
SAP HANA DBA's Guide to the Actifio Virtual Data Pipeline

Copyright, Trademarks, and other Legal Matter

Copyright © 2009 - 2019 Actifio, Inc. All rights reserved.

Actifio[®], AnyIT[®], Dedup Async[®], OnVault[®], Enterprise Data-as-a-Service[®], FlashScan[®], AppFlash DEVOPS Platform[®], Copy Data Cloud[®], and VDP[®] are registered trademarks of Actifio, Inc.

Actifio Sky[™], Actifio One[™], and Virtual Data Pipeline[™] are trademarks of Actifio, Inc.

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

Actifio, Inc., is a provider of data protection and availability products. Actifio's technology is used in products sold by the company and products and services sold and offered by its commercial partners. The current list of Actifio patents is available online at: <http://www.actifio.com/patents/>

Actifio believes the information in this publication is accurate as of its publication date. Actifio reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." ACTIFIO, INC. MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This software and the associated documentation are proprietary and confidential to Actifio. Use, copying, and distribution of any Actifio software described in this publication requires an applicable software license. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

Actifio strives to produce quality documentation and welcomes your feedback. Please send comments and suggestions to docs@actifio.com.

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	5
SAP HANA Backup Methods	6
References	7
Chapter 2 - Preparing the SAP HANA 1.0 Database	9
Creating the Database User Account.....	9
Get the SQL Port ID	11
Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 (single container system)	11
Chapter 3 - Preparing a HANA 2.0 Database	13
Creating the System Database and Tenant Database Users	13
Creating the System Database User Account from HANA STUDIO	13
Creating the User under the Tenant DB.....	15
Getting the Instance and SQL Port Numbers.....	16
Creating the SAP HANA Hdbuserstore Key	16
Creating the SAP HANA Hdbuserstore Key for the System Database and Each Tenant Database in a Single Node System	16
Creating the SAP HANA Hdbuserstore Key for the System Database and each Tenant Database in a Scale-Out Multi-Node SAP HANA System	17
Chapter 4 - Adding a SAP HANA Database Host and Discovering the Database	19
Adding the Host from the Domain Manager	19
Discovering the HANA Database Application from the Application Manager.....	21
Finding the Discovered HANA Database in the Application Manager	22
Chapter 5 - Configuring the SAP HANA Backup Method	23
Ensure that the Disk Preference on the Host is Set Correctly	24

Ensure that the Backup Capture Method in the Application Settings is Set Correctly	26
Chapter 6 - Protecting the HANA Database	29
Chapter 7 - Protecting SAP HANA Database Logs	31
Setting up the Log Mode and Log Backup in HANA Studio.....	31
Setting up the Log Backup in Actifio AGM.....	33
Chapter 8 - Restoring and Recovering an SAP HANA Database	37
Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application	37
Restoring and Recovering a SAP HANA Database.....	38
Recovering from Block-Based LVM Snapshot with CBT.....	38
Recovering from a File-Based Backup with NFS.....	39

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.

Your Actifio appliance's Documentation Library contains detailed, step-by-step, application-specific instructions on how to protect and access your data. Each guide is in PDF format and may be viewed online, downloaded, or printed on demand. The following guides will be of particular interest:

- ***Introducing Actifio Copy Data Management***
- ***Connecting Hosts to Actifio Appliances***

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
- 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

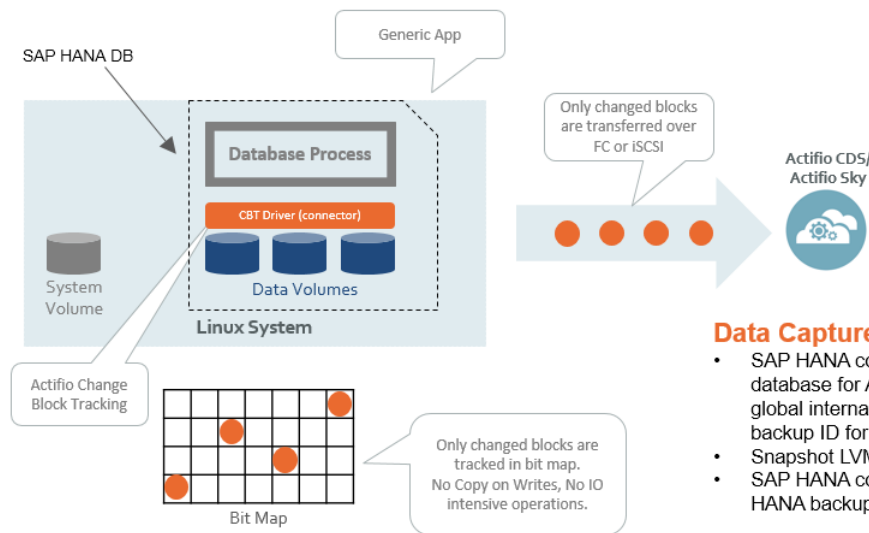
[Introduction to Actifio SAP HANA Administration](#) on page 5

[SAP HANA Backup Methods](#) on page 6

[References](#) on page 7

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.



Data Capture

- SAP HANA commands to prepare database for Actifio snapshot (creates a global internal savepoint with SAP HANA backup ID for data backup)
- Snapshot LVM
- SAP HANA commands to associate SAP HANA backup ID to Actifio snapshot ID

SAP HANA for LVM with Linux Change Block Tracking

Application data is captured at the block level, in application native format, according to a specified SLA. A Golden copy of that data is created and stored once, 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 simple 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.
- Encrypts data using the AES-256 encryption standard. Authentication between Actifio appliances is performed using 1024-bit certificates.

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

The Actifio appliance can instantly present a copy of the database rolled forward to a specific point of time. The roll forward operation is performed from the Actifio Desktop (user interface).

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. Application Aware mounts are described in [To mount the database image as a virtual application \(an application aware mount\) to a new target](#): on page 37.

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 8, Restoring and Recovering an SAP HANA Database](#).

Workflows

While SLAs govern the automated *capture* of a production database, 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 application aware) 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.

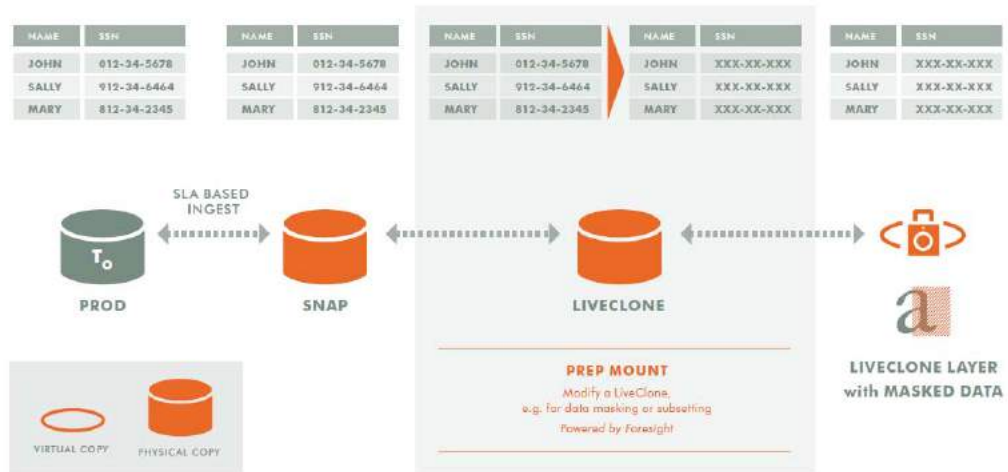
Combining Actifio's automated data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now, instead of having to wait for DBAs to update test and development environments, users can provision their own environments almost instantly.

For example, an Actifio administrator can create an SLA Template Policy that captures data according to a specified schedule. Optionally, the administrator can mark the captured production data as sensitive and only accessible by users with the proper access rights.

After access rights have been defined and data has been captured, the administrator 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 with proper access can provision their environments with the LiveClone or mountable data via the Actifio Desktop.



Workflow With Masked Social Security Data

Introduction to Actifio SAP HANA Administration

Actifio can virtualize and protect:

- **Single Container system (HANA 1.0) Dedicated:** In single-container system the system database and tenant database are perceived as a single unit and are therefore 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.

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		Yes	Yes
Scale-Out MDC: Multiple-Container Systems (HANA 2.0) with one or more tenant databases			Yes
Scale-Out MDC Local HA (N Active Host + 1 or More Standby Nodes)			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 handled automatically in all options and integrated with database backup policies

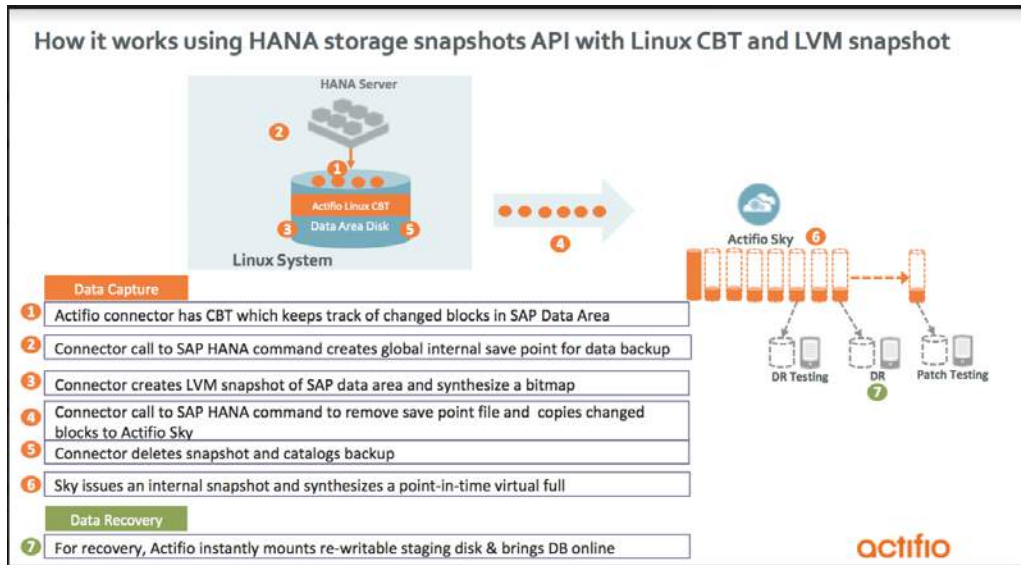
SAP HANA Backup Methods

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

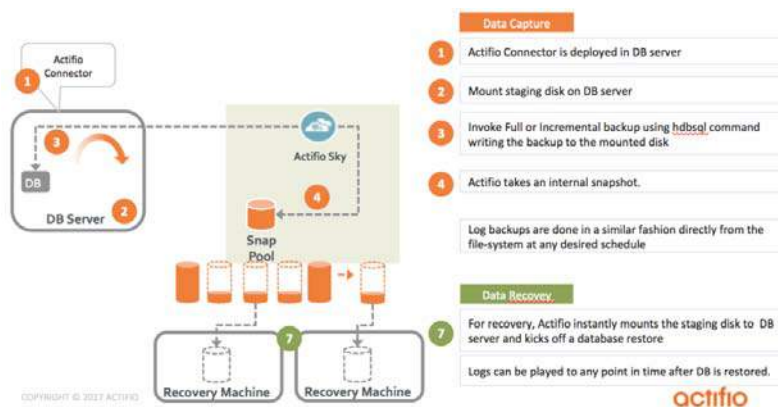
The SAP HANA database creates a database internal snapshot based on a system wide save point executed during the PREPARE step. The database internal snapshot is stored in the data volumes area.



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
2. Storage Snapshots: 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¤t_toc=/en/00/0ca1e3486640ef8b884cdf1a050fbb/plain.htm&node_id=773&show_children=false

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 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 userstore key needs to be created for a single container under database.

Preparing the HANA 1.0 database requires:

[Creating the Database User Account](#) on page 9

[Get the SQL Port ID](#) on page 11

[Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 \(single container system\)](#) on page 11

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.

Naming Convention for Database User Account

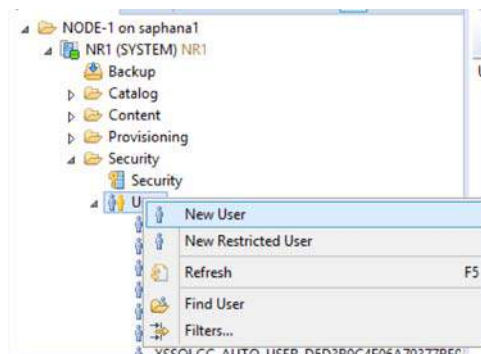
Actifio standard naming convention is DATABASE BACKUP USERNAME: Choose a database user name based on company's standard. Make sure to create this user account under SYSTEMDB.

For example if the DATABASE BACKUP USERNAME is ACTBACKUP then this must be created under SYSTEMDB.

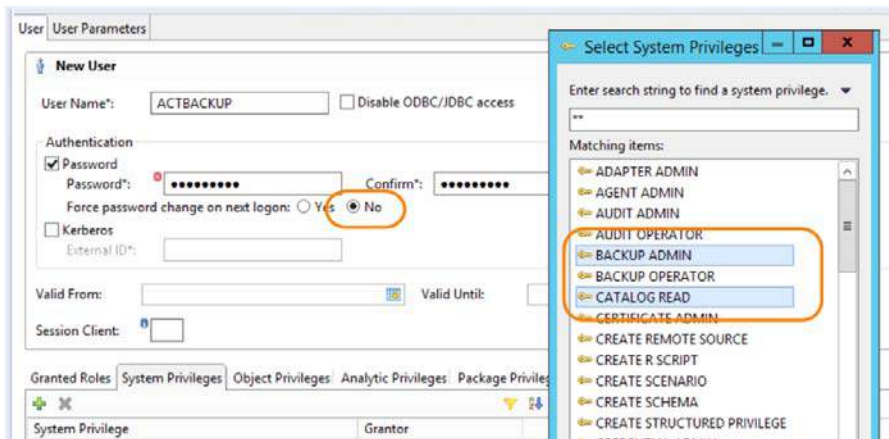
Procedure

To create the user:

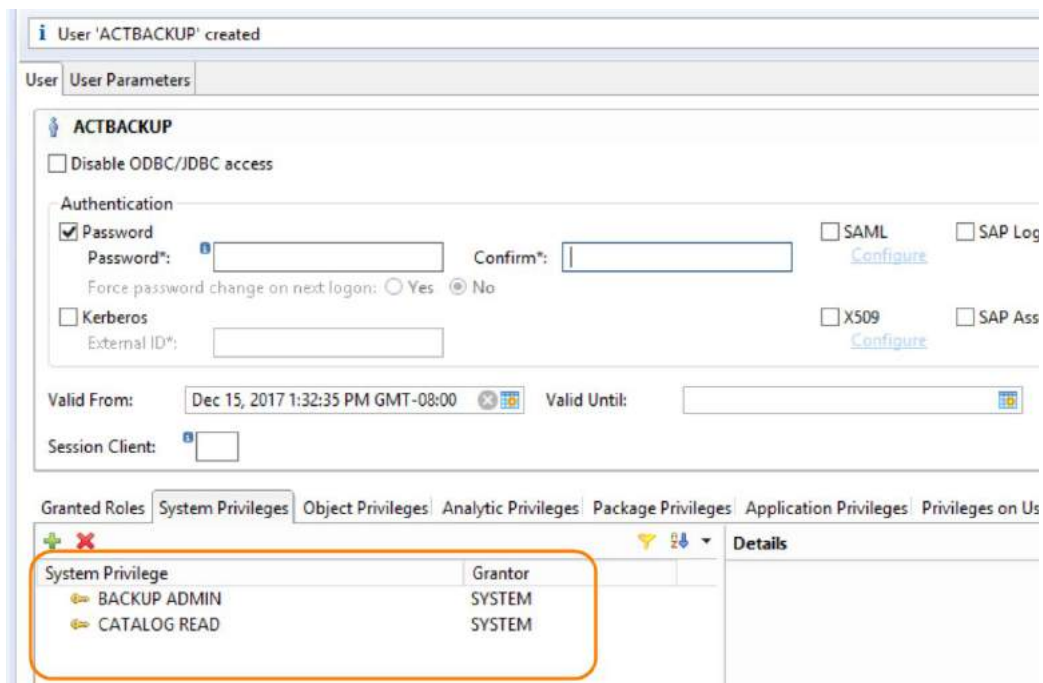
1. From SAP HANA Studio go to System > Security > Users > New User.



2. Assign a user name and a password.
3. Select 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.

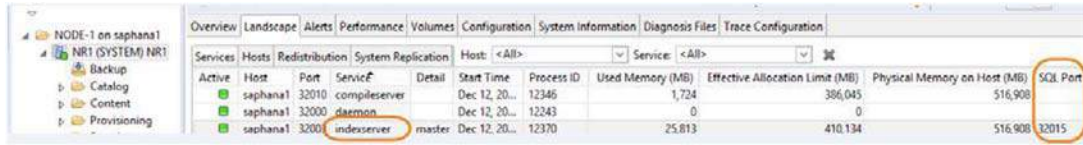


You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN and CATALOG READ privileges.



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.



The screenshot shows the SAP HANA Studio Landscape view. The 'Services' tab is active, displaying a table of services. The 'indexserver' service is highlighted with a red circle. The 'SQL Port' column for this service is also highlighted with a red circle, showing the value 32015.

Active	Host	Port	Service	Detail	Start Time	Process ID	Used Memory (MB)	Effective Allocation Limit (MB)	Physical Memory on Host (MB)	SQL Port
✓	saphana1	32010	compileserver		Dec 12, 20...	12346	1,724	386,045	516,908	
✓	saphana1	32000	daemon		Dec 12, 20...	12243	0	0	0	
✓	saphana1	32001	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

For SYSTEMDB 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: 32015
- Hostname : saphana3

```
./hdbuserstore SET ACTBACKUP saphana3:32015 ACTBACKUP <database backup user password>  
*****>
```

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

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 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 13

[Getting the Instance and SQL Port Numbers](#) on page 16

[Creating the SAP HANA Hdbuserstore Key](#) on page 16

Creating the System Database and Tenant Database Users

[Creating the System Database User Account from HANA STUDIO](#) on page 13

[Creating the User under the Tenant DB](#) on page 15

Creating the System Database User Account from HANA STUDIO

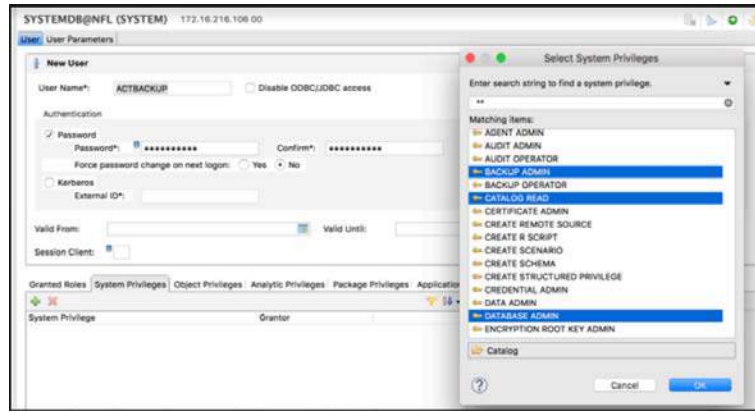
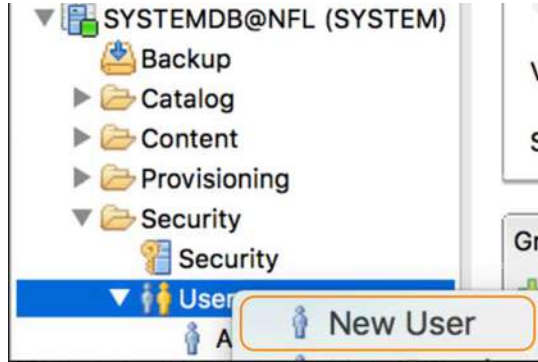
Naming convention for database user account

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


Procedure

To create the system database user account:

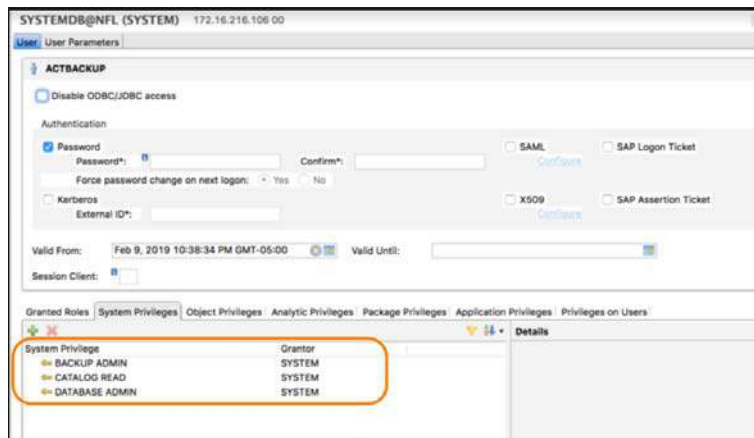
1. Create the USER under SYSTEMDB
 - o Assign a User Name and a Password.
 - o Select Force password change on next logon to No.
 - o Click on the System Privilege tab and assign privileges by selecting BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN
 - o From SAP HANA Studio SYSTEMDB, go to System >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



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



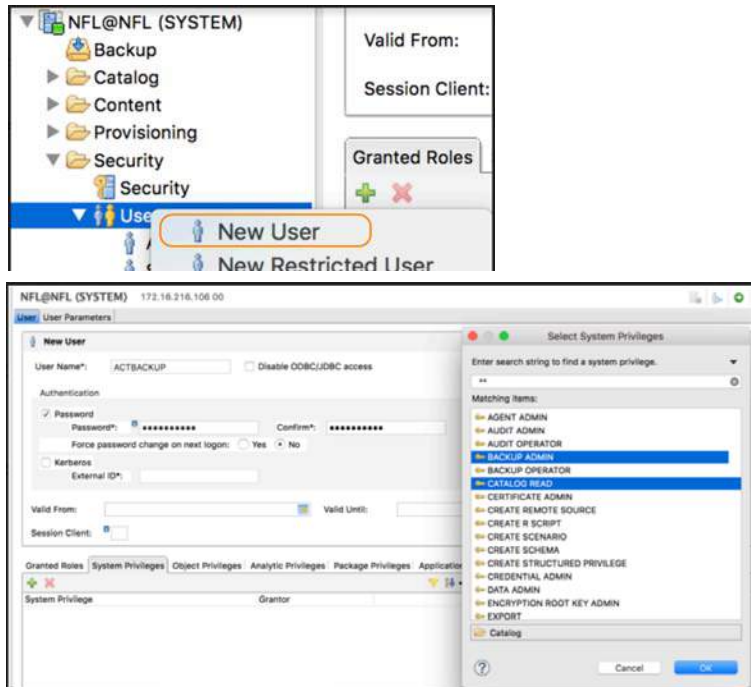
You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN privileges.



Creating the User under the Tenant DB

To create the tenant database user account:

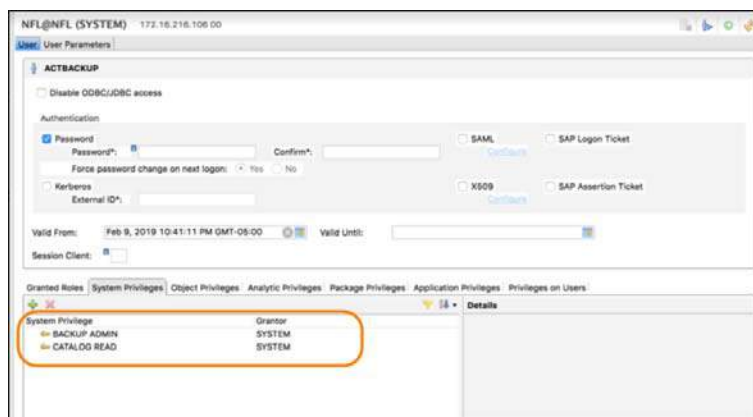
1. Create the USER under TENANTDB
 - o Assign a User Name and a Password.
 - o Select Force password change on next logon to No.
 - o Click on the System Privilege tab and assign privileges by selecting BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN
 - o From SAP HANA Studio SYSTEMDB, go to TENANTDB > 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	30006	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	compileserv		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 M	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				
saphana-autovm6	30001	nameserver	master	Mar 13,...	14989			4,925	4,925	2...	40,...	
saphana-autovm6	30010	compileserv		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 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 16](#)

[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 17](#)

Hdbuserstore Key Naming Convention

For SYSTEMDB set the key name = DATABASE BACKUP USERNAME.

For TENANTDB set the key name = DATABASE BACKUP USERNAME<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 postfix 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 *****>
```



# 4 Adding a SAP HANA Database Host and Discovering the Database


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

1. [Adding the Host from the Domain Manager](#) on page 19
2. [Discovering the HANA Database Application from the Application Manager](#) on page 21
3. [Finding the Discovered HANA Database in the Application Manager](#) on page 22

## Adding the Host from the Domain Manager

Add the host to Domain Manager. If the host is already added then edit the host and make sure to set the Disk Preference correctly.

1. From the AGM Domain Manager, Hosts tab, click +Add Host.



2. On the Add Host page:
  - o Name: Provide the HANA database server name.
  - o IP Address: Provide the HANA database server IP and click the + sign on the right corner.
  - o Appliances: Select the check box for the appliance.
  - o Host Type: Make sure this is Generic.
  - o Click Add at bottom right to add the host.

The Host will get added.

3. Right-click the host and select Edit.
4. On the Edit Host page: Select the disk preference:
  - o For block-based backup with CBT: select **Block**
  - o For file-based backup with Full+Incremental file system backup: select **NFS**

DOMAIN MANAGER Filter by Appliances Organizations

Appliances Users Roles Organizations LDAP Hosts Storage Pools

**SAPAHAN-AUTOVM4.sqa.actifio.com**

IP 172.16.216.104

FRIENDLY PATH SAPHANA-AUTOVM4

UNIQUE NAME e21a8a21-e519-42c0-8efb-9a4ac64b0f75\_6778

OS RELEASE Red Hat Enterprise Linux

OS VERSION 3.10.0-514.26.2.el7.x86\_64

OS TYPE Linux

DISK PREFERENCE BLOCK

### Edit Host

Name \* SAPAHAN-AUTOVM4.sqa.actifio.com

Friendly Name SAPHANA-AUTOVM4

IP Address \*

Description

Appliances \*   
   

| APPLIANCE                                    | IP            |
|----------------------------------------------|---------------|
| <input type="checkbox"/> saphana-remote      | 172.16.200.22 |
| <input checked="" type="checkbox"/> sky-hana | 172.16.201.44 |

Host Type Generic

Disk Preference NFS   
 Block   
 NFS

Enable Auto Discovery

5. Select save at the bottom of Edit Host page

## Discovering the HANA Database Application from the Application Manager

To discover the HANA database:

1. From the AGM Application Manager, Applications tab, select Add Application in the upper right corner.
2. On the Add Application page, select Discover connector supported applications and Using existing host, then select the HANA database host. If you have many hosts, you can use the search feature or use the filter to see only hosts that are managed by a specific Actifio appliance.

| Host                           | IP             | Friendly Path   | Appliance |
|--------------------------------|----------------|-----------------|-----------|
| SAPHANA-AUTOVM4.sqa.actifio... | 172.16.216.104 | SAPHANA-AUTOVM4 | sky-hana  |
| SAPHANA-AUTOVM3.SQA.ACTI...    | 172.16.216.103 | SAPHANA-AUTOVM3 | sky-hana  |
| SAPHANA-AUTOVM2.sqa.actifio... | 172.16.216.102 | SAPHANA-AUTOVM2 | sky-hana  |
| SAPHANA-AUTOVM1.sqa.actifio... | 172.16.216.101 | saphana-autovm1 | sky-hana  |

3. Select the host and click Add Applications in the bottom right corner. This will run the discovery on the HANA database host and will discover all HANA databases running on it.

| Host                           | IP             | Friendly Path   | Appliance |
|--------------------------------|----------------|-----------------|-----------|
| SAPHANA-AUTOVM4.sqa.actifio... | 172.16.216.104 | SAPHANA-AUTOVM4 | sky-hana  |
| SAPHANA-AUTOVM3.SQA.ACTI...    | 172.16.216.103 | SAPHANA-AUTOVM3 | sky-hana  |
| SAPHANA-AUTOVM2.sqa.actifio... | 172.16.216.102 | SAPHANA-AUTOVM2 | sky-hana  |
| SAPHANA-AUTOVM1.sqa.actifio... | 172.16.216.101 | saphana-autovm1 | sky-hana  |

## Finding the Discovered HANA Database in the Application Manager

To find the newly-discovered database, go to the AGM Application Manager Applications tab. All applications known to the AGM of all types are listed. Use the Type application filter on left pane to show only SAP HANA databases.

The new HANA database will appear in the list as unmanaged (the red shield icon).

The screenshot shows the Application Manager interface. The left sidebar contains a filter menu with 'SAP HANA' selected. The main area displays a table of applications. The 'sl1' application is highlighted with a red shield icon, indicating it is unmanaged.

|                          | APPLICATION | TEMPLATE            | PROFILE      | FRIENDLY PATH    | HOST NAME        | APPLIANCE       |
|--------------------------|-------------|---------------------|--------------|------------------|------------------|-----------------|
| <input type="checkbox"/> | ha6         | HANABackup          | LocalProfile | saphana-autovm10 | saphana-autovm10 | SAK-SKY-upgrade |
| <input type="checkbox"/> | has         | HANADBTemplate1     | LocalProfile | saphana-autovm11 | saphana-autovm11 | saphanasky      |
| <input type="checkbox"/> | ipl         | TESTSAPHANATEMPL... | LocalProfile | saphana-autovm5  | saphana-autovm5  | SAK-SKY-upgrade |
| <input type="checkbox"/> | md1         | SAPHANALogSmart     | LocalProfile | md1_cluster      | md1_cluster      | SAK-SKY-upgrade |
| <input type="checkbox"/> | nfl         | SAPHANALogSmart     | LocalProfile | saphana-autovm6  | saphana-autovm6  | SAK-SKY-upgrade |
| <input type="checkbox"/> | pqt         | SAPHANALogSmart     | LocalProfile | saphana6         | saphana6         | SAK-SKY-upgrade |
| <input type="checkbox"/> | sl1         |                     |              | Hana-Sles        | Hana-Sles        | saphana-remote  |

# 5 Configuring the SAP HANA Backup Method

You can back up the database:

- Using Block-Based Database Storage Snapshots with CBT
- Using File-Based Traditional Backup and Recovery in NFS

| Setting                                     | Block-Based LVM Snapshot with CBT                                                                                                                                                                                                                                                                                                                  | File-Based Backup in NFS                                                                  |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Percentage of Reserve Space in Volume Group | This is needed for LVM snapshot temporary space. Recommended value is 20%                                                                                                                                                                                                                                                                          | Not applicable                                                                            |
| Backup Capture Method                       | Use Changed block tracking based backup                                                                                                                                                                                                                                                                                                            | Use full+incremental filesystem backup                                                    |
| Force Full Filesystem Backup                | Not applicable                                                                                                                                                                                                                                                                                                                                     | Use for an ad hoc full backup                                                             |
| Database Filesystem Staging Disk Size in GB | Not applicable                                                                                                                                                                                                                                                                                                                                     | Use the default calculation: (database size * 1.5)+ 10%. The disks will grow dynamically. |
| Log Backup Staging Disk Size in GB          | By default Actifio calculates this as daily log generation * retention of log backup SLA plus 20% buffer. Default is recommended.<br>Providing a value will override the default calculation and the log disk will not grow dynamically. This will become a fixed size                                                                             |                                                                                           |
| Retention of Production DB Logs in Days     | This value 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 the log will be purged older than the data backup id. Default value is 0 days. With default value all logs prior to last data backup will be purged. |                                                                                           |
| HANA DB User Store Key                      | This is the SAP HANA hdbuserstore key for the system database created in earlier. This field is mandatory.                                                                                                                                                                                                                                         |                                                                                           |
| Script Timeout                              | This value is applied to internal backup and recovery scripts called by connector. Default value is recommended.                                                                                                                                                                                                                                   |                                                                                           |

File-based backup also requires that the CLI command DB dump schedule be configured. See **Actifio CLI Reference**.

Whichever method you select, you must:

[Ensure that the Disk Preference on the Host is Set Correctly](#) on page 24

[Ensure that the Backup Capture Method in the Application Settings is Set Correctly](#) on page 26

## Ensure that the Disk Preference on the Host is Set Correctly

Choose between:

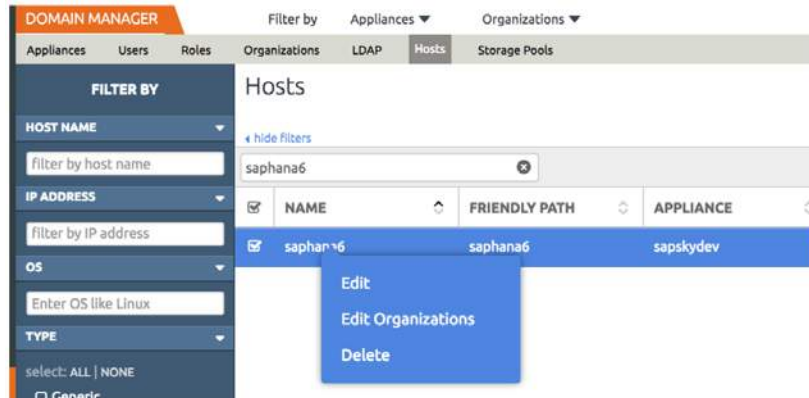
[Setting Disk Preference for Block-Based Database Storage Snapshots with CBT](#) on page 24

[Setting Disk Preference for File-Based Traditional Backup and Recovery in NFS](#) on page 25

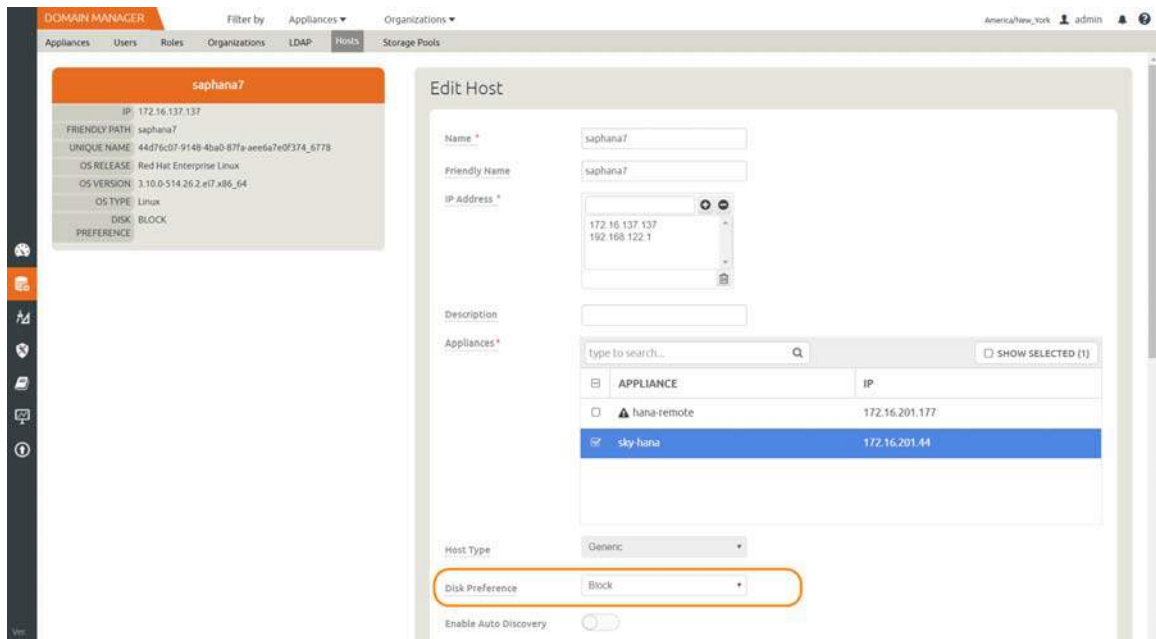
### Setting Disk Preference for Block-Based Database Storage Snapshots with CBT

To set disk preference for block-based database storage snapshots with CBT:

1. From AGM Domain Manager, Hosts tab, right-click the host and select Edit.



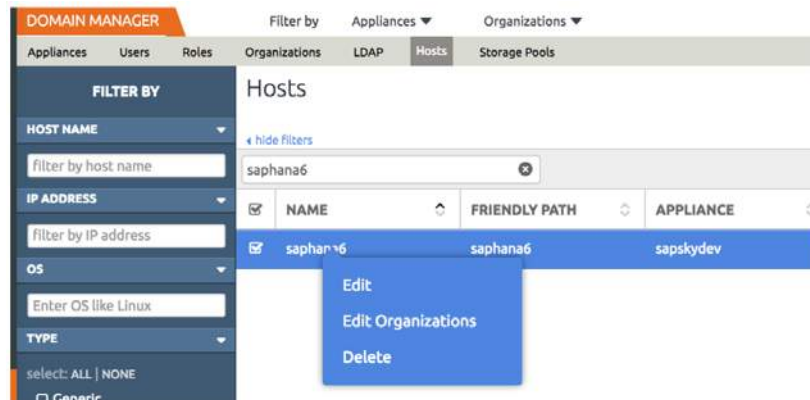
2. In the Edit Host pane, set Disk Preference to *Block* and click Save at the bottom of the page.



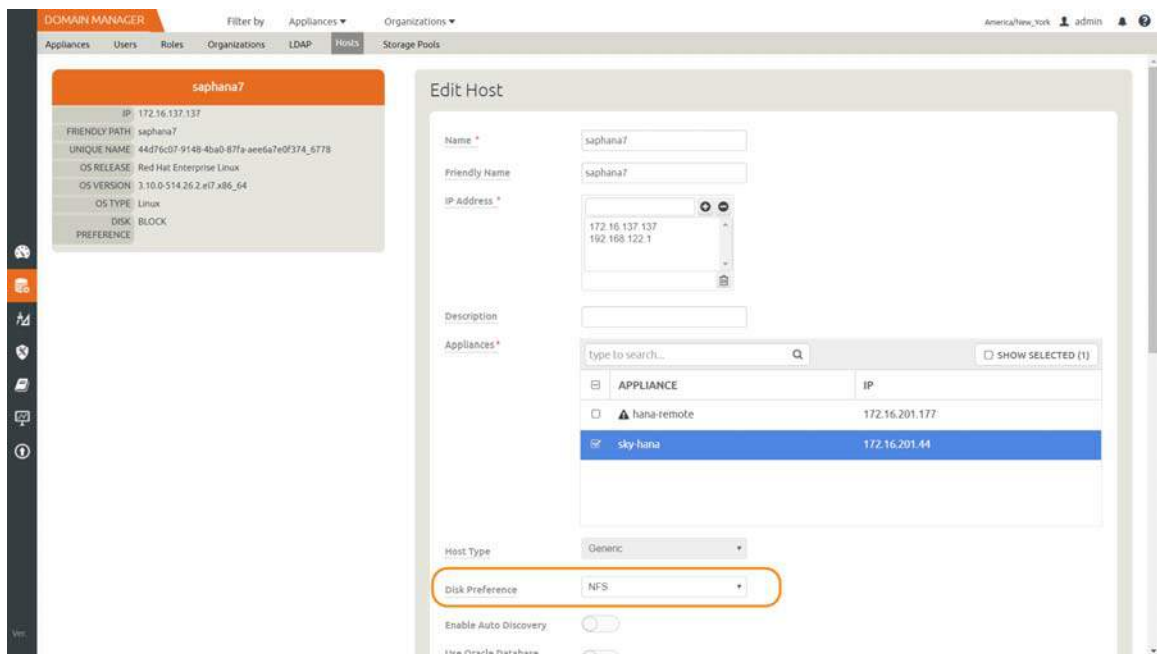
## Setting Disk Preference for File-Based Traditional Backup and Recovery in NFS

To set disk preference for File-Based Traditional Backup and Recovery in NFS:

1. From AGM Domain Manager, right-click the host and select Edit.



2. In the Edit Host pane, set Disk Preference to *NFS* and click Save at the bottom of the page.



## Ensure that the Backup Capture Method in the Application Settings is Set Correctly

Choose between:

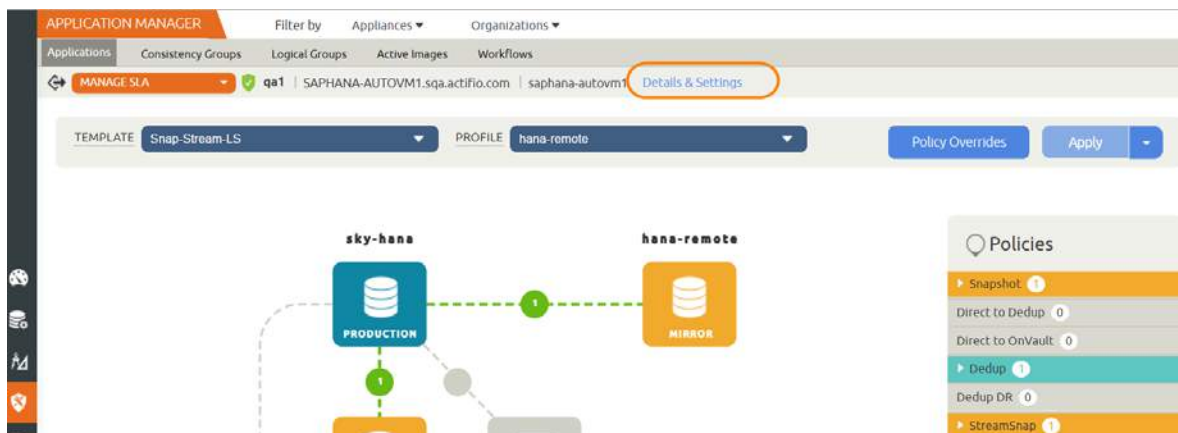
[Setting Backup Capture Method for Block-Based Database Storage Snapshots with CBT](#) on page 26

[Setting Backup Capture Method for File-Based Backup and Recovery in NFS](#) on page 27

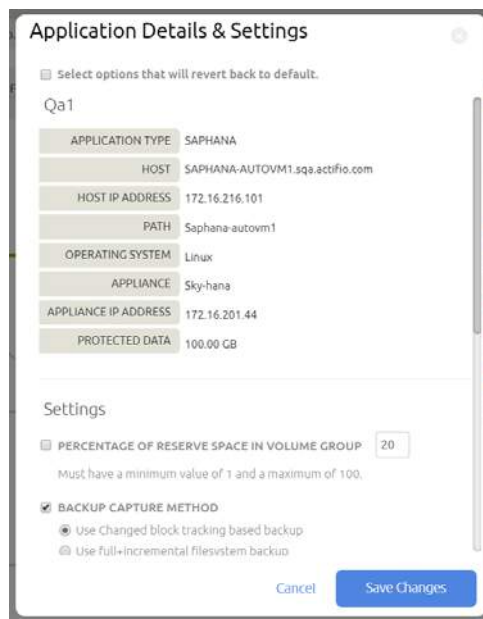
### Setting Backup Capture Method for Block-Based Database Storage Snapshots with CBT

To set the backup capture method for block-based database storage snapshots with CBT:

1. Go to the Application Manager. In the Applications tab, right-click the application and select Manage SLA. At the top of the page, click the blue Details & Settings link.



2. Set the Backup Capture Method to Use Changed block tracking based backup and click Save Changes. For details on the other settings, see [Configuring the SAP HANA Backup Method](#) on page 23.





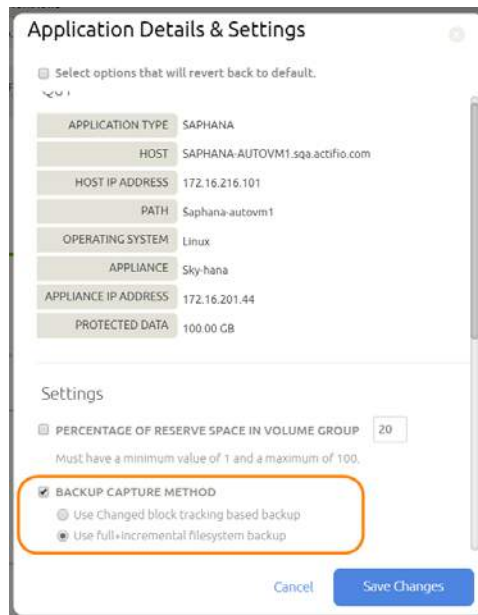
## Setting Backup Capture Method for File-Based Backup and Recovery in NFS

To set the backup capture method for file-based backup and recovery in NFS:

1. Go to the Application Manager. In the Applications tab, right-click the application and select Manage SLA. At the top of the page, click the blue Details & Settings link.



2. Set the Backup Capture Method to Use full+incremental filesystem backup and click Save Changes. For details on the other settings, see [Configuring the SAP HANA Backup Method](#) on page 23.

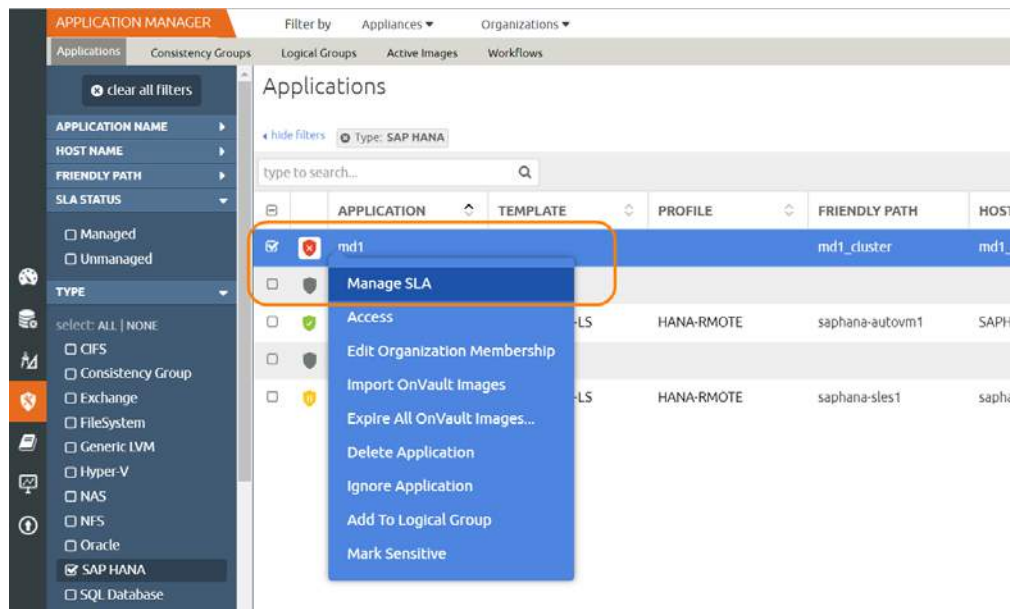




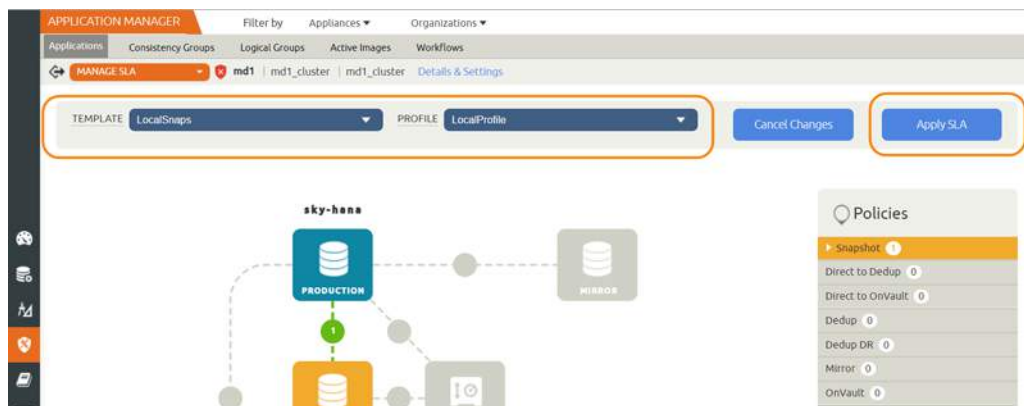
# 6 Protecting the HANA Database

To protect the database:

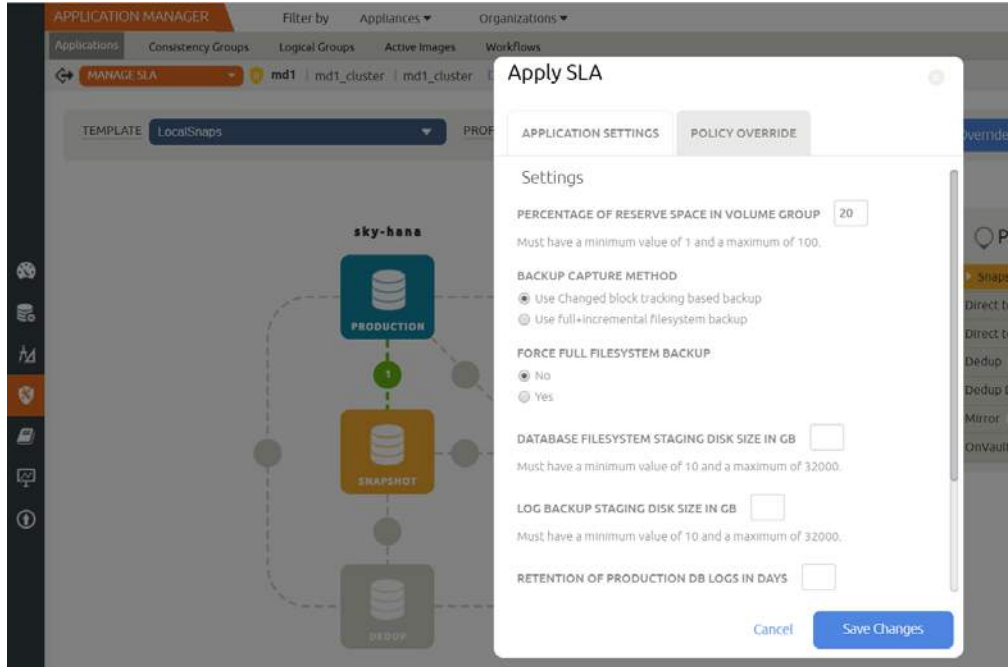
1. Right-click the HANA database and select Manage SLA.



2. On the Manage SLA page, select your desired template and profile, then click Apply SLA.



3. On the Apply SLA page, fill in the required field based on type of backup as detailed in [Configuring the SAP HANA Backup Method](#) on page 23. Click Save Changes.



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.

# 7 Protecting SAP HANA Database Logs

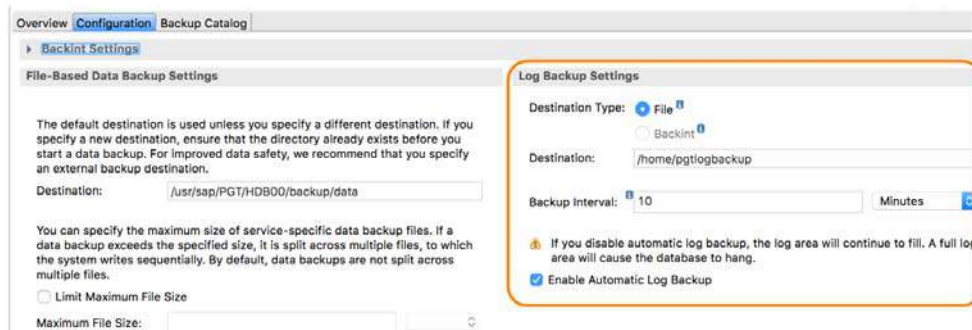
There are two parts to configuring protection of SAP HANA database logs:

[Setting up the Log Mode and Log Backup in HANA Studio on page 31](#)

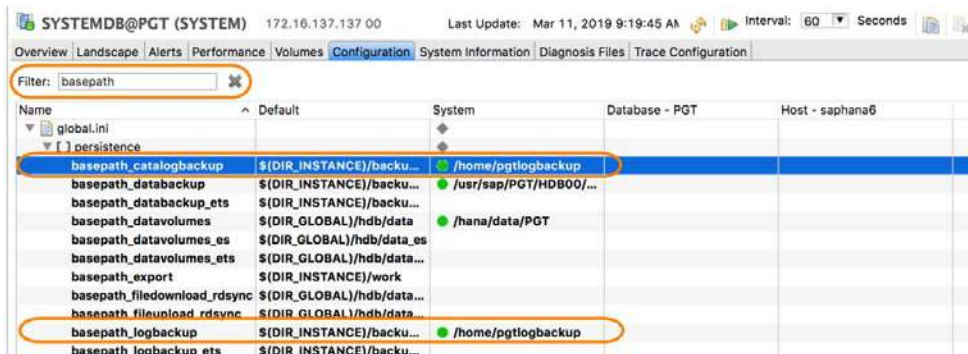
[Setting up the Log Backup in Actifio AGM on page 33](#)

## 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 DATABASE (SYSTEMDB FOR HANA 2.0) - Backup - Configuration page
  - a. Destination Type is File.
  - a. Destination is set to a local file system mount path.
  - a. Backup Interval is set to required RPO.
  - a. Automatic Log Backup is enabled.



2. Check under Database configuration: DATABASE (SYSTEMDB FOR HANA 2.0) - Configuration page. In the filter, type basepath.



3. Verify basepath\_logbackup is set correctly:
  - a. Set the basepath\_catalogbackup to same as basepath\_logbackup.

- o Open the basepath\_catalogbackup edit page.
- o Set the New Value to 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

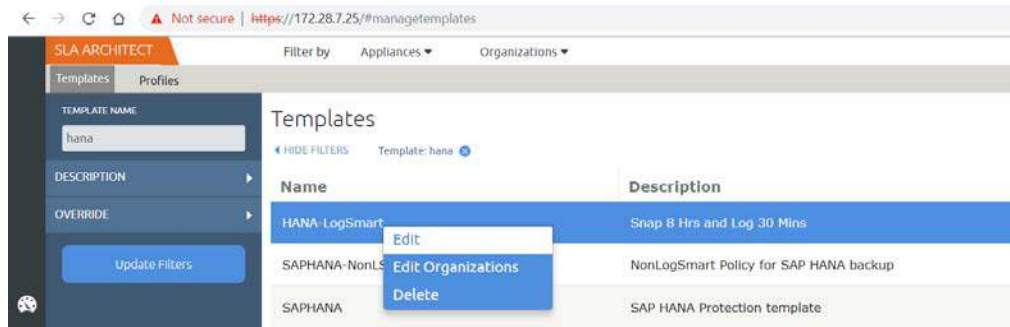
Active Value:

New Value:

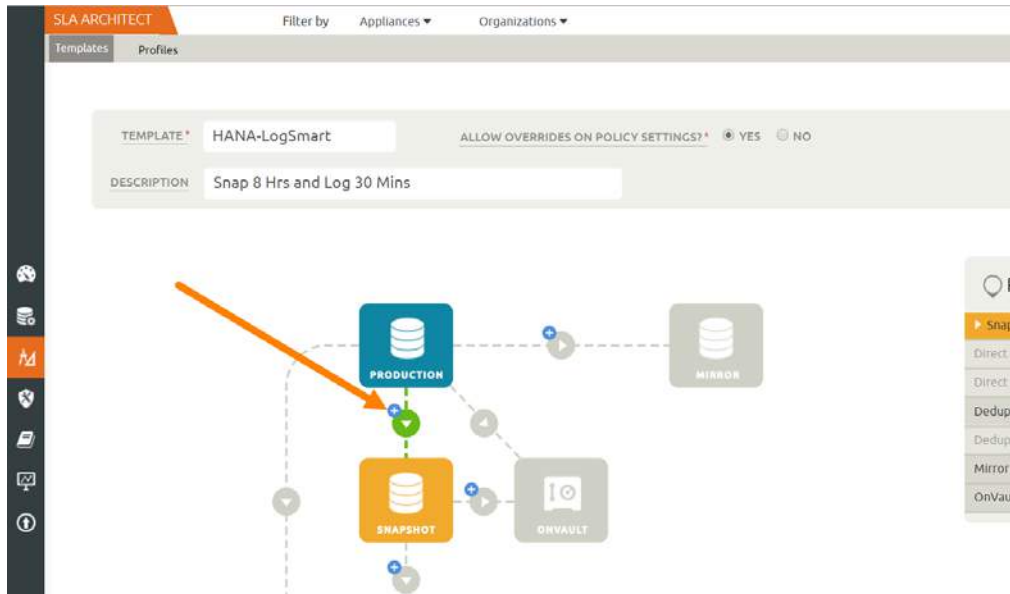
## Setting up the Log Backup in Actifio AGM

To enable and set up the HANA database log backup:

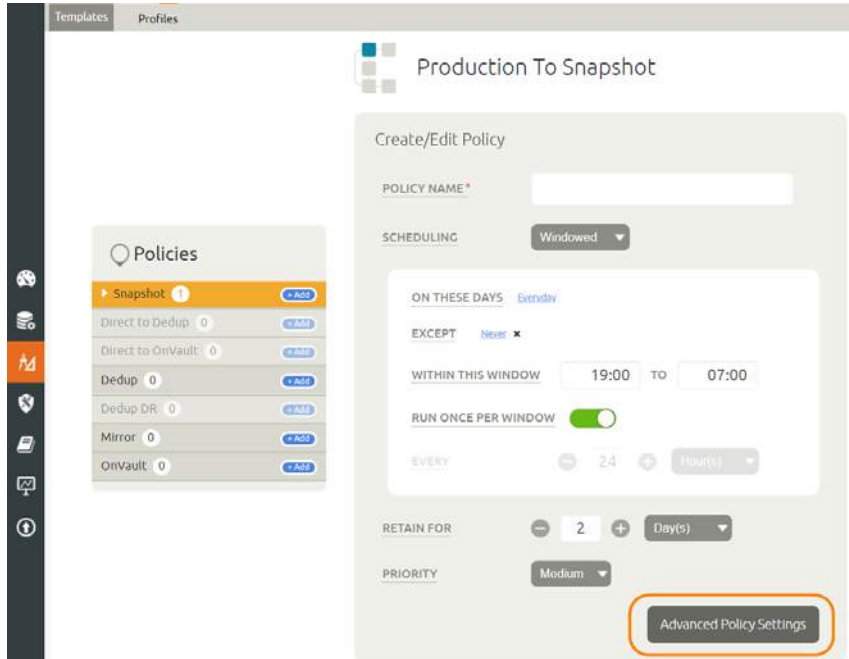
1. From the SLA Architect page, edit the template created for HANA database protection:



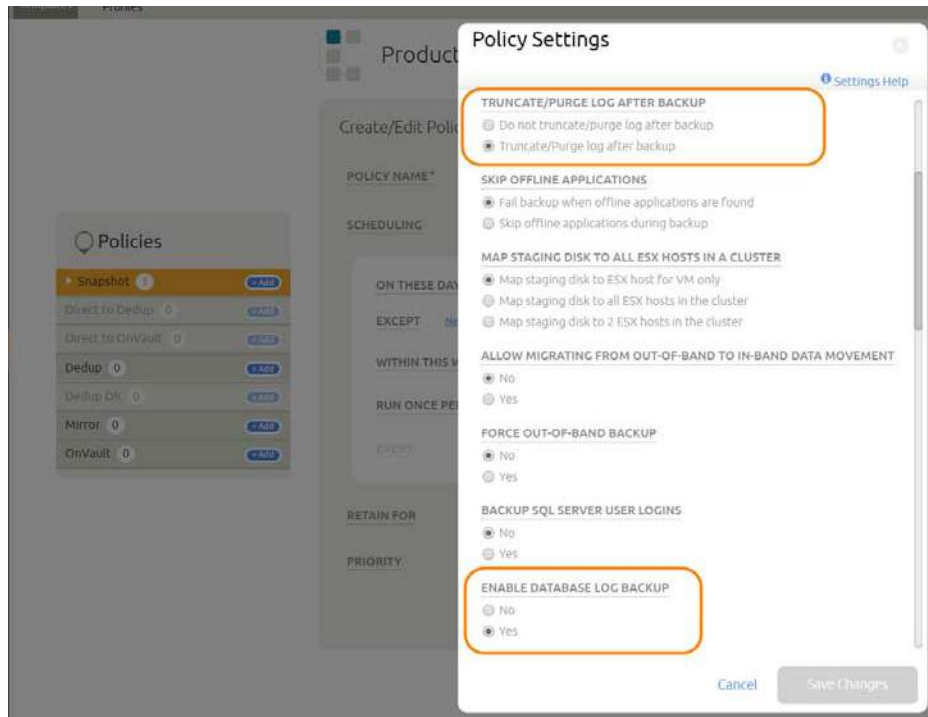
2. Click the Production to Snapshot "+".



3. Select Advanced Policy Settings.

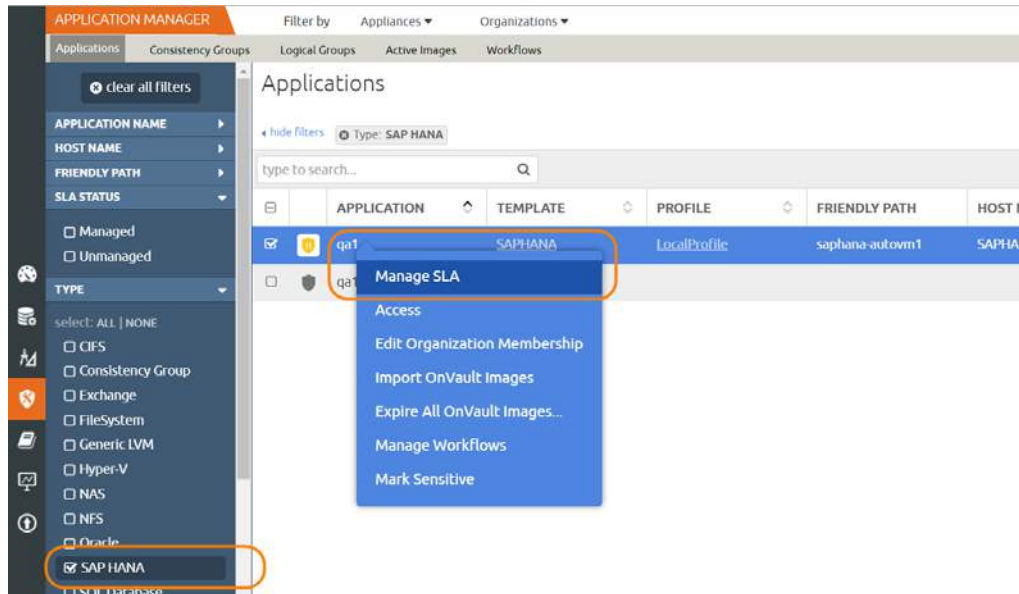


4. Set the log policy options (you will have to scroll to see them all):
  - o Truncate/Purge Log After Backup: Select this.
  - o Enable Database Log Backup: Select this.
  - o RPO (Minutes): Enter the desired frequency of log backup
  - o Log Backup Retention Period (in Days): the SLA to retain the backup of log for point in time recovery.
  - o Replicate Logs (Uses StreamSnap Technology): Select this to enable StreamSnap replication of log backup to a DR site.

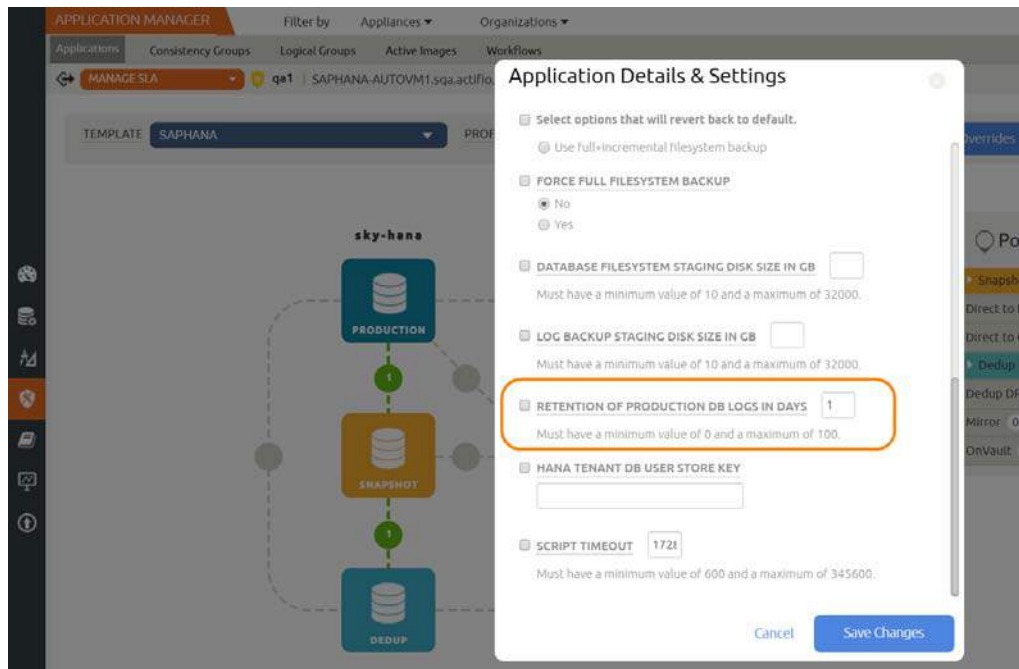




- From Application Manager, select the HANA database. You can use the SAP HANA checkbox to filter the list. Select Manage SLA.



- At the top of the screen, select Details & Settings.



- Set the Retention of Production DB Logs in Days. This value 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 the log will be purged older then the data backup id. Default value is 0 days. With the default value, all logs prior to last data backup are purged.



# 8 Restoring and Recovering an SAP HANA Database

This section includes:

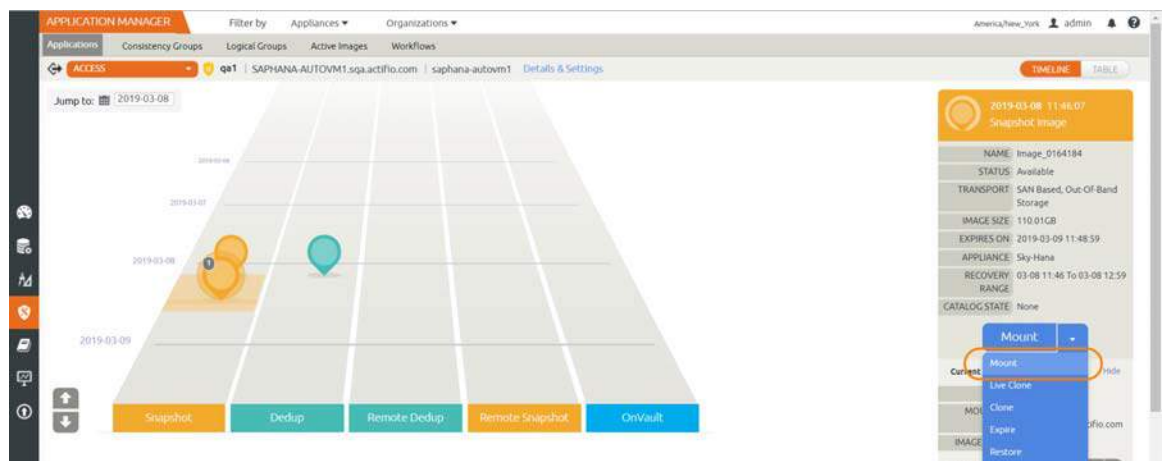
[Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application on page 37](#)

[Restoring and Recovering a SAP HANA Database on page 38](#)

## Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application

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

1. From Application 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 Choose a point in time on the slider bar for a database protected with log roll-forward to recover to.
  - 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).
  - o Mount Location > Specify a Mount Point to mount to new target.
  - o Manage New Application > To reprotect, click and enable Manage New Application.
  - o Template > Choose a template to protect the database.
  - o Profile > Choose a profile.
4. Click Submit.

## Restoring and Recovering a SAP HANA Database

Depending on how you protected the database, you need the procedure for:

[Recovering from Block-Based LVM Snapshot with CBT](#) on page 38

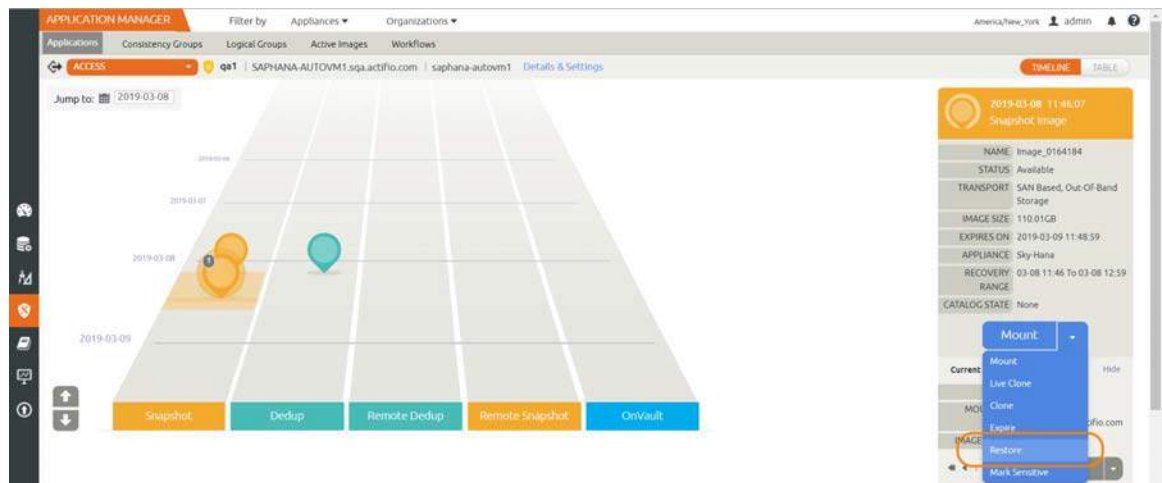
[Recovering from a File-Based Backup with NFS](#) on page 39

### Recovering from Block-Based LVM Snapshot with CBT

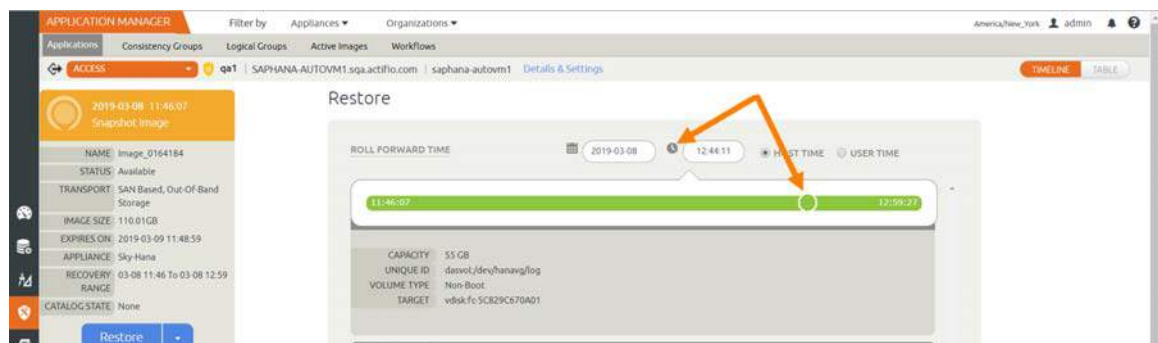
Use this procedure to restore and recover the source HANA database. This procedure uses physical recovery of the source data area.

To recover back to the source:

1. From the Application Manager > Protected Application > Access, from the latest snapshot to recover, choose Restore.



2. On the Restore page choose point in time on the slider bar for database protected with log to recover to desired point in time.



3. Click Submit.

## Recovering from a File-Based Backup with NFS

You have two options:

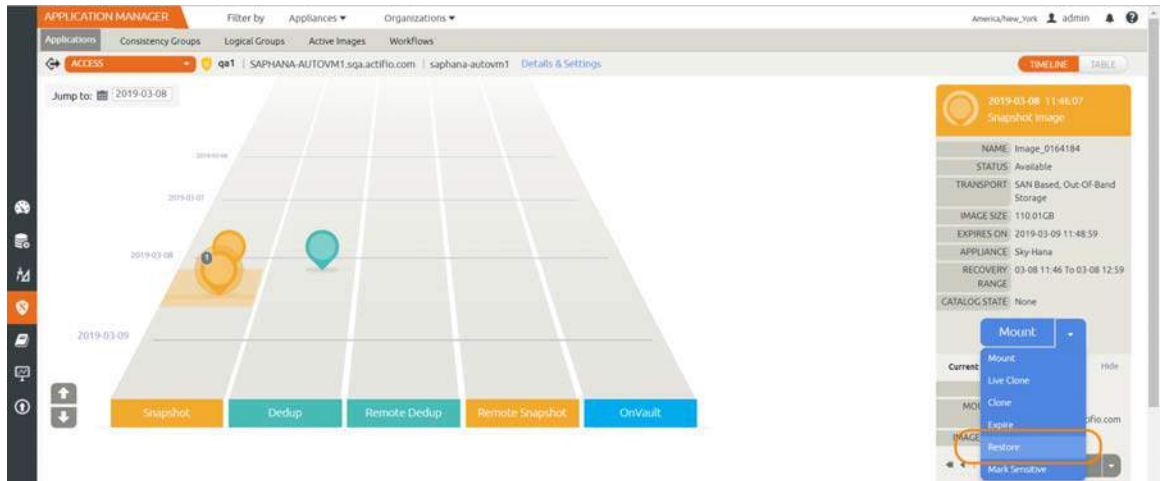
**Recovering Back to the Source:** Use this procedure to restore and recover the source HANA database. This procedure overwrites the source data.

**Recovering to a New Target:** Use this procedure to restore and recover to a new target server.

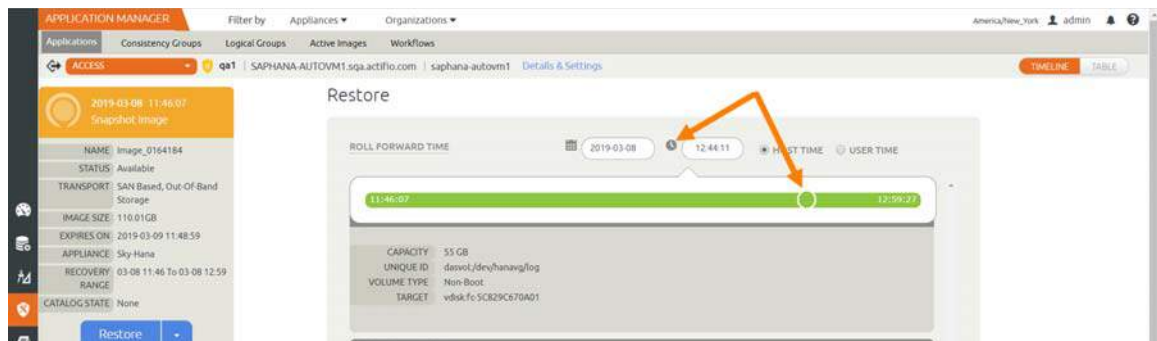
When you are finished, you must bring up the database, as detailed in [Bringing up the HANA Database](#) on page 41.

### Recovering Back to the Source

1. From Application Manager > Protected Application > 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 then a point in time on the slider.



### 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

### Restore

4. Click Submit. This will start the source database physical recovery using HANA recover commands.

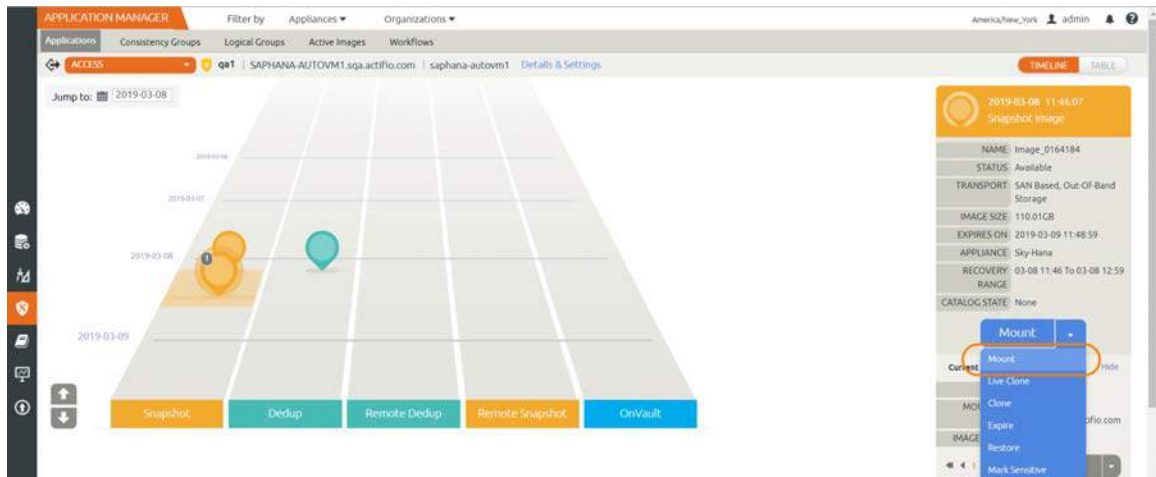
## Recovering to a New Target

Before You Begin:

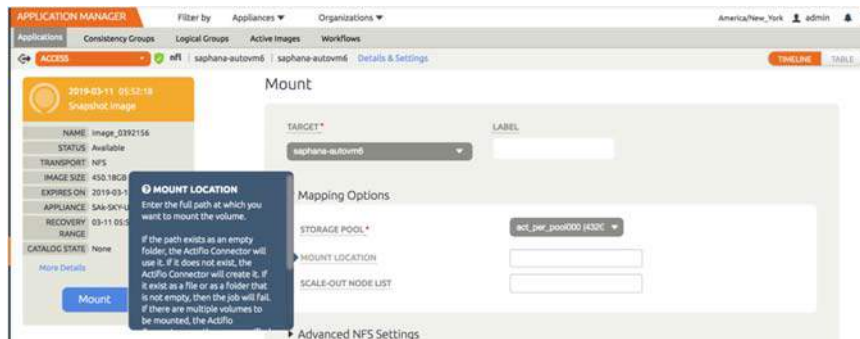
- Make sure target HANA server is set up as same as source HANA server (OS version, CPU and memory, HANA version)
- Make sure HANA database on target server is configured same as source i.e. global.ini, nameserver.ini

To recover:

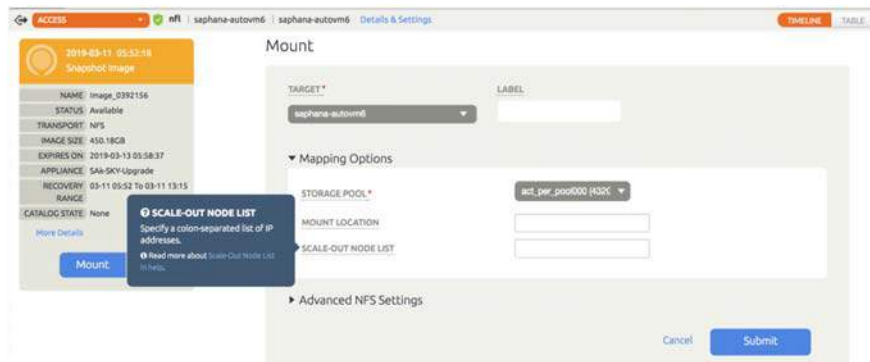
1. From Application Manager > Protected Application > Access, select the latest snapshot to recover, and choose Mount.



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 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.

### Bringing up the HANA Database

To bring up the HANA database from the mounted image, modify and configure this script:

1. Configure `/act/custom_apps/saphana/dump/restoreDumpToNewTarget.conf` parameter

```
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"
```

2. `cd /act/custom_apps/saphana/dump/`
3. Run `ACT_HANADB_newtargetdumprestore.sh`:
 

```
./ACT_HANADB_newtargetdumprestore.sh
```

 or
 

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

