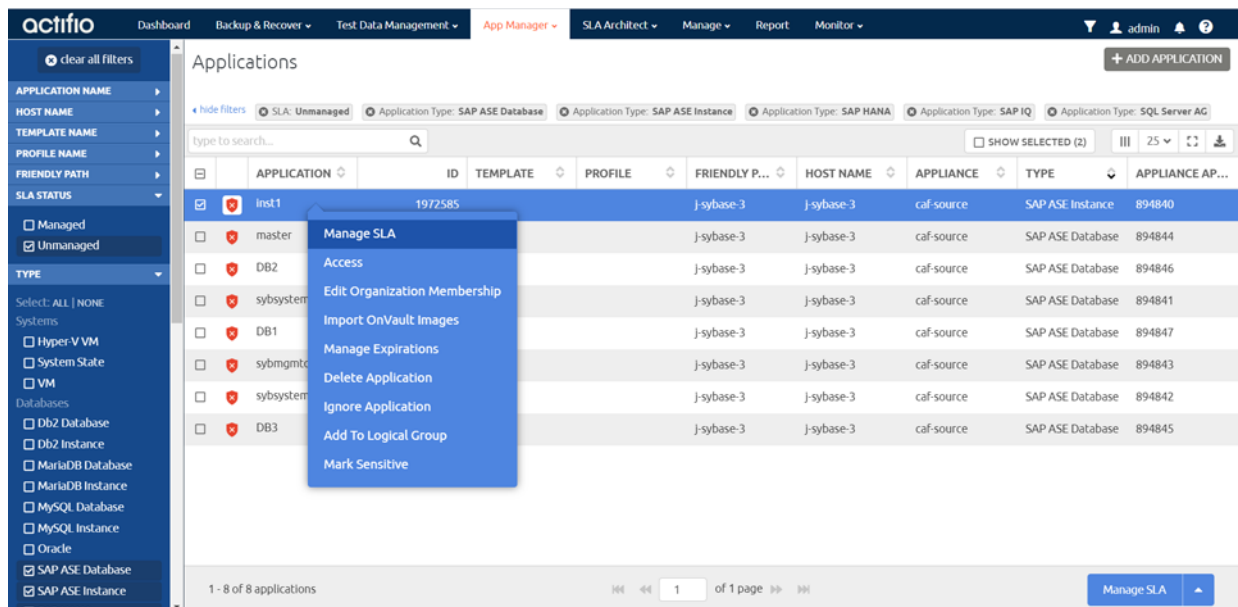
[Protecting the HANA Database on page 25](#)

[Protecting HANA Database Logs on page 27](#)


## Protecting the HANA Database

To protect the database:

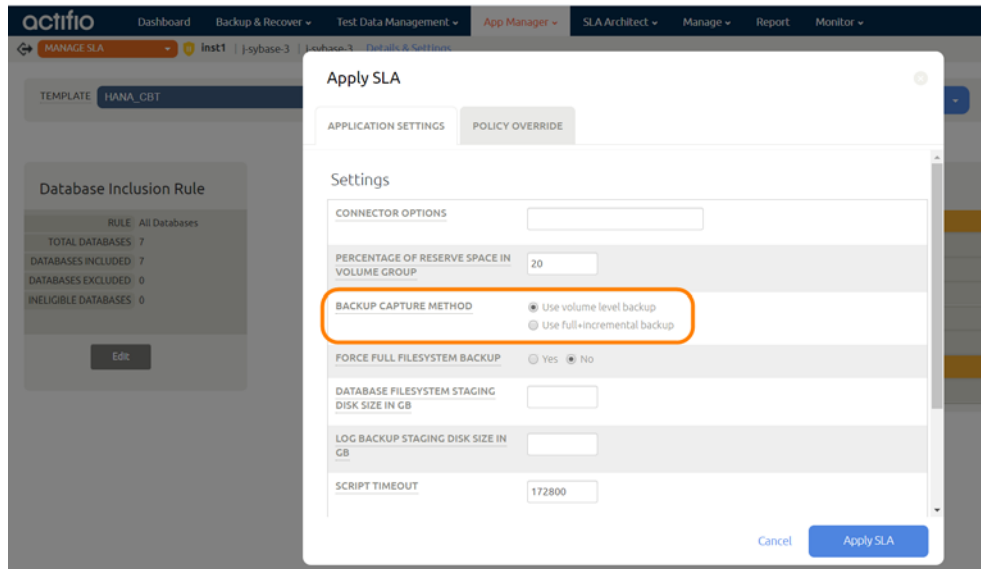
1. From the AGM App Manager, Applications list, right-click the HANA database and select **Manage SLA**.



2. On the Manage SLA page, select a template and a profile, then click **Apply SLA**.



3. On the Apply SLA page, fill in the required field based on type of backup as detailed in [Apply Policy Templates and Resource Profiles to Discovered Databases](#) on page 11. Click **Apply SLA**.



The database will be protected when the snapshot job runs according to the schedule in the template. After the first successful snapshot job, the database will appear in the Application Manager as protected, with a green shield icon.



# Protecting HANA Database Logs

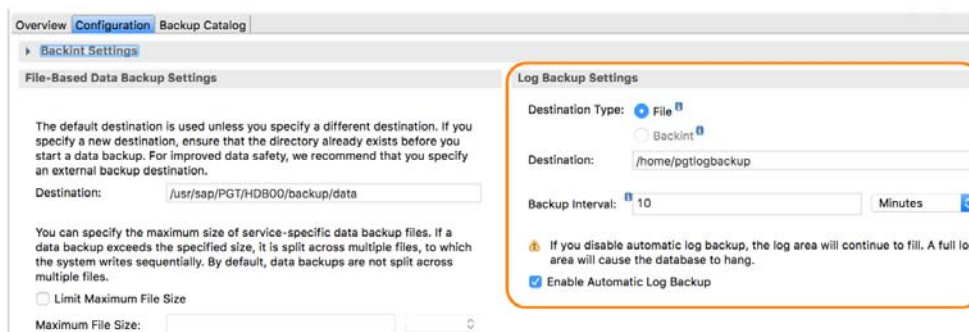
There are two parts to configuring protection of SAP HANA database logs:

[Setting up the Log Mode and Log Backup in HANA Studio on page 27](#)

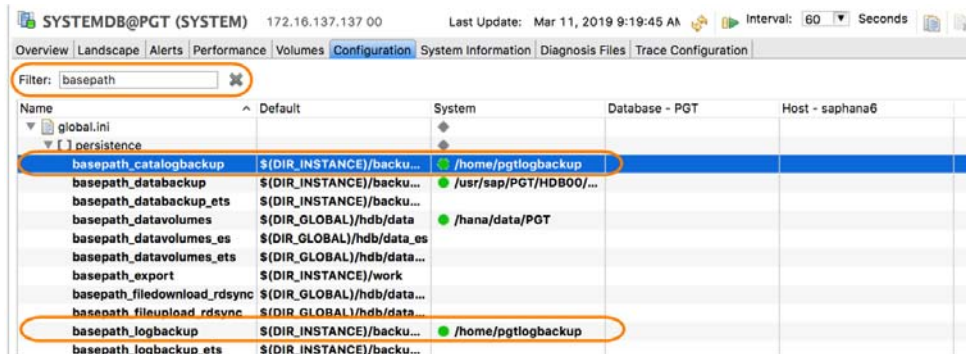
[Setting up the Log Backup in Actifio AGM on page 29](#)

## Setting up the Log Mode and Log Backup in HANA Studio

1. In SAP HANA HDB studio, make sure log backup is set correctly under the DATABASE (HANA 1.0) or SYSTEMDB (HANA 2.0) > Backup > Configuration page:
  - o Destination Type is File.
  - o Destination is set to a local file system mount path.
  - o Backup Interval is set to required RPO.
  - o Automatic Log Backup is enabled.



2. Check under Database configuration: DATABASE (HANA 1.0) or SYSTEMDB (HANA 2.0) > Configuration page.  
In the filter, type **basepath**.



3. Verify that basepath\_logbackup is set correctly:
  - o Set the basepath\_catalogbackup to the same as basepath\_logbackup.
  - o Open the basepath\_catalogbackup edit page.
  - o Set the New Value to the same as basepath\_logbackup and click **Save**. This will ensure the backup of catalog with log backup for point in time recovery.

**basepath\_catalogbackup**  
global.ini [persistence]

Default Value:

System

Active Value:

New Value:

4. Make sure tenant db log backup is set correctly under DATABASE (HANA 1.0) or TENANTDB (HANA 2.0) > Backup > Configuration page:
  - o Destination Type is File.
  - o Destination is set to a local file system mount path.
  - o Backup Interval is set to required RPO.
  - o Automatic Log Backup is enabled.

Overview **Configuration** Backup Catalog

Backint Settings

File-Based Data Backup Settings

The default destination is used unless you specify a different destination. If you specify a new destination, ensure that the directory already exists before you start a data backup. For improved data safety, we recommend that you specify an external backup destination.

Destination:

You can specify the maximum size of service-specific data backup files. If a data backup exceeds the specified size, it is split across multiple files, to which the system writes sequentially. By default, data backups are not split across multiple files.

Limit Maximum File Size

Maximum File Size:

Log Backup Settings

Destination Type:  File  Backint

Destination:

Backup Interval:  Minutes

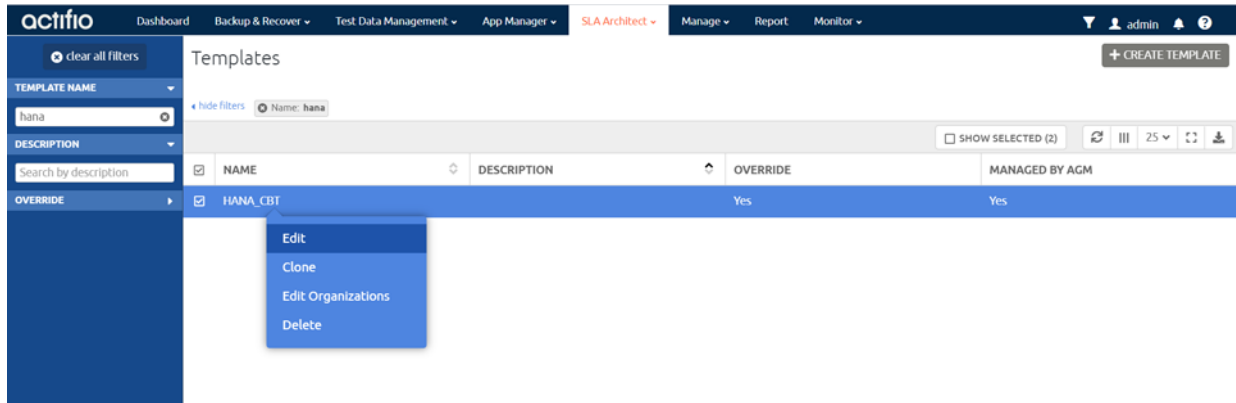
Enable Automatic Log Backup

**!** If you disable automatic log backup, the log area will continue to fill. A full log area will cause the database to hang.

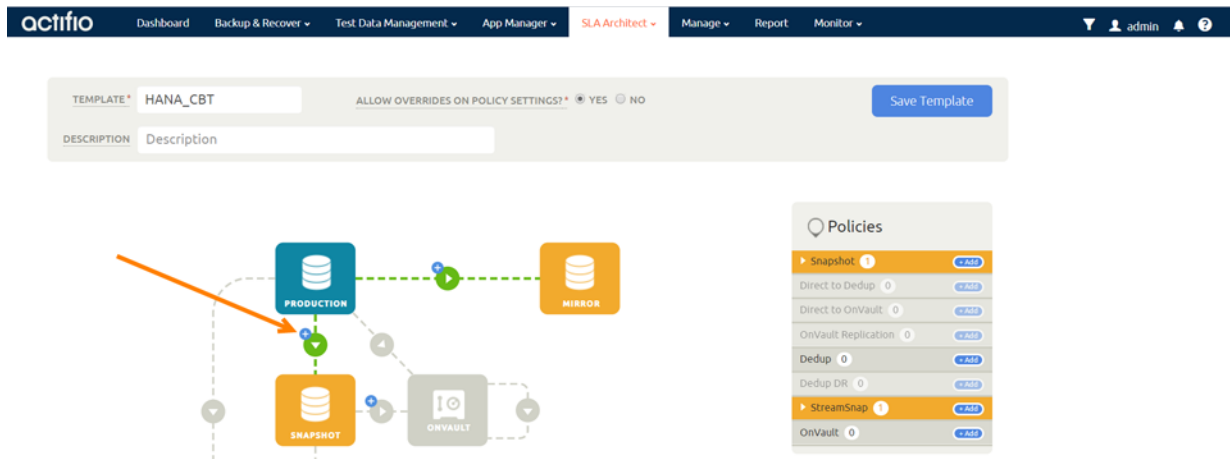
# Setting up the Log Backup in Actifio AGM

To enable and set up the HANA database log backup:

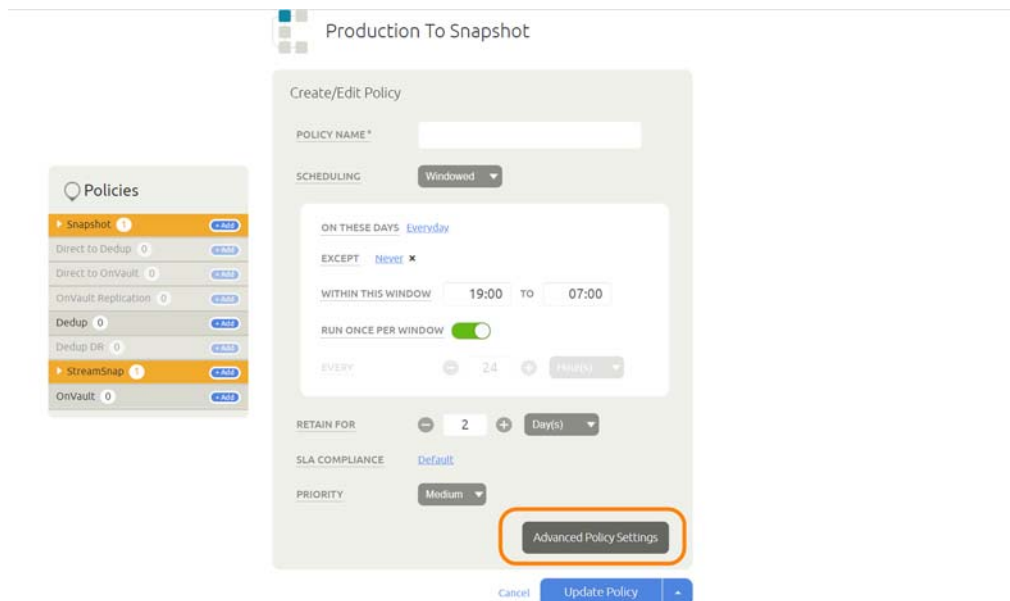
1. From the SLA Architect page, edit the template created for HANA database protection:




2. Click the Production to Snapshot "+".




3. Select **Advanced Policy Settings**.




4. Set the log policy options (you will have to scroll to see them all):
  - o Truncate/Purge Log After Backup: Select this.
  - o Enable Database Log Backup: Select this.
  - o RPO (Minutes): Enter the desired frequency of log backup
  - o Log Backup Retention Period (in Days): SLA to retain log backup for point in time recovery.
  - o Replicate Logs (Uses StreamSnap Technology): Select this to enable StreamSnap replication of log backup to a DR site.
  - o Send Logs to OnVault Pool: Set this to Yes if you want the database logs to be sent to an OnVault Pool, enabling point-in-time recoveries from OnVault on another site.



5. There is one more application-specific setting. From the App Manager, Applications list, select the HANA database. You can use the SAP HANA checkbox to filter the list. Select **Manage SLA**.



6. At the top of the screen, select **Details & Settings**.
7. Set the **Retention of Production DB Logs in Days**. This is used to purge the HANA log backup from basepath\_logbackup destination. Based on this setting the last data backup id will be selected (CURRENT\_TIMESTAMP - the # days set) and logs older than the data backup id will be purged. The default is 0 days: all logs prior to last data backup are purged. Click **Save Changes**.



---

# 7 Restoring or Recovering an SAP HANA Database

---

This section includes:

- [Instant Recovery of a HANA Database from a Volume-Based Snapshot to the Source on page 31](#)
- [Instant Recovery of a HANA Database from a Block-Based LVM Snapshot to a New Target on page 34](#)
- [Recovering a Single Tenant Database from a Volume-Based Snapshot on page 36](#)
- [Restoring a HANA Database from a Volume-Based Snapshot Back to the Source on page 37](#)
- [Restoring a HANA Database from a Full+Incremental Snapshot Back to the Source on page 38](#)
- [Recovering from a Full+Incremental Snapshot to a New Target on page 39](#)

## Instant Recovery of a HANA Database from a Volume-Based Snapshot to the Source

Instant recovery from block-based LVM snapshot protection has three steps:

- Mount the backup snapshot image from Actifio to source HANA server as a standard mount.
- Recover and bring up the HANA database from the Actifio mounted image.
- Migrate the data online from the mounted Actifio image to the HANA production disk while the database is up and running.

To recover from a volume-based snapshot:

1. From the AGM App Manager, Applications list, right-click the application and perform a standard mount (do not create a virtual application).
2. Once the mount job is completed, run this script with the parameters mentioned below:

```
/act/custom_apps/saphana/clone/ACT_HANADB_mountrecover.sh <DBSID> <TARGET MOUNT POINT> <DB USER> <HANA VERSION> <DATA PATH> <LOG PATH> [OLD DBSID] [LOGMOUNT PATH] [RECOVERY TIME]
```

DATA PATH and LOG PATH can also be obtained from the `/act/touch/<dbsid_mount_params>` file, which gets created after mount job is finished.

### Parameters

DBSID = Target Database SID  
TARGET MOUNT POINT= <Mount Location specified in AGM>  
DB USER = < USERSTORE KEY >  
HANA VERSION = <2.0>  
DATA PATH = <Target Data Volume>  
LOG PATH = <Target Log volume>  
OLD DBSID = <Source DB SID>  
LOGMOUNT PATH = <Transaction log Mount point>  
RECOVERY TIME = '<Recovery time YYYY-MM-DD HH24:MI:SS in UTC>

## Example

```
/act/custom_apps/saphana/clone/ACT_HANADB_mountrecover.sh p01 /testmnt ACTBACKUP 2.0 /testmnt/hana/data/P01 /testmnt/hana/log/P01 p01
```

---

**Note:** The HANA database will be recovered to the chosen point in time and will be up and running and available for application access.

---

3. The migration script can be started when the production volumes are available, which moves all the data from Actifio volumes to the production volumes.

---

**Note:** The HANA database will be restarted to start the migration. The HANA database will be available after restart and the migration will continue online in the background.

---

```
/act/custom_apps/saphana/lvm_migrate/ACT_HANADB_lvm_migrate.sh <DBSID> <PROD DATA VOLUME> <PROD LOG VOLUME>
```

Where:

DBSID= <Database SID>

PROD DATA VOLUME= <Data volume for migration>

PROD LOG VOLUME=<Log volume for migration>

4. Comment PROD DATA VOLUME and PROD LOG VOLUME from /etc/fstab until migration is finished.

For example:

```
df -kh
/dev/mapper/hanavg-data      443G   11G  415G   3% /hana/data
/dev/mapper/hanavg-log      50G    7.5G  40G  16% /hana/log
```

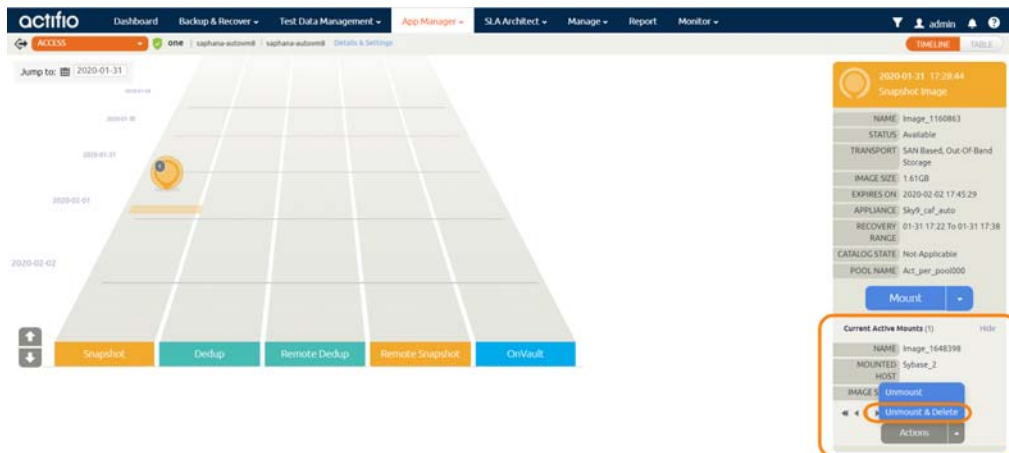
Where:

PROD DATA VOLUME= /dev/mapper/hanavg-data

PROD LOG VOLUME= /dev/mapper/hanavg-log

```
/act/custom_apps/saphana/lvm_migrate/ACT_HANADB_lvm_migrate.sh p01 /dev/mapper/hanavg-data /dev/mapper/hanavg-log
```

5. Once this script completes, go to AGM and Unmount + Delete the image.



6. The migration process is finished. Uncomment the /etc/fstab entries for DATA and LOG volumes.

## Clean-up Procedure After a Failed Migrate Job

If a Migrate job fails, the host logical volume groups may be in an inconsistent state with respect to production volumes and Actifio staging volumes that were used during the migration process. Before starting another migration job or to put the HANA database back to the state preceding the migration job, you must make the host volume groups consistent again:

1. Check the contents of `/act/hana_recovery.log` to learn the failure location. The contents of this file point out at what step the migration failed. Look for "ERRORMSG:" string to determine the exact failure reason.
  - If the failure is before "Starting Device migration steps" then skip to [Step 6](#).
  - If the failure is during or after "Starting Device migration steps" then continue to [Step 2](#).
2. Check to see if Actifio volumes are part of the production volume group. Run:

```
vgdisplay -v <production volume group name>
```

Examine the output to determine if Actifio volumes appear in the output under the section called "Physical Volumes". Actifio devices have a "PV Name" starting with **`/dev/mapper/3638`**.

- If you see raw block device names like `/dev/sdc`, for example, then Actifio volumes *are* part of the volume group, and you must fetch the LUN UUID of the device. Run:

```
/opt/act/bin/udsagent fetchlunid <block device name>
```

For example: `/opt/act/bin/udsagent fetchlunid /dev/sdc`. Then proceed to [Step 3](#).

- If Actifio volumes are *not* part of the volume group, then from the AGM, unmount and delete the mounted image and the cleanup procedure is finished.
3. If Actifio volumes *are* part of the volume group, then see if there are partial extents on target volumes (this is an uncommon case). Refer back to the output of `vgdisplay` from [Step 2](#). For each device identified as an Actifio device, locate the "Total PE / Free PE" entry.
    - If the entries for both are the same, then there are no partial extents on target physical volumes, so skip to [Step 4](#).
    - If they are not the same, then partial extents have been moved from the Actifio device into the target physical volumes. You must deactivate the production volume group and remove the Actifio device from the volume group:

```
vgchange -an <production volume group name>
```

```
lvremove <Actifio device name identified in Step 2>
```

4. In the more common scenario, all the physical extents have been moved into the target physical volumes. The next step is to remove the Actifio device identified in [Step 2](#) from the production volume group, then remove the Actifio physical volume:

```
vgreduce <prod volume group> <Actifio device name>
```

```
pvremove <Actifio device name>
```

5. If there were no partial extents then you did not need to deactivate the production volume group, so skip to [Step 6](#). If you did have to deactivate the production volume group, then reactivate the production volume group with the following command:

```
vgchange -ay <prod volume group>
```

6. Revert the `global.ini` file to its original saved contents:
  - a. Locate the `global.ini` file that was created as a backup. This can be found in the `.ini` file location and will be named as `global.ini.<date-and-timestamp>`.
  - b. Copy this file to `global.ini` file (overwriting any existing `global.ini` file).
7. Clean up and remove any folders that have names starting with the mount job name (`Job_xxx_/act/mnt`).
8. From the AGM, unmount and delete the mounted image.
9. The cleanup procedure is finished. You can start the database from its original configuration.



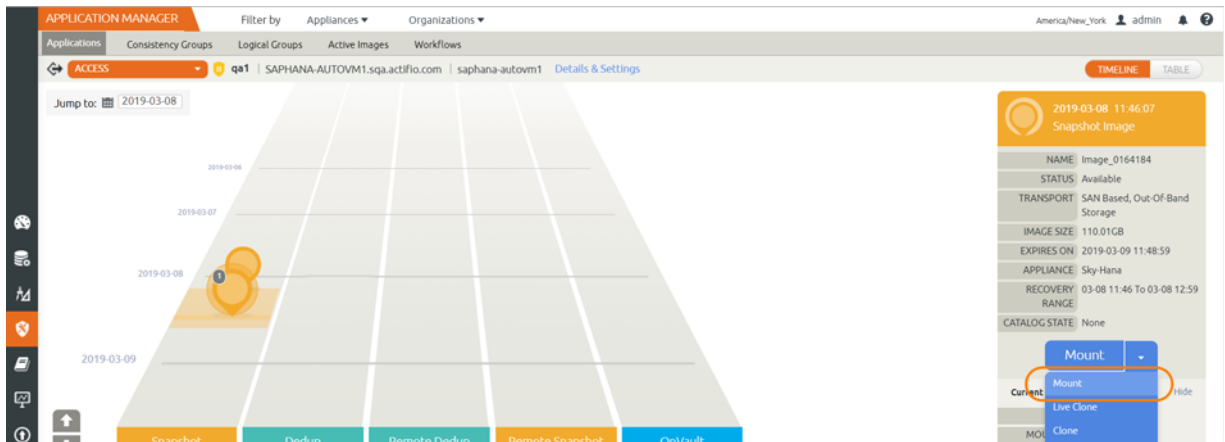
# Instant Recovery of a HANA Database from a Block-Based LVM Snapshot to a New Target

There are two stages to this recovery process:

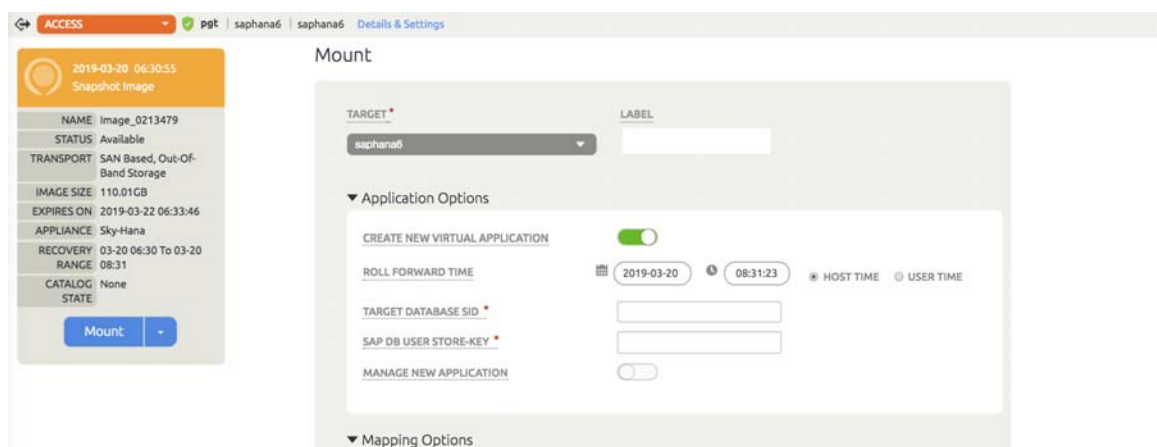
1. Mount the Snapshot as a New Virtual Database from to the New HANA Target Server
2. Migrate the Data to Production Storage

## Mount the Snapshot as a New Virtual Database from to the New HANA Target Server

1. From App Manager > Protected Application > Access, from the latest snapshot, choose **Mount**.



2. On the Mount page, from Target, choose the desired target HANA server from the dropdown.
3. Under Application Options:
  - o Select **Create New Virtual Application**.
  - o If the database was protected with log roll-forward, choose a point in time on the slider bar for recovery.
  - o For Target Database SID, provide the target HANA database name.
  - o For SAP DB User Store-Key, provide the hdbuserstore key for the target database (for HANA 2.0: SYSTEMDB).
  - o For Mount Location, specify a Mount Point to mount to new target.
  - o At Manage New Application, if you want to reprotect the database, click and enable Manage New Application.
  - o Choose a template and a profile to protect the database.



4. Click **Submit**.



---

**Note:** The HANA database will be recovered to the chosen point in time and will be up and running and available for application access.

---

## Migrate the Data to Production Storage

This will require valid Logical Volumes for data and log.

---

**Note:** The HANA database will be stopped momentarily prior to the migration. HANA database will be available once migration is initiated.

---

---

**Note:** The HANA database will be restarted to start the migration. The HANA database will be available after restart and the migration will continue online in the background.

---

```
/act/custom_apps/saphana/lvm_migrate/ACT_HANADB_lvm_migrate.sh <DBSID> <PROD DATA VOLUME> <PROD LOG VOLUME>
```

Where:

DBSID= <Database SID>

PROD DATA VOLUME= <Data volume for migration>

PROD LOG VOLUME= <Log volume for migration>

For example:

```
df -kh
```

```
/dev/mapper/hanavg-data
```

```
/dev/mapper/hanavg-log
```

Where:

```
443G 11G 415G 3% /hana/data
```

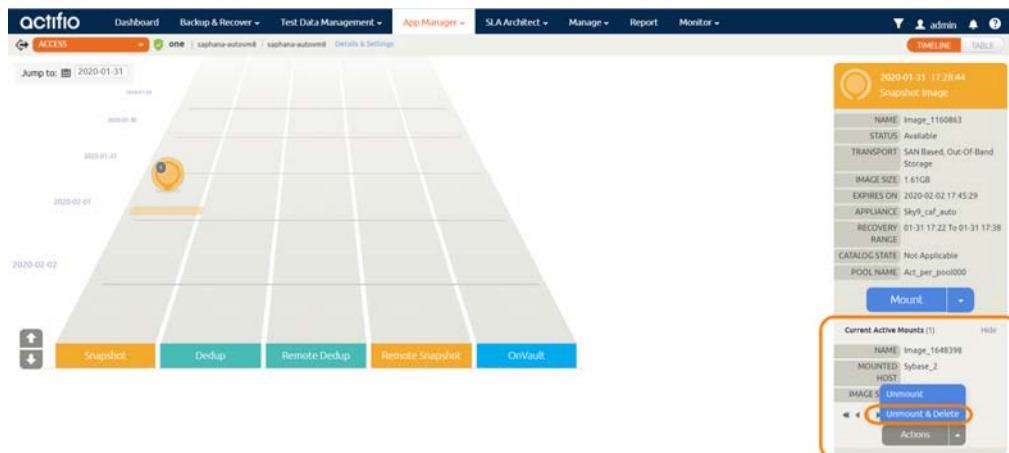
```
50G 7.5G 40G 16% /hana/log
```

```
PROD DATA VOLUME= /dev/mapper/hanavg-data
```

```
PROD LOG VOLUME= /dev/mapper/hanavg-log
```

```
/act/custom_apps/saphana/lvm_migrate/ACT_HANADB_lvm_migrate.sh p01 /dev/mapper/hanavg-data /dev/mapper/hanavg-log
```

5. Once this script completes, go to AGM and Unmount + Delete the image.

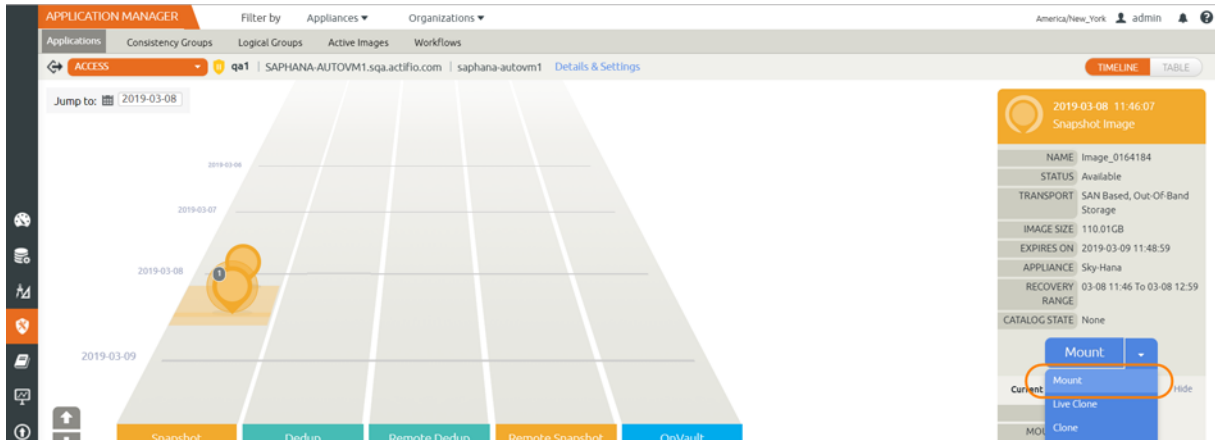


6. The migration process is finished. Uncomment the `/etc/fstab` entries for DATA and LOG volumes.

# Recovering a Single Tenant Database from a Volume-Based Snapshot

To recover a single-tenant database:

1. From the App Manager, Applications list, right-click the database and select **Access**. Then select the latest snapshot to recover, and choose **Mount**.



2. On the server, change the directory to `/act/custom_apps/saphana/restore`:

```
cd /act/custom_apps/saphana/restore
```

3. Execute the script for recovery:

```
./CALL_LVM_single_tenant_recover.sh <DBSID> <TENANT SID> <SYSTEMDB USERSTORE KEY> '<RECOVERY TIME-YYYY-MM-DD HH24:MI:SS>'
```

Description of arguments to the script:

DBSID = < Database SID>

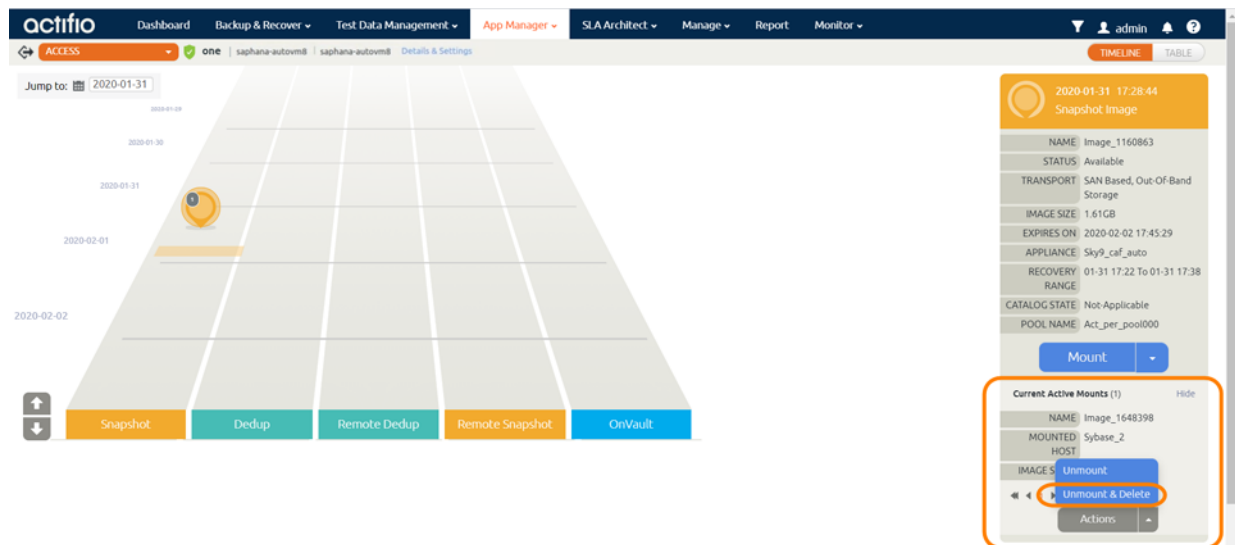
TENANT SID = < Name of the Tenant DB to be restored>

SYSTEM DB USERSTORE KEY = < System DB User store key>

RECOVERY TIME = < Recovery time YYYY-MM-DD HH24:MI:SS in UTC>

For example: `./CALL_LVM_single_tenant_recover.sh lv1 lv1 ACTBACKUP '2019-09-24 20:00:00'`

4. Once the script has completed, the Tenant DB is ready for validation. Unmount and delete the backup image.



# Restoring a HANA Database from a Volume-Based Snapshot Back to the Source

Use this procedure to restore and recover the source HANA database. This procedure uses physical recovery of the source data area.

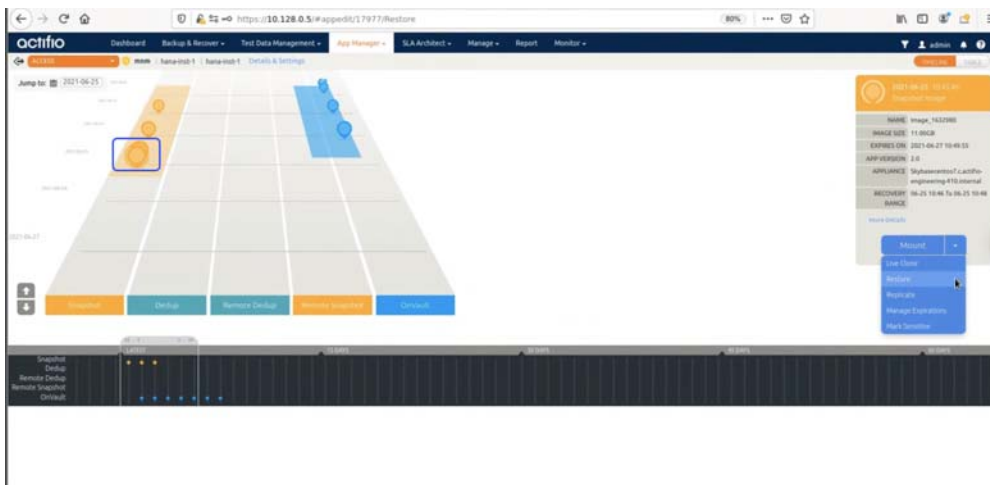
---

**Note:** System databases on a root partition backed up as LVM Snapshots can be mounted as virtual databases, but they cannot be used in a traditional Restore operation as the root partition cannot be unmounted. This will need manual restore and recovery from a simple mount back to the same host.

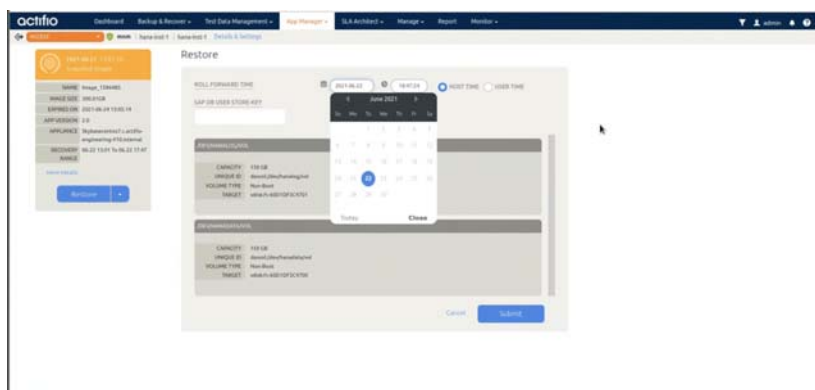
---

To recover back to the source:

1. From the App Manager, Applications list, right-click the database and select **Access**. From the latest snapshot to recover, choose **Restore**.



2. On the Restore page choose a date and time for a database protected with logs to recover to the desired point in time.

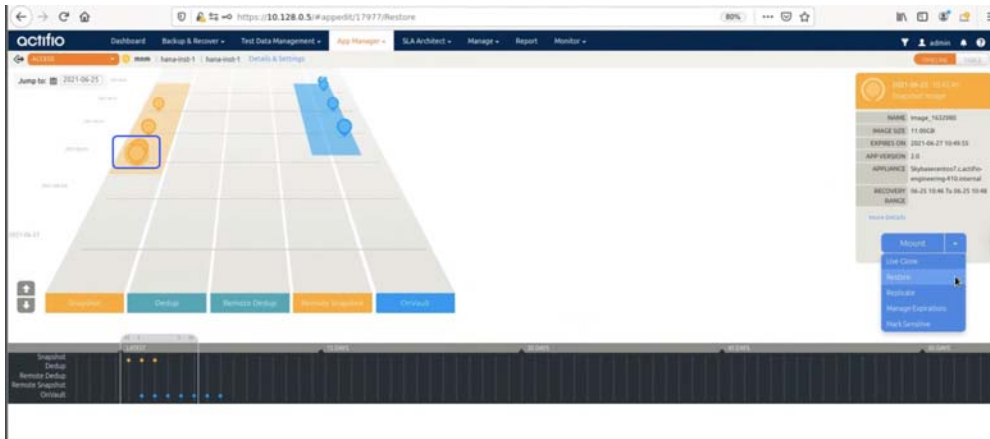


3. Click **Submit**.

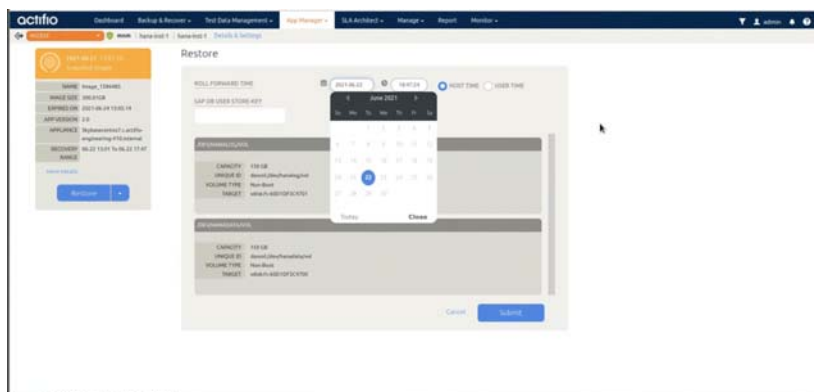
# Restoring a HANA Database from a Full+Incremental Snapshot Back to the Source

Use this procedure to restore and recover the source HANA database. This overwrites the source data.

1. From App Manager, Applications list, right-click the database and select **Access**.
2. Select the latest snapshot to recover, and choose **Restore**.



3. For a database protected with logs, on the Restore page, choose a date and a point in time.



## Notes

- HANA 1.0: EXCLUDE and INCLUDE db list do not apply.
- HANA 2.0
  - o Only one out of EXCLUDE and INCLUDE is applicable at a time.
  - o Complete HANA recovery leave EXCLUDE AND INCLUDE empty.
  - o INCLUDE LIST: For recovering one or more database out of n database: provide comma separated list of database under INCLUDE.
  - o EXCLUDE LIST: For excluding one or more database during recovery out of n database: provide comma separated list of database under EXCLUDE.



4. Click **Submit** to start the source database physical recovery using HANA recover commands.

# Recovering from a Full+Incremental Snapshot to a New Target

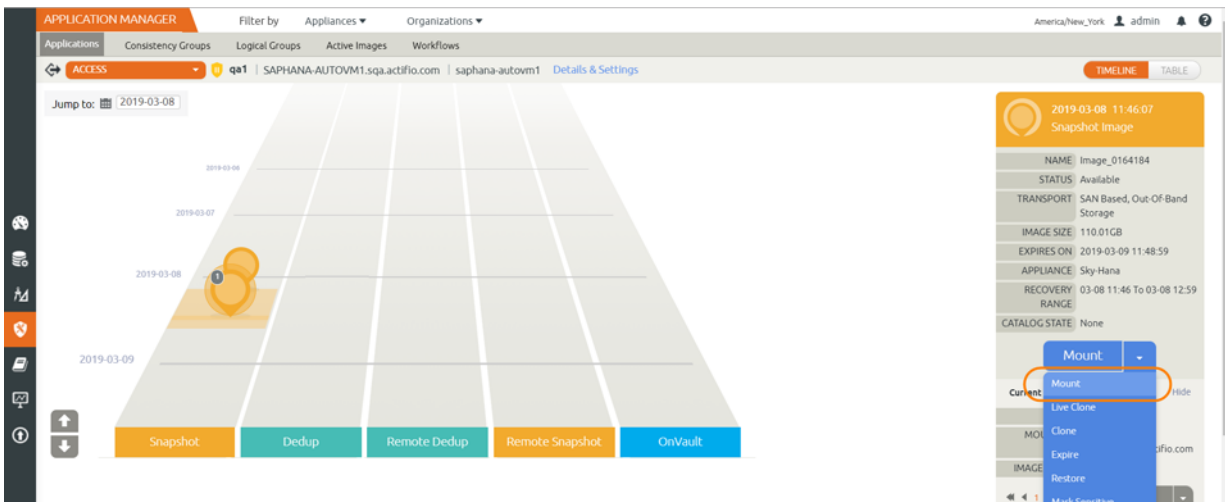
Use this procedure to restore and recover to a new target server.

Before you begin:

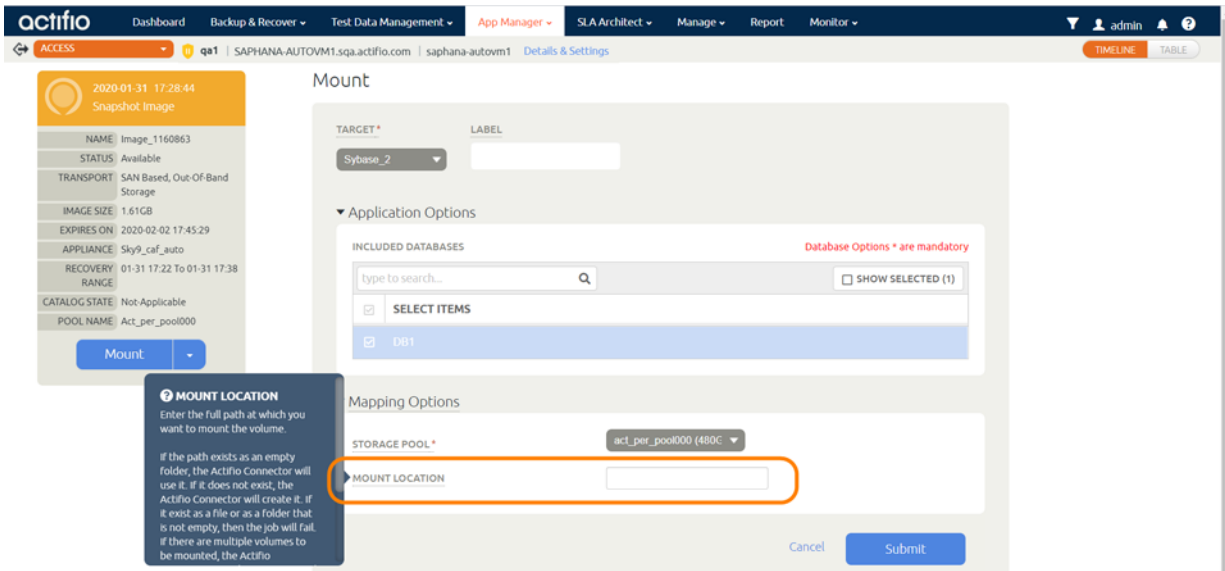
- Make sure the target HANA server is set up the same as the source HANA server (OS version, CPU and memory, HANA version).
- Make sure the HANA database on the target server is configured the same as the source, *i.e.* global.ini, nameserver.ini.

To recover to a new target server:

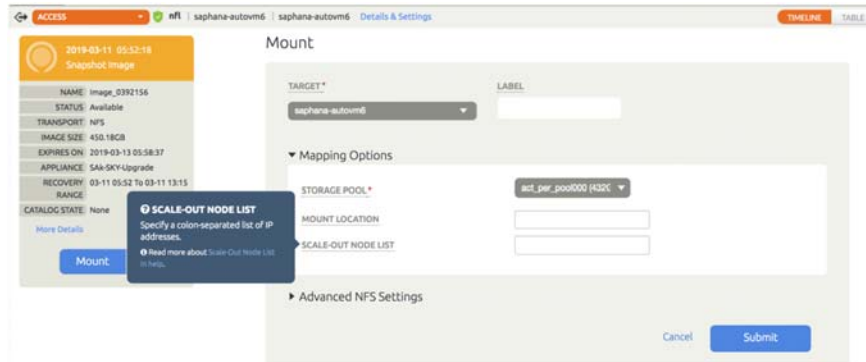
1. From App Manager, Applications list, right-click the database and select Access. Then select the latest snapshot to recover, and choose **Mount**.



2. On the Mount page, specify a mount location to mount to new target.



3. Enter scale-out information:
  - o For non-scale out HANA: leave SCALE-OUT NODE LIST empty.
  - o For scale-out HANA environment: provide a colon-separated list of target HANA servers.



4. Click **Submit**. This will mount the backup image to target server. In case of scale out, the image will be mounted to all nodes as NFS shared volume.
5. Next you must bring up the HANA database. To bring up the HANA database from the mounted image, modify and configure `/act/custom_apps/saphana/dump/restoreDumpToNewTarget.conf`:

```
DBSID=<source database sid>
DBPORT="HDB<instance #>" ex:for instane# 00 this will be "HDB00"
HANABACKUPPATH=<mount path from mount operation>
DBUSER=<userstore key or HANA 2.0: systemdb userstore key>
HANAVERSION="<HANA version: 1.0 or 2.0>"
# optional if rollforward is required
LOGMOUNTPATH="<mounted log backup mount point>"
RECOVERYTIME="2019-03-04 03:11:36"
# do not change below
EXCLUDE_DB_LIST="null"
INCLUDE_DB_LIST="null"
```

For example:

```
DBSID=ipl
DBPORT="HDB01"
HANABACKUPPATH=/iplmnt
DBUSER=ACTBACKUP
HANAVERSION="2.0"
# optional if rollforward is required
LOGMOUNTPATH="/iplmnt_archivelog"
RECOVERYTIME="2019-03-04 03:11:36"
# do not change below
EXCLUDE_DB_LIST="null"
INCLUDE_DB_LIST="null"
```

6. `cd /act/custom_apps/saphana/dump/`
7. Run `ACT_HANADB_newtargetdumprestore.sh`:

```
./ACT_HANADB_newtargetdumprestore.sh
```

or

```
/act/custom_apps/saphana/dump/ACT_HANADB_newtargetdumprestore.sh
```

# 8 Accessing an SAP HANA Database

This section includes:

[Mount a Virtual SAP HANA Database from a Volume-Based Snapshot to a Target SAP HANA Host on page 41](#)

[Workflow to Automate Mount and Refresh of a Virtual Database on page 43](#)

## Mount a Virtual SAP HANA Database from a Volume-Based Snapshot to a Target SAP HANA Host

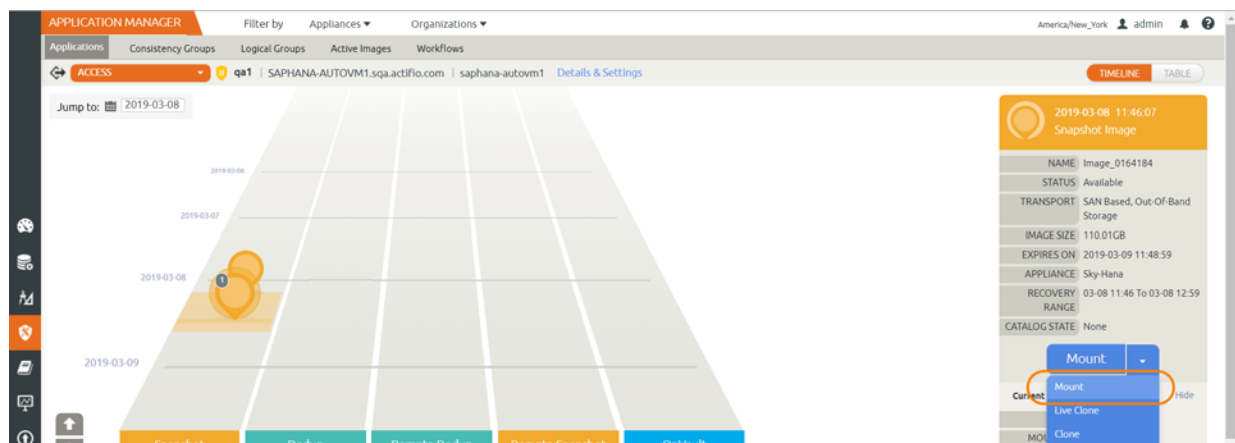
### Before You Begin

To create a virtual application on a target server, an empty HANA instance with required hana sid name must exist on the target. Refer to the HANA documentation for details.

### Procedure

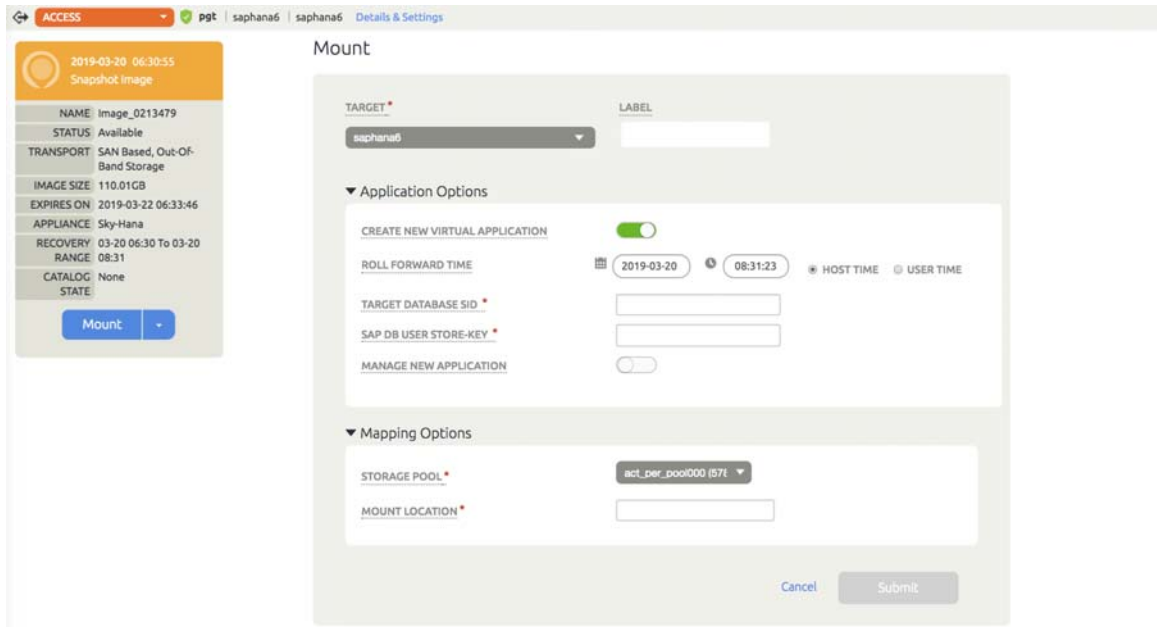
To mount the database image as a virtual application (an application aware mount) to a new target:

1. From App Manager, Applications list, right-click the database and select Access. From the latest snapshot, choose **Mount**.



2. On the Mount page, from Target, choose the desired target HANA server from the dropdown.
3. Under Application Options:
  - o Select **Create New Virtual Application**.
  - o Choose a point in time for a database protected with log roll-forward to recover to.
  - o For Target Database SID, provide the target HANA database name.

- o For SAP DB User Store-Key, provide the hdbuserstore key for the target database (HANA 2.0: SYSTEMDB).
- o For Mount Location, specify a Mount Point to mount to a new target.
- o For Manage New Application, if you want to protect the new database, enable Manage New Application.
- o At Template, choose a template to protect the database.
- o For Profile, choose a profile to use to protect the database.



4. Click **Submit**.

## SAP HANA Database Restore/Recovery License Requirements, Impact on Restore

The license key for an SAP HANA database is based on the system ID and the hardware ID. After a recovery, an SAP HANA license key becomes invalid if the SID or hardware ID has changed.

During recovery, a temporary license key is installed automatically if the backup used for recovery has a permanent license, which is still valid. You can work with the automatically installed temporary license for up to 90 days. During this time, you need to apply to SAP to have the license from the source database transferred to a new license key. You then need to install the new license key in the recovered SAP HANA database.

If the backup that was used for recovery only had a temporary license, the database is in lockdown mode immediately after recovery.

### Source database with temporary license backups taken with temporary licenses

**Restore back to source** – It will be 90 days from the time of database creation and the database will be in lockdown mode.

**Restore to the new target** – It will fail as SAP temp license does not allow the restore to new target.

### Source database with permanent license Backups taken with permanent licenses

**Restore back to source** – no issue

**Restore to the new target** – It will have 90 days trial license. Backup will succeed but they will not be able to use the backup to restore.

## SAP References

<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/19a0f5a85685453080f00087bb9b9c98.html>


<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.02/en-US/bddd0b28bb571014bd9592d247dcd403.html>



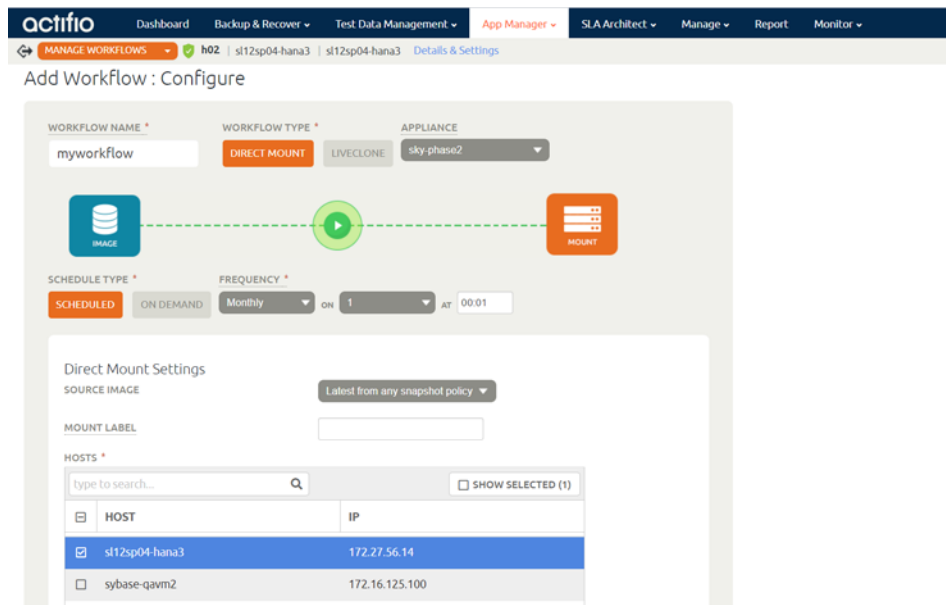
# Workflow to Automate Mount and Refresh of a Virtual Database

To create a workflow to automate the process of mounting and refreshing a database from a snapshot:

1. From the AGM App Manager, Applications list, right-click the HANA database and select **Manage Workflows**.
2. In the upper right corner of the Workflows: Application Dashboard page, click **+ Add Workflow**.



3. Specify:
  - o Workflow Name: Enter a name for this workflow.
  - o Workflow Type: Select Direct Mount.
  - o Schedule Type: Choose Scheduled or On Demand. For a scheduled workflow, specify the frequency as well.



- o Mount Label: (Optional) Specify a mount label for the mounted image.
- o Hosts: Select the target host where the virtual HANA database copy will be created.
- o Mount Location: Specify a mount point for the data volume and log volume of the target.
- o Post-Script: Specify a post script name to be run virtual HANA database copy at the end of refresh. Post scripts are detailed in Network Administrator’s Guide to Actifio VDP.
- o Create New Virtual Application: Enable Create New Virtual Application.
- o Target Database SID: Provide the target HANA database name.
- o SAP DB User Store-Key: Provide the hdbuserstore key for the target database (HANA 2.0: SYSTEMDB).

Optionally, if you want to protect the new virtual database:

- o Manage New Application: Enable Manage New Application.
- o Template: Choose a template to protect the database.
- o Profile: Choose a profile.

4. Click **Add**. This will create an on-demand or scheduled work flow to create or refresh the HANA database virtual copy.



---

# 9 HANA Database Management Using actHANADBМ

---

DBAs and developers can use actHANADBМ.pl to perform database access tasks using the command line interface. ActHANADBМ is a set of Perl scripts that let you automate all essential tasks with a simple language that needs no SSH keys, doesn't store passwords in the clear and takes almost no effort to learn. ActHANADBМ.pl is installed on the database server automatically along with the Actifio Connector.

This section includes:

[Installing and Configuring actHANADBМ.pl](#) on page 46

[actHANADBМ Commands](#) on page 47.

- [agmconfig](#) on page 47

- [createTemplate](#) on page 47

- [hostDiscovery](#) on page 48

- [protectApp](#) on page 49

- [backup](#) on page 49

- [listImageDetails](#) on page 50

- [mount](#) on page 51

- [unmountdelete](#) on page 52

- [restore](#) on page 53

- [runwf](#) on page 54

## Installing and Configuring actHANADB.M.pl

There are four steps to installing and configuring actHANADB.M.pl:

[Installing actHANADB.M.pl with the Actifio Connector](#) on page 46

[Enabling and Verifying Port 443](#) on page 46

### Installing actHANADB.M.pl with the Actifio Connector

The actHANADB.M script library is automatically installed on the host when you install the Actifio Connector. It is available on the host under `/act/custom_apps/saphana/acthanadb.m`. To install the Actifio Connector, see **Network Administrator's Guide to Actifio VDP**.

### Enabling and Verifying Port 443

actHANADB.M uses https port 443 for communication between the host and the appliance. Port 443 should be enabled for the host where the actHANADB.M tool is configured. To test whether the port 443 is enabled, run telnet from the actHANADB.M configured host:

```
telnet <Appliance IP address> 443
```

If port 443 is enabled then the sample output looks like this:

```
[root@zoravmn4 ~]# telnet <Actifio Appliance IP> 443
Trying 172.16.15.200...
Connected to 172.16.15.200.
```

---

**Note:** The escape character is '^'.

---

### Running actHANADB.M.pl

To run the actHANADB.M tool, CD to `/act/custom_apps/saphana/acthanadb.m` folder and invoke `./actHANADB.M.pl`.

To run the script from any other directory, include the script directory in the Perl library path by using the `-I` switch in the command line argument: `perl -I /act/custom_apps/saphana/acthanadb.m/ /act/custom_apps/saphana/acthanadb.m/actHANADB.M.pl`

### Usage of actHANADB.M.pl

When you run actHANADB.M.pl, you must use the `--type` parameter and a type option such as backup:

```
actdbm.pl -type backup
```

The type options for actHANADB.M.pl are:

Usage: actHANADB.M

`--type`

```
<agmconfig>
<createTemplate>
<hostDiscovery>
<protectApp>
<backup>
<listImageDetails>
<mount>
<unmountdelete>
<restore>
<runwf>
```

## actHANADBM Commands

The actHANADBM commands are:

- [agmconfig](#) on page 47
- [createTemplate](#) on page 47
- [hostDiscovery](#) on page 48
- [protectApp](#) on page 49
- [backup](#) on page 49
- [listImageDetails](#) on page 50
- [mount](#) on page 51
- [unmountdelete](#) on page 52
- [restore](#) on page 53
- [runwf](#) on page 54

### agmconfig

#### Storing the Login Credentials for an Actifio Global Manager (agmconfig)

This is one time setup to create and store the Actifio username and password (encrypted). This configuration file is used to access the AGM for invoking different operations using the API.

#### Example

```
perl actHANADBM.pl --type agmconfig
--username <AGM username>
--password <AGM password>
--AGM <AGM IP>
```

**Table 1: agmconfig Parameters**

Parameter	Use
--username	AGM username to access the appliance. This is a required parameter.
--password	Password to access the appliance. This is a required parameter.
--AGM	The name or IP address of the AGM

### createTemplate

To create SLA template, use --type createTemplate

#### Example

```
perl actHANADBM.pl --type createTemplate
--appliancename <appliance name>
--templatename <template name>
[--snappolicyname <Snapshot policy name>]
[--snapRPO <snapshot RPO, default 24 hours>]
--logbackupenable <true|false>
[--logbackupfrequency <Log Backup frequency RPO in minutes>]
[--logbackupretention <Log Backup Retention period in Days>]
[--onVaultPolicyname <onVault policy name>]
[--onVaultRPO <onVault RPO, default 24 hours>]
--profileName <profile name>
--AGM <AGM name|ip>
```

**Table 2: createTemplate Parameters**

Parameter	Use
--applianceName	Sky Appliance name or IP address. This is a required parameter.
--templateName	Name of the SLA template. This is a required parameter.
--snapPolicyName	Name of the Snapshot Policy. this is optional parameter.
--snapRPO	Snapshot Interval. This is optional parameter. Default value 24 hrs.
--logBackupEnable	Enable log backup. This is a required parameter. Input value must be true or false.
--logBackupFrequency	Log backup frequency in minutes. This is optional parameter.
--logBackupRetention	Logbackup retention period (in Days) in Actifio staging disk. This is optional.
--onVaultPolicyName	OnVault policy name. This is optional parameter.
--onVaultRPO	OnVault interval, default 24 hrs. This is optional parameter.
--profileName	Profile name to create the template. This is a required parameter.
--AGM	AGM name or IP address. This is a required parameter.

## hostDiscovery

To discover SAP HANA database host, use --type hostDiscovery

### Example

```
perl actHANADBMS --type hostDiscovery
--applianceName <appliance name>
--hostname <source hostname>
--hostip <source host ip>
--stagingDiskPreference <Type of disk for backup: Block|NFS>
--AGM <AGM name|ip>
```

**Table 3: hostdiscovery Parameters**

Parameter	Use
--applianceName	Name of the appliance. This is a required parameter.
--hostname	Source database hostname. This is a required parameter.
--hostip	Source database host IP. This is a required parameter.
--stagingDiskPreference	Staging disk type for backup, Block or NFS. This is a required parameter.
--AGM	AGM name or IP address. This is a required parameter.

## protectApp

To protect the application, use `-type protectApp`.

### Example

```
perl actHANADBM.pl --type protectApp
--appname <application name>
--hostname <source hostname>
--templatename <Template Name>
--profilename <Profile Name>
--backupType <CBT|filesystemDump>
--hanaSystemDbKey <Hana SystemDB HDB userstore Key Name>
[--hanaTenantDbkeyPrefix <Prefix for Hana Tenant DB Key Name>]
[--volumegrouppreservespace <volume group snap reserve space in percentage: default 20%>]
[--forcefulldbump <true|false>]
[--productionlogretention <production log purging retention in days>]
--AGM <AGM name|ip>
```

**Table 4: protectApp Parameters**

Parameters	Use
<code>--appname</code>	Name of the application to be protected. This is a required parameter.
<code>--hostname</code>	Name of the source host. This is a required parameter.
<code>--templatename</code>	SLA template name to be applied. This is a required parameter.
<code>--profilename</code>	Resource Profile name. This is a required parameter.
<code>--backupType</code>	Type of the backup. CBT or Filesystem Dump. This is a required parameter.
<code>--hanaSystemDbKey</code>	HANA SYSTEM database user store key name. This is a required parameter.
<code>--hanaTenantDbkeyPrefix</code>	Tenant database user store key prefix. This is an optional parameter.
<code>-- volumegrouppreservespace</code>	Volume group snap reserve space. If not specified, default value is 20%. This is an optional parameter.
<code>--forcefulldbump</code>	Force full database dump backup. Input values are true/false. This is optional.
<code>--productionlogretention</code>	Production log retention period in number of days. This is optional.
<code>--AGM</code>	AGM name or IP address. This is a required parameter.

## backup

To create a database backup, use `--type backup, backuptype <db|log|dblog>`.

Use this for:

- db backup
- log backup
- dblog backup

### Example

```
actHANADBM --type backup
```

```

--appname <application name>
[--hostname <hostname>]
[--backuptype <db|log|dblog>]
[--jobpriority <low|medium|high>]
--AGM <AGM name|ip>
[--wait <yes|no>]

```

**Table 5: backup Parameters**

Parameters	Use
--appname	Name of the application. This is a required parameter.
--hostname	Name of the application host. If not specified, host where script is running will be used.
--backuptype	Type of backup operation. This is optional. If not specified, the default is database backup (db)
--jobpriority	The priority for the job. This is optional parameter. Valid inputs are low, medium or high.
--AGM	AGM name or IP address. This is a required parameter.
--wait	Wait until the job completed. This is optional parameter, if not specified default value is yes.

## listImageDetails

To return a list of snapshot images with recovery range for a protected database, use --type listImageDetails

### Example

```

perl actHANADBM .pl--type listImageDetails
--appname <application name>
--hostname <source hostname>
--AGM <AGM name|ip>

```

**Table 6: listImageDetails Parameters**

Parameters	Use
--appname	Name of the application. This is a required parameter.
--hostname	Name of the source host. This is a required parameter.
--AGM	AGM name of IP address. This is a required parameter.



## mount

To mount a backup image or to perform app aware mount, use `--type mount`.

### Example

```
perl actHANADB.M.pl --type mount
--appName <Source Database Name or Source File System Mount Point>
[--image <Image name>]
--sourceHost <source Host Name>
--targetHost <Target Host name>
[--scaleoutnodelist <Scaleout node list seperated by colon>]
[--mountpoint <mount location '/act/mnt'>]
[--appawaremount <true|false default: false>]
[--targetdbuser <Target database Database user store key>]
[--targetdbsid <Target Database SID>]
[--recoverytime <'yyyy-mm-dd hh24:mi:ss'>]
--AGM <AGM name|ip>
[--wait <yes|no>]
```

**Table 7: mount Parameters**

Parameters	Use
<code>--appname</code>	Source application name or Source file system mount point. This is a required parameter.
<code>--image</code>	Name of the image to be mounted. This is optional parameter. If not specified, latest image will be used.
<code>--sourceHost</code>	Name of the source host. This is a required parameter.
<code>--targetHost</code>	Name of the target host. This is a required parameter.
<code>--scaleoutnodelist</code>	In case of cluster, specify the other nodes separated by colon. This is an optional parameter; if not specified, mount will be done only on target node.
<code>--mountpoint</code>	Target mount point name where backup image will be mounted. If not specified, Actifio naming convention will be used.
<code>--appawaremount</code>	Mount and Recover the application on target node. This is optional parameter, if not specified, default value is false.
<code>--targetdbuser</code>	Target database user store key required for recovery. This is optional parameter, required only when <code>--appawaremount</code> is true.
<code>--targetdbsid</code>	Target database SID. This is optional parameter, required only when <code>--appawaremount</code> is true.
<code>--recoverytime</code>	Recovery range to roll forward the logs. Must be specified in the format 'yyyy-mm-dd hh24:mi:ss'. This is optional parameter. If not specified, all available logs will be applied.
<code>--AGM</code>	AGM name or IP address. This is a required parameter.
<code>--wait</code>	Wait until the job completed. This is optional parameter; if not specified default is yes.

## unmountdelete

To perform an unmount and delete operation on an image, use `--type cleanup`. This operation will stop and remove any copy of a database running out of a mounted image and remove the filesystem mount as part of the cleanup.

### Example

```
actHANADBMS --type unmountdelete
  --appName <Source Database Name or Source File System Mount Point>
  --sourceHost <source host name>
  --targetHost <target host name>
  [--imageName <Mounted Image Name>]
  [--targetAppName <Target Cloned Database Name or Target Mounted File System MountPoint>]
  --AGM <AGM name|ip>
  [--wait <yes|no>]
```

**Table 8: unmountdelete Parameters**

Parameters	Use
<code>--appname</code>	Source application name or Source file system mount point. This is a required parameter.
<code>--sourceHost</code>	Name of the source host. This is a required parameter.
<code>--targetHost</code>	Name of the target host. This is a required parameter.
<code>--image</code>	Name of the image to be mounted. This is optional parameter. If not specified, latest image will be used.
<code>--targetAppName</code>	Target application name or Target mounted file system mountpoint.
<code>--AGM</code>	AGM name or IP address. This is a required parameter.
<code>--wait</code>	Wait until the job has completed. This is optional parameter; if not specified default value is yes.

## restore

To restore back to source server, use the option `--type restore`.

### Example

```
perl actHANADBMS --type restore
--appName <Source Database Name or Source File System Mount Point>
[--image <Image name>]
--sourceHost <source Host Name>
--targetHost <Target Host name>
[--recoverytime <'yyyy-mm-dd hh24:mi:ss'>]
[--excludedblist <Exclude DB list name seperated by name>]
[--includedblist <Include DB list name seperated by comma>]
--AGM <AGM name|ip>
[--wait <yes|no>]
```

**Table 9: restore Parameters**

Parameters	Use
<code>--appname</code>	Source application name or Source file system mount point. This is a required parameter.
<code>--image</code>	Name of the image to be mounted. This is an optional parameter. If not specified, latest image will be used.
<code>--sourceHost</code>	Name of the source host. This is a required parameter.
<code>--targetHost</code>	Name of the target host. This is a required parameter.
<code>--recoverytime</code>	Recovery range to roll forward the logs. Must be specified in the format 'yyyy-mm-dd hh24:mi:ss'. This is optional parameter. If not specified, all the available logs will be applied.
<code>--excludedblist</code>	Specify the list of databases separated by comma to exclude during the restore operation. This is an optional parameter; if not specified all tenant databases will be restored.
<code>--includedblist</code>	Specify the list of databases separated by comma to include during the restore operation. This is an optional parameter, if not specified all tenant databases will be restored.
<code>--AGM</code>	AGM name or IP address. This is a required parameter.
<code>--wait</code>	Wait until the job completed. This is an optional parameter; if not specified, the default is yes.

## runwf

Run Workflow creates a new database copy or refreshes an existing database copy based on the re-provision option. To run a workflow, use `--type runwf`,

### Example

```
actHANADB.M.pl --type runwf
--appName <source database name>
--hostname <sourcehostname>
--wfname <workflow name>
--reprovision <yes|no>
[ --image <Image name>]
--AGM <AGM name|ip>
[ --wait <yes|no>]
```

**Table 10: runwf Parameters**

Parameters	Use
<code>--appname</code>	Source application name or Source file system mount point. This is a required parameter.
<code>--sourceHost</code>	Name of the source host. This is a required parameter.
<code>--wfname</code>	Name of the workflow. This is a required parameter.
<code>--reprovision</code>	Reprovision flag to indicate new application aware mount or reprovision application aware mount.
<code>-image</code>	Image name to use for provision the database. This is an optional parameter. If not specified, the latest image will be used for database provision.
<code>--AGM</code>	AGM name or IP address. This is a required parameter.
<code>--wait</code>	Wait until the job completed. This is an optional parameter; if not specified default value is yes.