
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	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+	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	17
Creating the SAP HANA Hdbuserstore Key for the System Database and each Tenant Database in a Scale-Out Multi-Node SAP HANA System	18
SAP HANA Database Application Details and Settings.....	19
Using the Tenant DB User Store Key Prefix	20
Chapter 4 - Adding a SAP HANA Database Host and Discovering the Database	21
Adding the Host from the Domain Manager	21
Discovering the HANA Database Application from the Application Manager.....	23
Finding the Discovered HANA Database in the Application Manager	24

Chapter 5 - Configuring the SAP HANA Backup Method	25
Ensure that the Disk Preference on the Host is Set Correctly	26
Ensure that the Backup Capture Method in the Application Settings is Set Correctly	28
Setting the Schedule for Dumps	30
Chapter 6 - Protecting the HANA Database	31
Chapter 7 - Protecting SAP HANA Database Logs	33
Setting up the Log Mode and Log Backup in HANA Studio	33
Setting up the Log Backup in Actifio AGM	35
Chapter 8 - Restoring, Accessing, or Recovering an SAP HANA Database	39
Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application	39
Workflow to Automate Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application	41
Restoring and Recovering an SAP HANA Database	43
Recovering from Block-Based LVM Snapshot with CBT	43
Recovering from a File-Based Backup with NFS	44
Single Tenant DB Recovery using LVM snapshot	47
Bringing up the HANA Database	48
Chapter 9 - HANA Database Management Using actHANADBМ	49
Installing and Configuring actHANADBМ.pl	50
actHANADBМ Commands	51
agmconfig	51
createTemplate	52
hostDiscovery	53
protectApp	54
backup	55
listImageDetails	56
mount	57
unmountdelete	58
restore	59
runwf	60

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 4

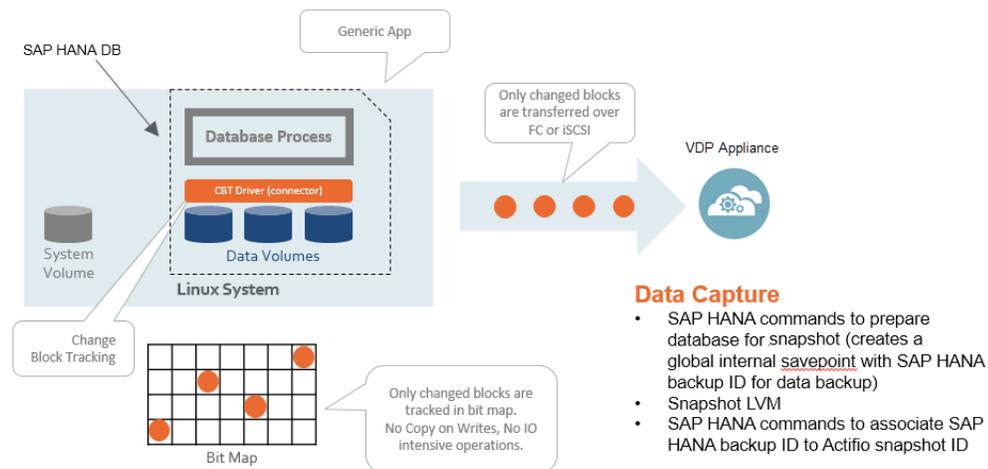
[SAP HANA Backup Methods](#) on page 5

[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



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.

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

The Actifio Appliance 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 8, Restoring, Accessing, 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 8, Restoring, Accessing, or 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.

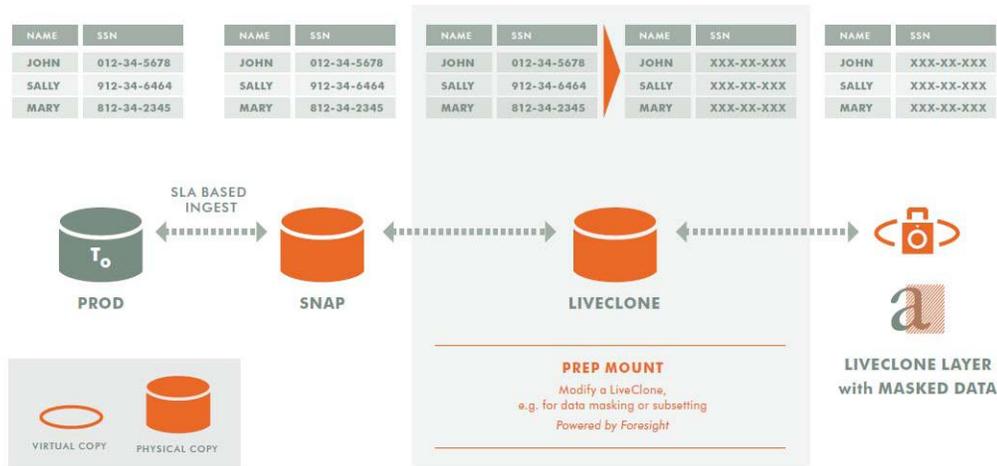
Combining Actifio's automated data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now, instead of having to wait for DBAs to update test and development environments, users can provision their own environments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

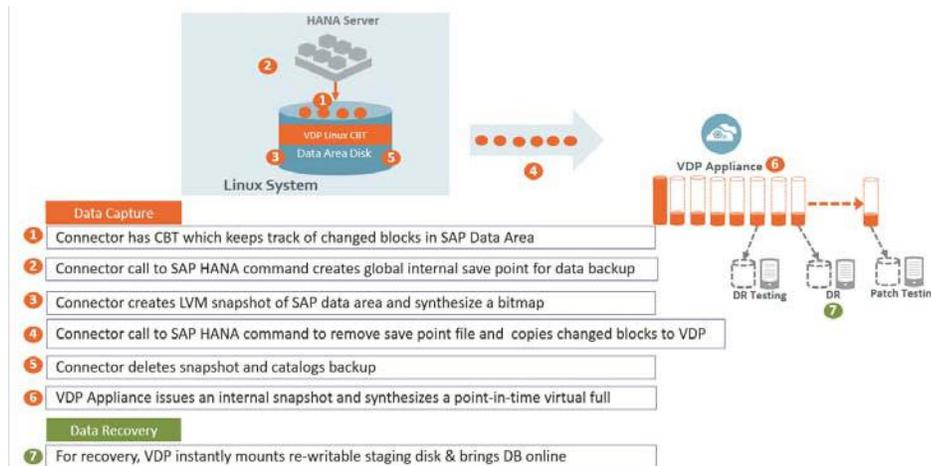
[Block-Based LVM Snapshot with CBT Integrated with SAP HANA Database Storage Snapshot API](#)

[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

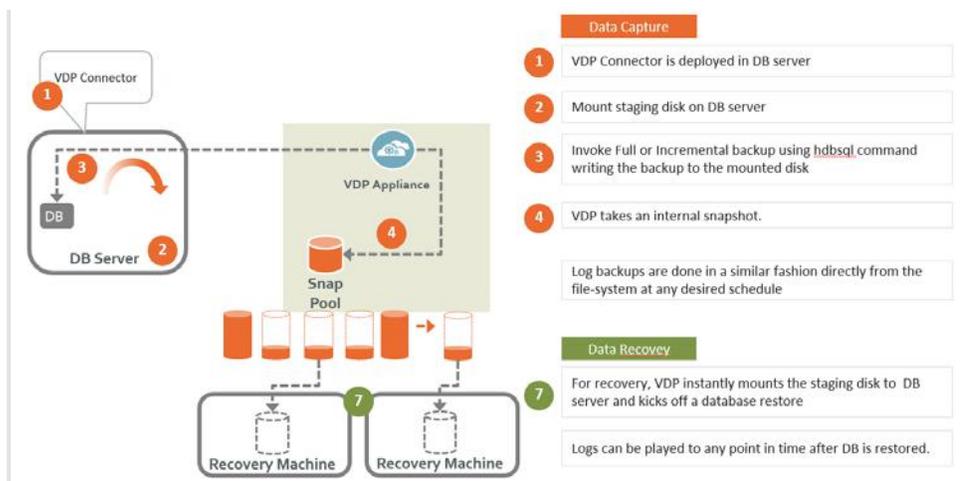
The SAP HANA database creates a database internal snapshot based on a system wide save point executed during the PREPARE step. The database internal snapshot is stored in the data volumes area.



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 7

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

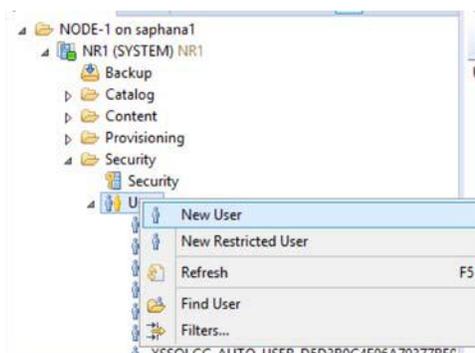
[Adding SAP HANA Hdbuserstore Key in SAP HANA 1.0 \(single container system\)](#) on page 9

Creating the Database User Account

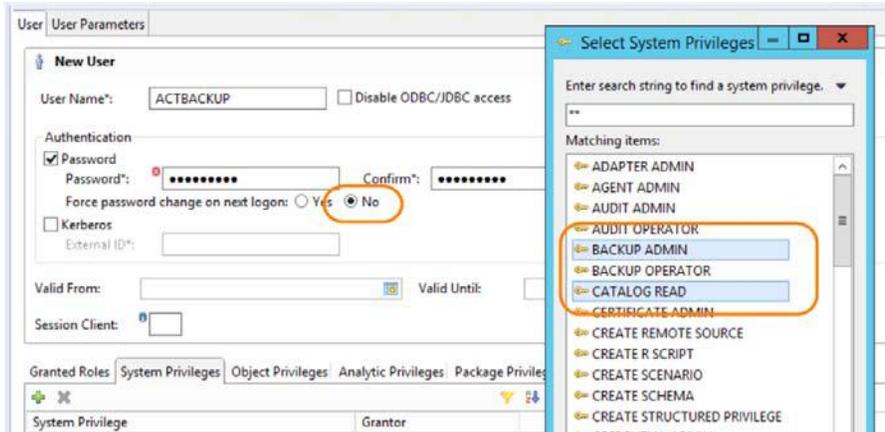
Make sure to create this user account under a single container database. Make sure to provide BACKUP ADMIN and CATALOG READ to back up the user created under database. Choose a database user name based on company's standard.

To create the user:

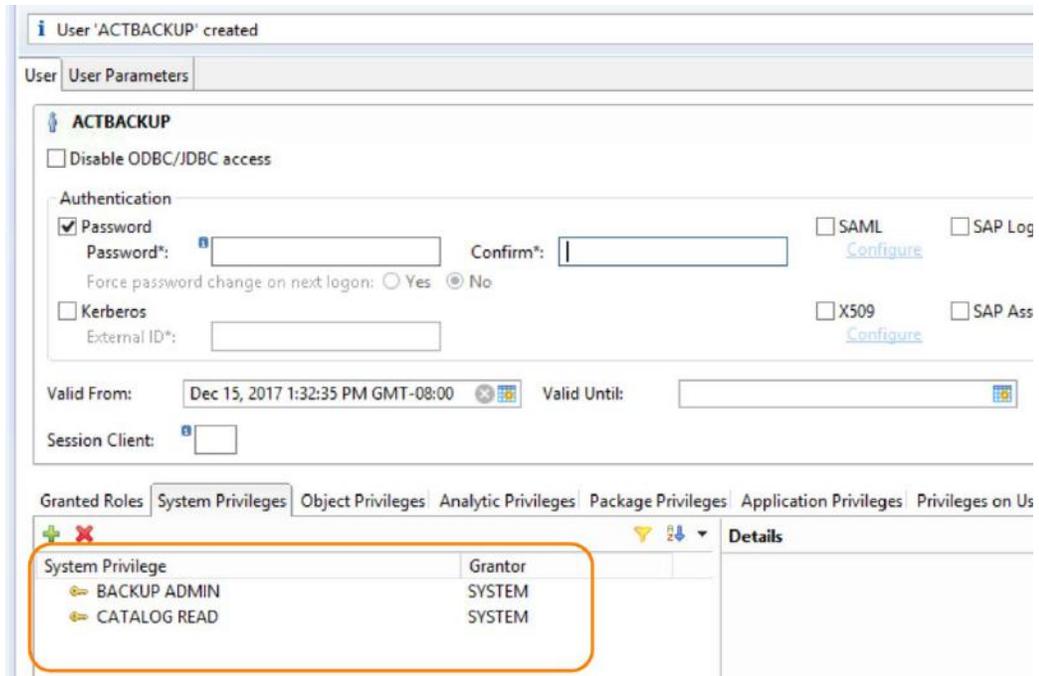
1. From SAP HANA Studio go to System > Security > Users > New User.



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**.

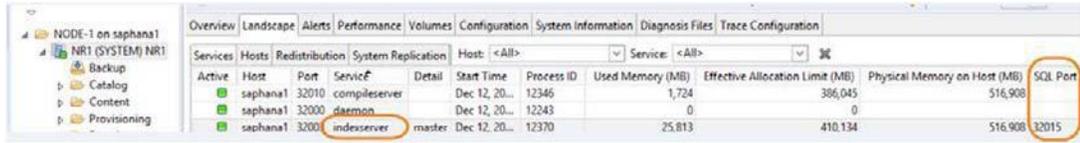


You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN and CATALOG READ privileges.



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.



The screenshot shows the SAP HANA Studio interface. The 'Landscape' 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	Hosts	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...	12346	1,724	386,045	516,908	
Active	saphana1	32000	daemon		Dec 12, 20...	12243	0	0		
Active	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

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

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 11

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

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

[SAP HANA Database Application Details and Settings](#) on page 19

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+](#) on page 13

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

Creating the System Database User Account from HANA STUDIO SPS01

Naming convention for database user account

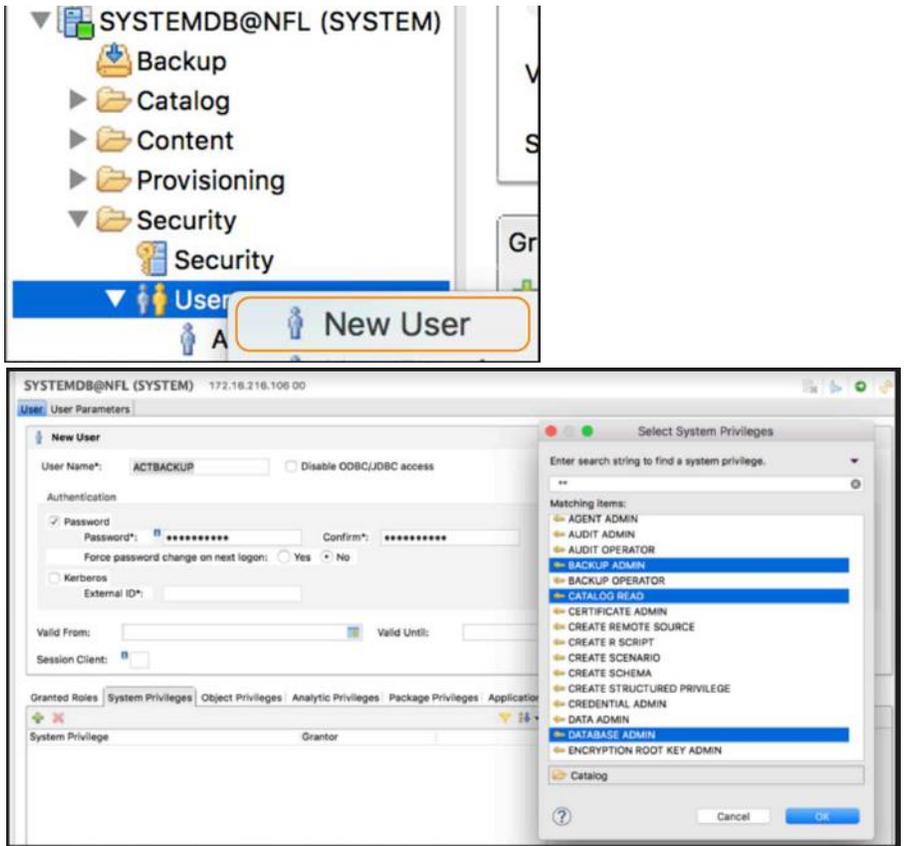
Choose a database user name based on company's standard. Make sure to create this user account under SYSTEMDB. Make sure to provide BACKUP ADMIN, CATALOG READ, and DATABASE ADMIN to the backup user created under database.

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**.

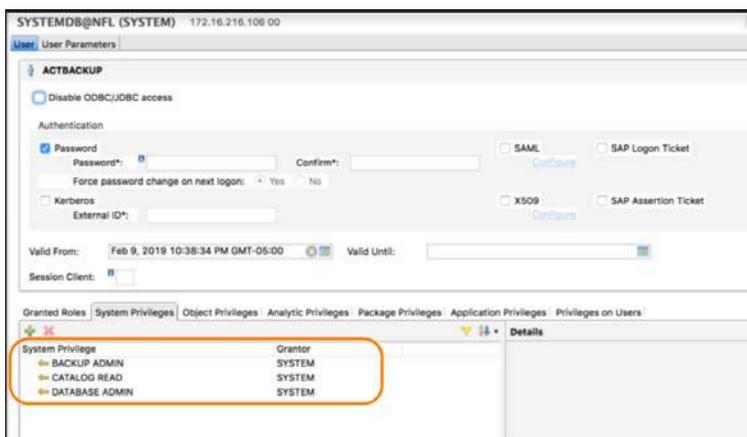
- o From SAP HANA Studio SYSTEMDB, go to System > Security > Users > New User.



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+

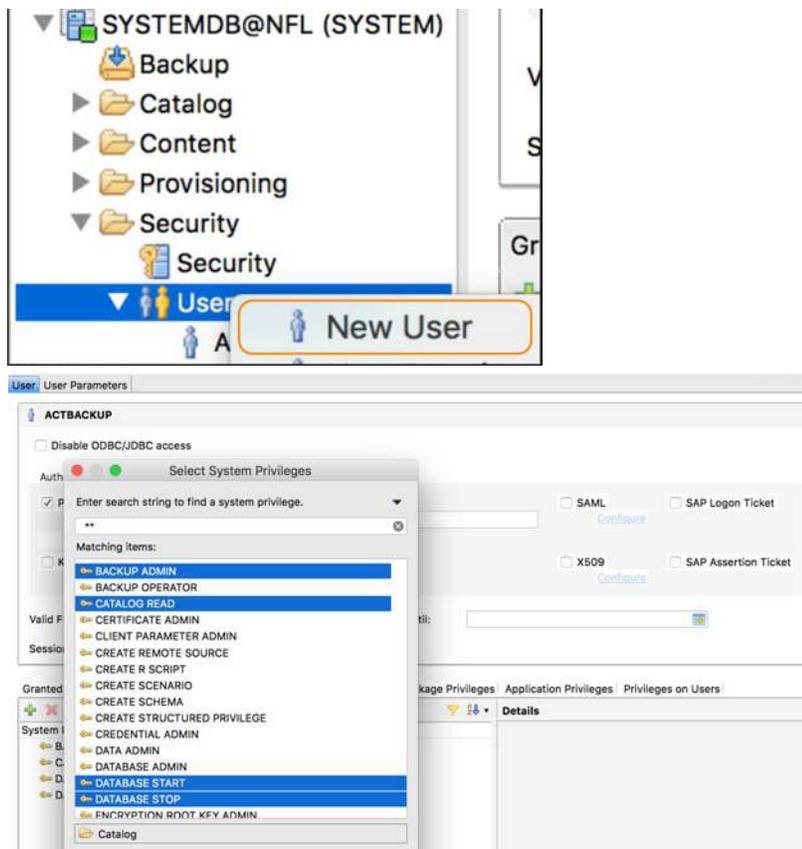
Naming convention for database user account

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.

User Parameters

ACTBACKUP

Disable ODBC/JDBC access

Authentication

Password
 Password*: Confirm*: [Configure](#) SAML SAP Logon Ticket

Force password change on next logon: Yes No

Kerberos
 External ID*: X509 [Configure](#) SAP Assertion Ticket

Valid From: Jul 27, 2019 8:02:00 PM GMT-04:00 Valid Until:

Session Client:

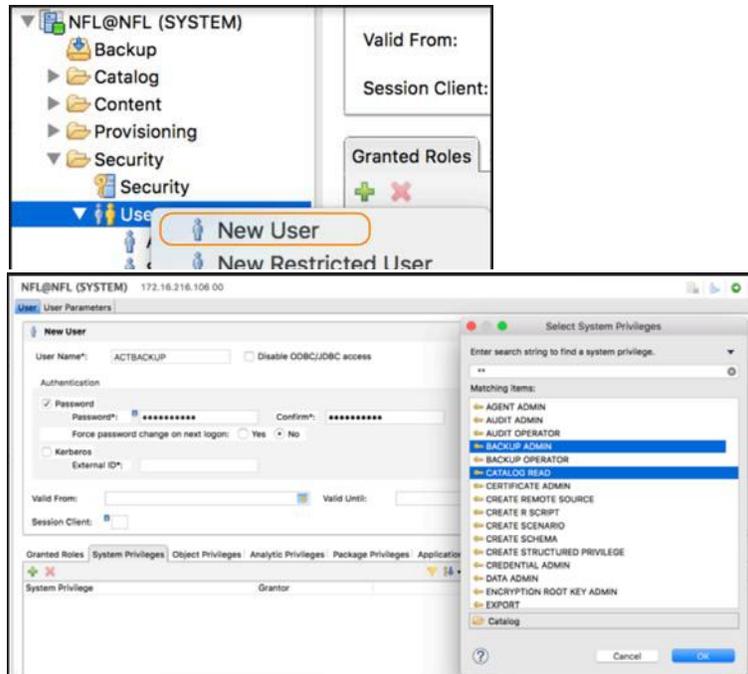
Granted Roles: System Privileges | Object Privileges | Analytic Privileges | Package Privileges | Application Privileges | Privileges on Users

System Privilege	Grantor	Details
BACKUP ADMIN	SYSTEM	
CATALOG READ	SYSTEM	
DATABASE START	SYSTEM	
DATABASE STOP	SYSTEM	

Creating the User under the Tenant DB

To create the tenant database user account:

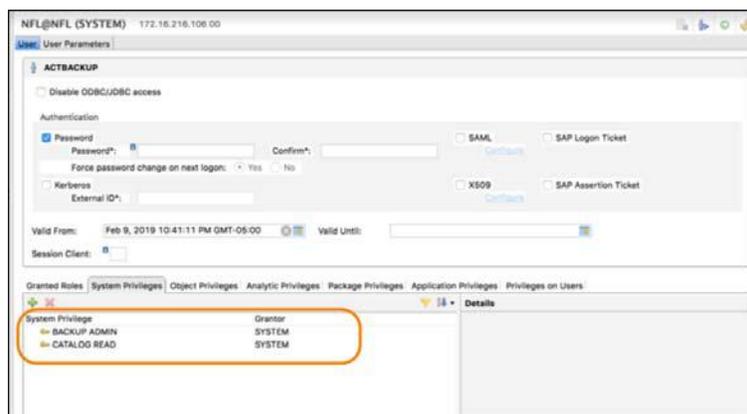
1. Create the USER under TENANTDB
 - o Assign a User Name and a Password.
 - o Set Force password change on next logon to **No**.
 - o Click on the System Privilege tab and assign privileges by selecting **BACKUP ADMIN** and **CATALOG READ**.
 - o From SAP HANA Studio SYSTEMDB, go to TENANTDB > Security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



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

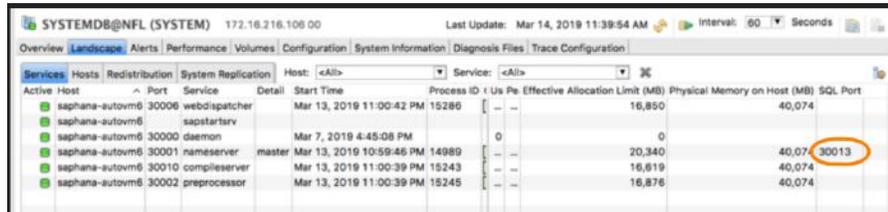


You will get a User Created message and the System Privileges will show the user has been granted BACKUP ADMIN and CATALOG READ privileges.



Getting the Instance and SQL Port Numbers

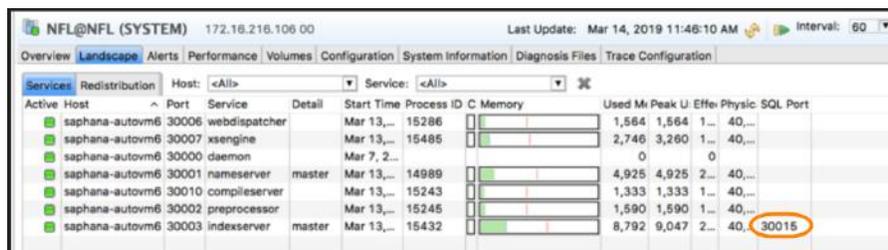
SYSTEMDB: From SYSTEMDB go to System > Landscape and get the value of SQL port for the nameserver. In the example below, 30013 is the SQL port, and the instance number is 00.



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	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	compleserver		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 indexeserver. 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	compleserver		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	indexeserver	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 17](#)

[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 18](#)

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 *****>
```

SAP HANA Database Application Details and Settings

From the Application Details & Settings dialog box (accessed through Details & Settings), you can modify application-specific settings for capturing Microsoft SQL Server databases. Application settings may be useful or required in certain circumstances. After you configure your application settings, click **Save Changes**.

Note: To reset one or more application settings back to its default state, click the check box to the left of the selection you want to reset.

To reset all application selections back to their default state, click *Select options that will revert back to default*.

Application Setting	Description
HANA DB User Store Key	This is the SAP HANA hdbuserstore key for the system database created in earlier. This field is mandatory.
Percentage of Reserve Space in Volume Group	For Block-Based Capture with CBT: This is needed for LVM snapshot temporary space. Recommended value is 20%. For File-Based Backup in NFS: Not applicable.
Backup Capture Method	For Block-Based Capture with CBT: Select Changed block tracking based backup. For File-Based Backup in NFS: Select full+incremental filesystem backup. File-based backup also requires additional CLI configuration. It is important to ensure that the Backup Capture Method is configured correctly. Do not skip this step.
Force Full FileSystem Backup	For Block-Based Capture with CBT: Not applicable. For File-Based Backup in NFS: Use for an ad hoc full backup.
Database Filesystem Staging Disk Size (GB)	For Block-Based Capture with CBT: Not applicable. For File-Based Backup in NFS: Default calculation is based on (database size * 1.5) + 10% and the disks will grow dynamically.
Log Backup Staging Disk Size (GB)	By default, Actifio calculates the daily log generation * retention of log backup SLA(+20% overhead). Keeping the default is recommended. Providing a fixed 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.
Tenant DB User Store Key Prefix	See Using the Tenant DB User Store Key Prefix on page 20.
Script Timeout	The timeout value is applied to internal backup and recovery scripts called by connector. The default value is recommended.

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 [SAP HANA Database Application Details and Settings](#) on page 19).

Use case 1

You have created a user store key and you have a tenant database TD1:

SYSTEMDB user store key = ACTBACKUP

TENANT DB user store key = ACTBACKUPTD1

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 TDBACKUPTD1 (tenant db name is different from system db name):

SYSTEMDB user store key = ACTBACKUP

TENANT DB user store key = TDBACKUPTD1

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 the Database

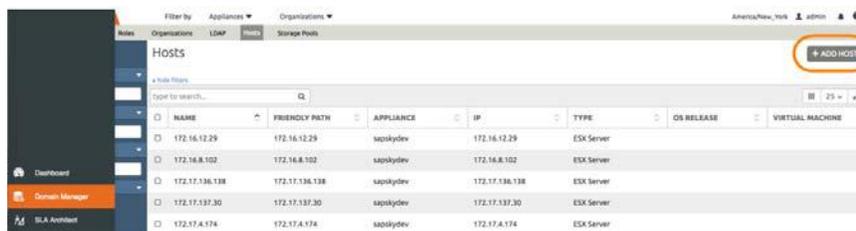
Before you can protect an SAP HANA database, you must add the host and discover the database. This requires:

1. [Adding the Host from the Domain Manager](#) on page 21
2. [Discovering the HANA Database Application from the Application Manager](#) on page 23
3. [Finding the Discovered HANA Database in the Application Manager](#) on page 24

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.

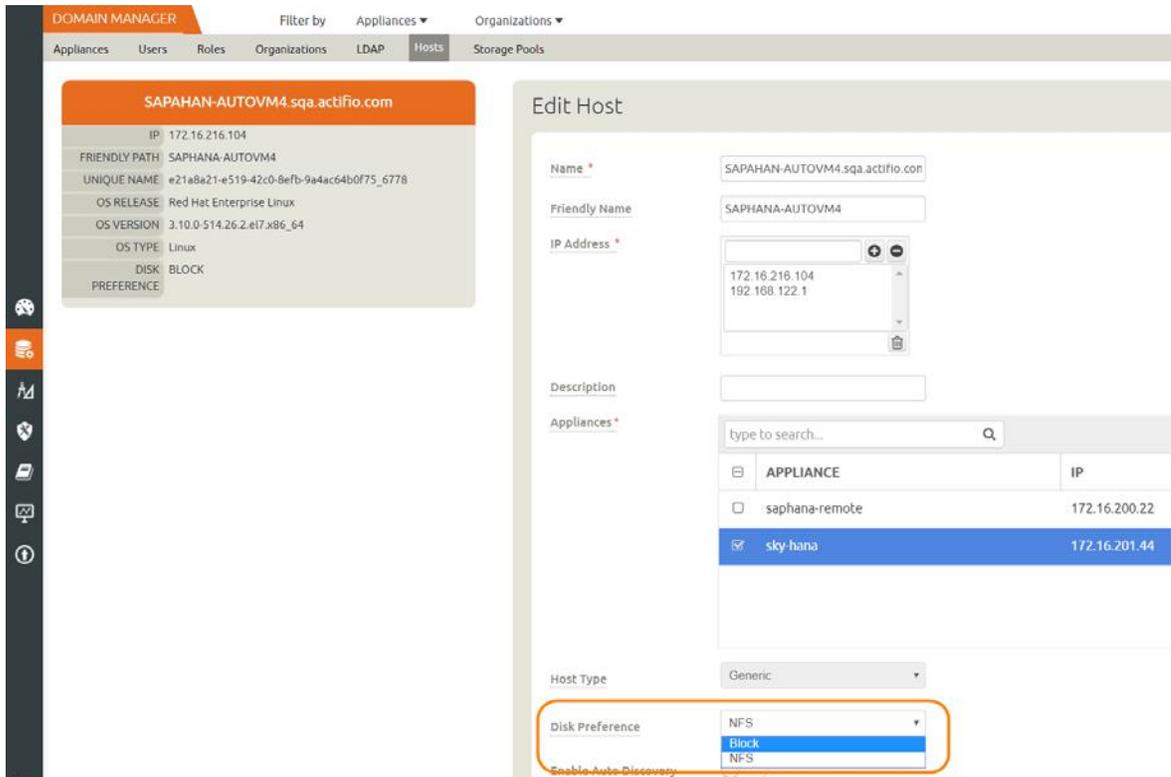


2. On the Add Host page:
 - o Name: Provide the HANA database server name.
 - o IP Address: Provide the HANA database server IP and click the + sign on the right corner.
 - o Appliances: Select the check box for the appliance.
 - o Host Type: Make sure this is Generic.
 - o Click Add at bottom right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

The Host will be 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**.



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 in the 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 with the following components:

- Header:** APPLICATION MANAGER, Filter by Appliances, Organizations, America
- Navigation:** Applications, Consistency Groups, Logical Groups, Active Images, Workflows
- Left Panel:**
 - clear all filters
 - APPLICATION NAME, HOST NAME, FRIENDLY PATH, SLA STATUS
 - Managed, Unmanaged
 - TYPE: select: ALL | NONE
 - CIFS
 - Consistency Group
 - Exchange
 - FileSystem
 - Generic LVM
 - Hyper-V
 - NAS
 - NFS
 - Oracle
 - SAP HANA** (checked)
 - SQL Database
 - SQL Instance
 - System State
 - VM

Applications Table:

	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/Block
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. 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 [Setting the Schedule for Dumps](#) on page 30.

Whichever method you select, you must:

[Ensure that the Disk Preference on the Host is Set Correctly](#) on page 26

[Ensure that the Backup Capture Method in the Application Settings is Set Correctly](#) on page 28

Ensure that the Disk Preference on the Host is Set Correctly

Choose between:

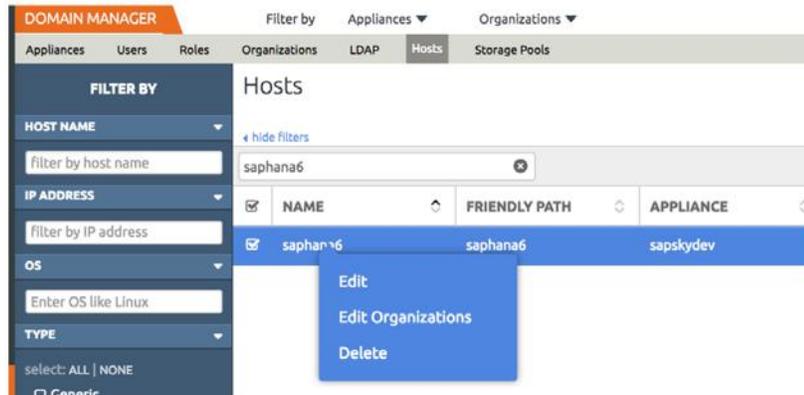
[Setting Disk Preference for Block-Based Database Storage Snapshots with CBT](#) on page 26

[Setting Disk Preference for File-Based Traditional Backup and Recovery in NFS](#) on page 27

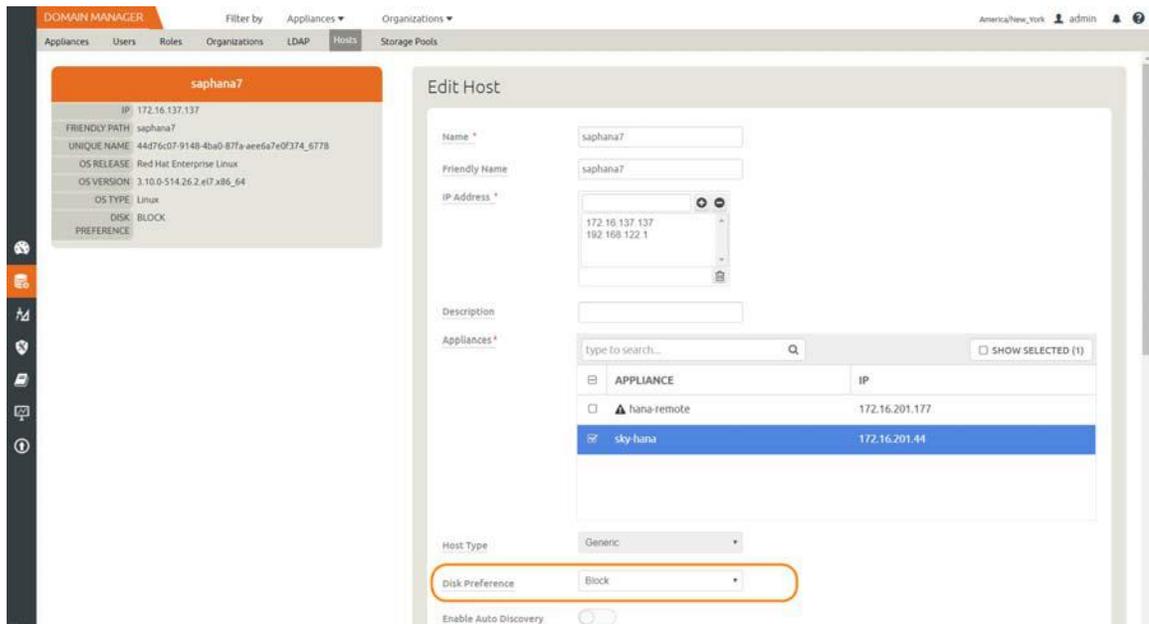
Setting Disk Preference for Block-Based Database Storage Snapshots with CBT

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.



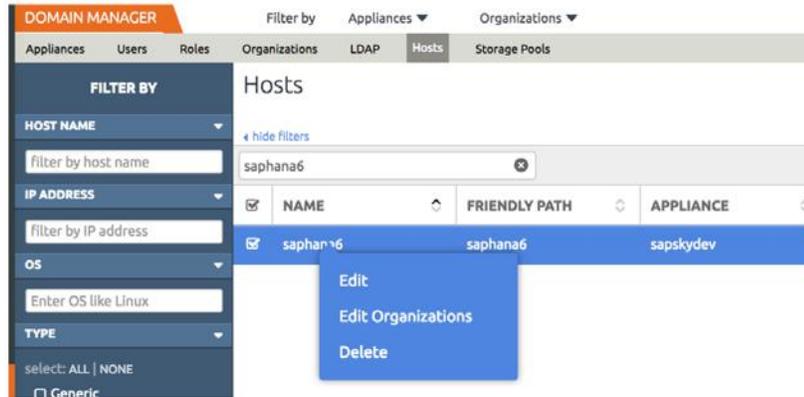
2. In the Edit Host pane, set Disk Preference to *Block* and click Save at the bottom of the page.



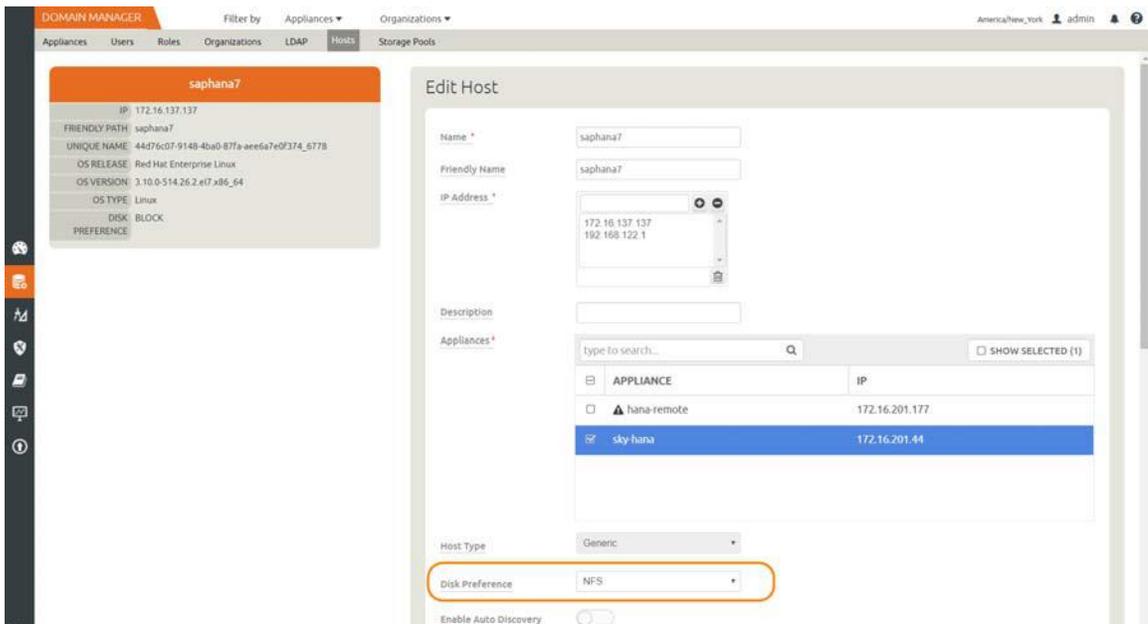
Setting Disk Preference for File-Based Traditional Backup and Recovery in NFS

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 28

[Setting Backup Capture Method for File-Based Backup and Recovery in NFS](#) on page 29

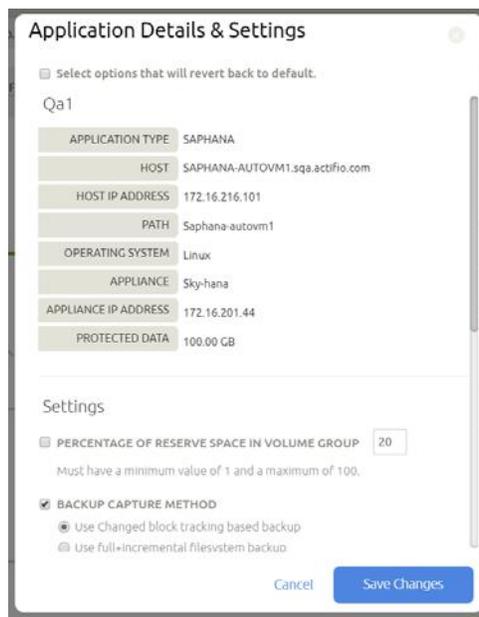
Setting Backup Capture Method for Block-Based Database Storage Snapshots with CBT

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 25.



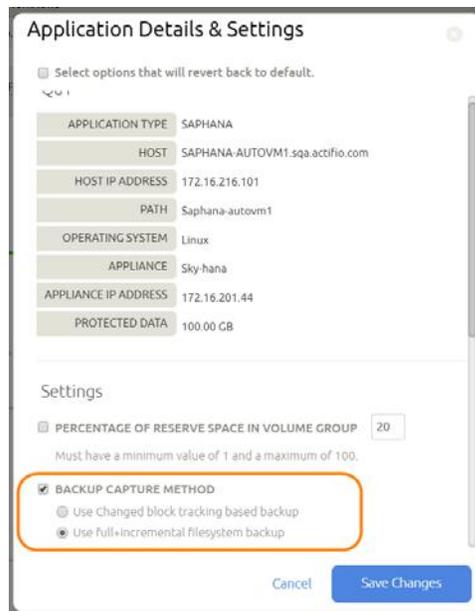
Setting Backup Capture Method for File-Based Backup and Recovery in NFS

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 25.



Setting the Schedule for Dumps

The database dump schedule is set by the Actifio CLI policy parameter `dumpschedule`. The default value of `dumpschedule="FIIIIII"`:

- The string must be seven characters - either an 'F' or an 'I'
- Each position within the string represents a weekday, starting with Sunday.
- **F** represents a full db dump
- **I** represents an incremental db dump

For example, "FIIIIII" results in:

- Sunday: Full backup
- Monday through Saturday: Incremental backups
- The following Sunday: Full backup again

To check the dump schedule, run this CLI command from the appliance:

```
udsinfo lspolicyoption -filtervalue appid=<appid> | grep dumpschedule
```

If this does not return any value, then the `dumpschedule` is set to default.

To modify the dump schedule run this CLI command from Appliance:

```
udstask mkpolicyoption -appid <appid> -name "dumpschedule" -value "FIIIIII"
```

Replace `<appid>` with the application id of the SAP HANA application.

Replace "FIIIIII" as needed.

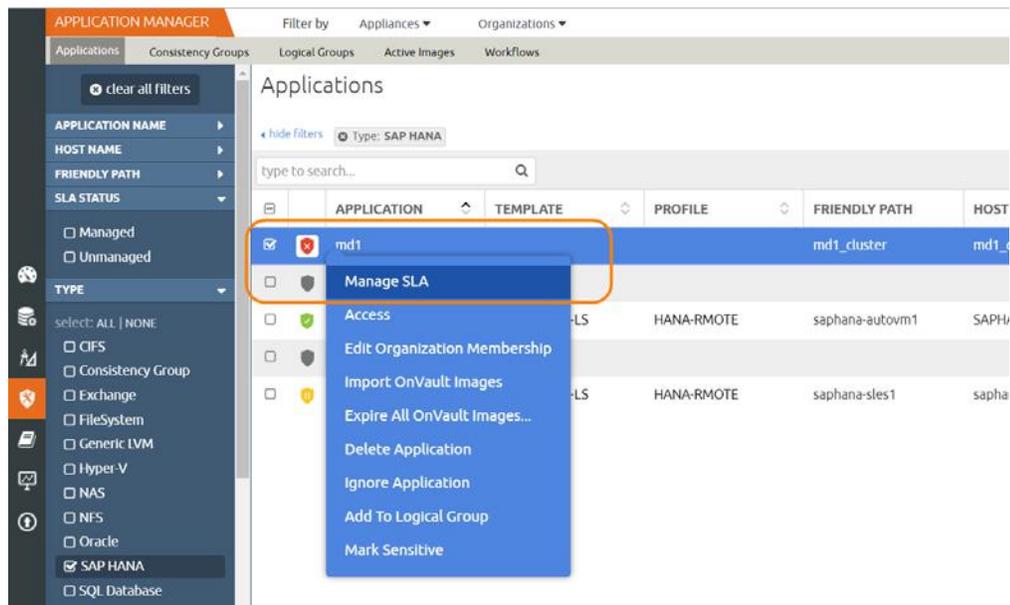
Example

To run full backup on Saturday and Tuesday, set `dumpschedule="IIFIIF"`

6 Protecting the HANA Database

To protect the database:

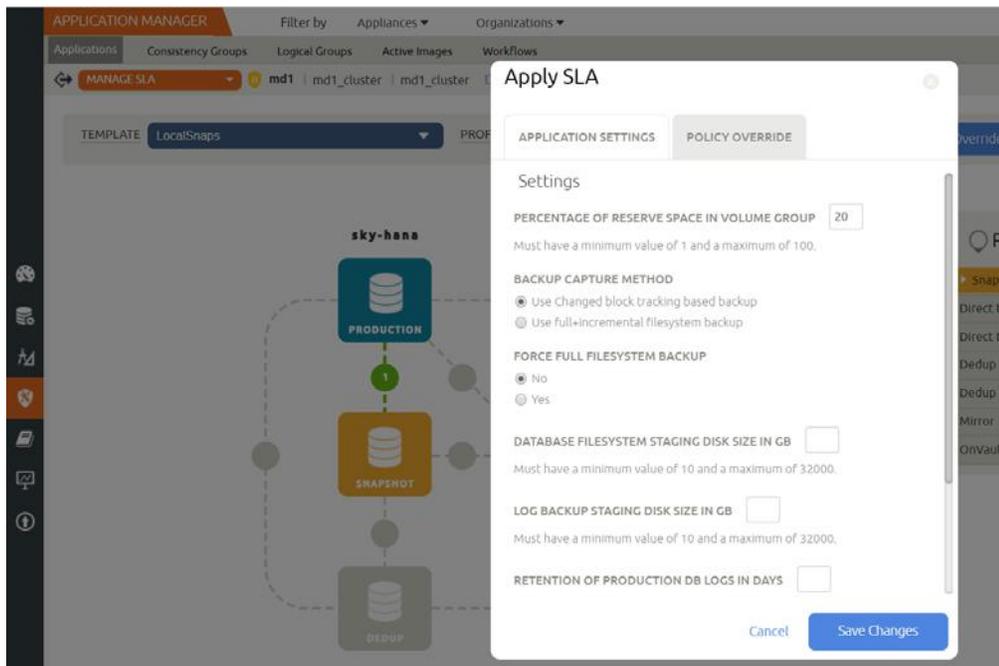
1. Right-click the HANA database and select **Manage SLA**.



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 25. 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

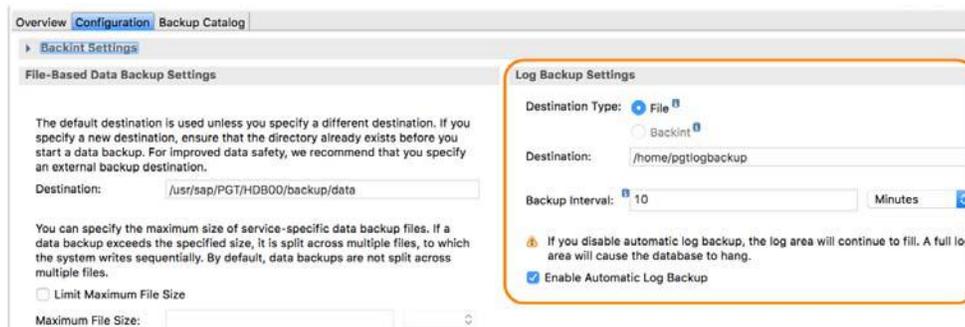
There are two parts to configuring protection of SAP HANA database logs:

[Setting up the Log Mode and Log Backup in HANA Studio on page 33](#)

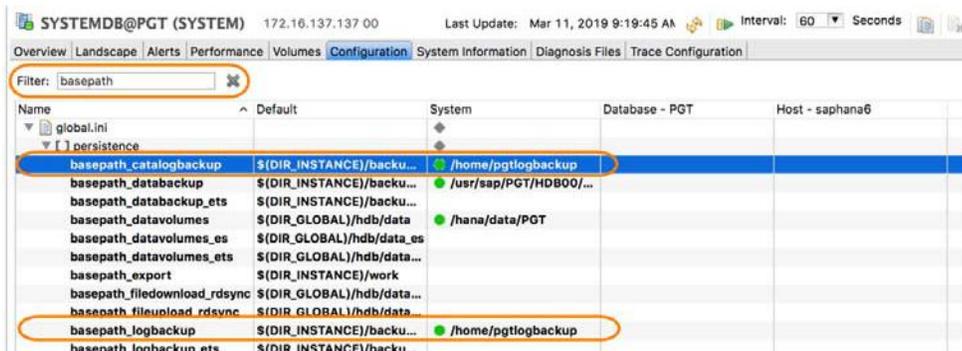
[Setting up the Log Backup in Actifio AGM on page 35](#)

Setting up the Log Mode and Log Backup in HANA Studio

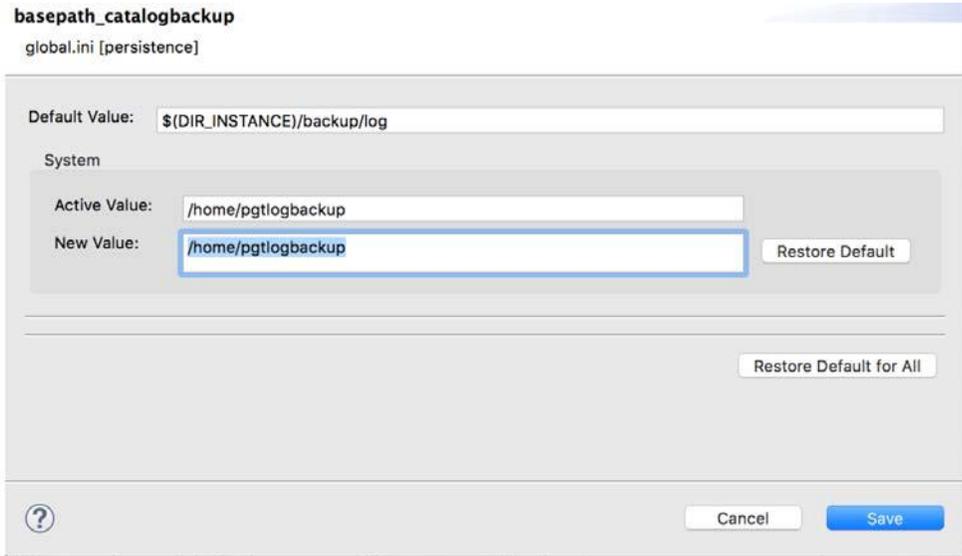
1. In SAP HANA HDB studio, make sure log backup is set correctly under DATABASE (SYSTEMDB FOR HANA 2.0) - Backup - Configuration page
 - o Destination Type is File.
 - o Destination is set to a local file system mount path.
 - o Backup Interval is set to required RPO.
 - o Automatic Log Backup is enabled.



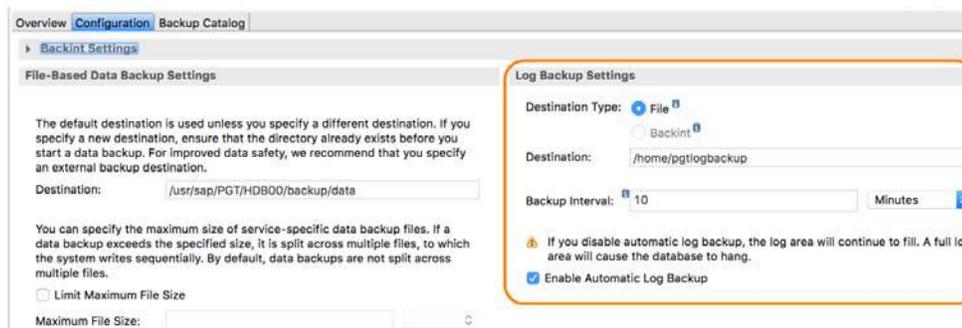
2. Check under Database configuration: DATABASE (SYSTEMDB FOR HANA 2.0) - Configuration page. In the filter, type **basepath**.



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.



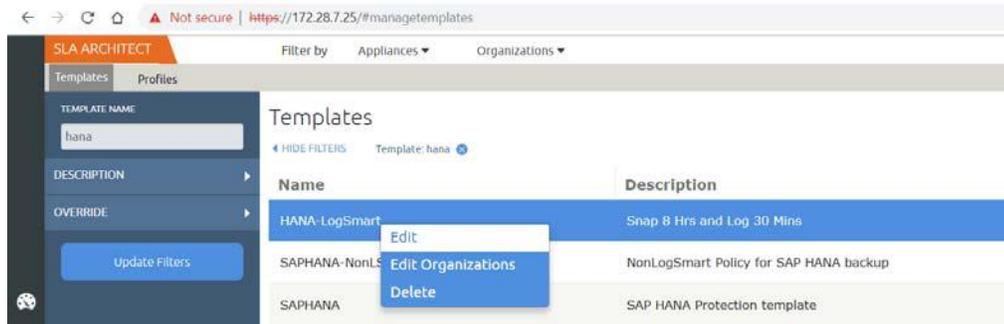
4. Make sure tenant db log backup is set correctly under DATABASE (TENANTDB FOR HANA 2.0) - Backup - Configuration page
 - o Destination Type is File.
 - o Destination is set to a local file system mount path.
 - o Backup Interval is set to required RPO.
 - o Automatic Log Backup is enabled.



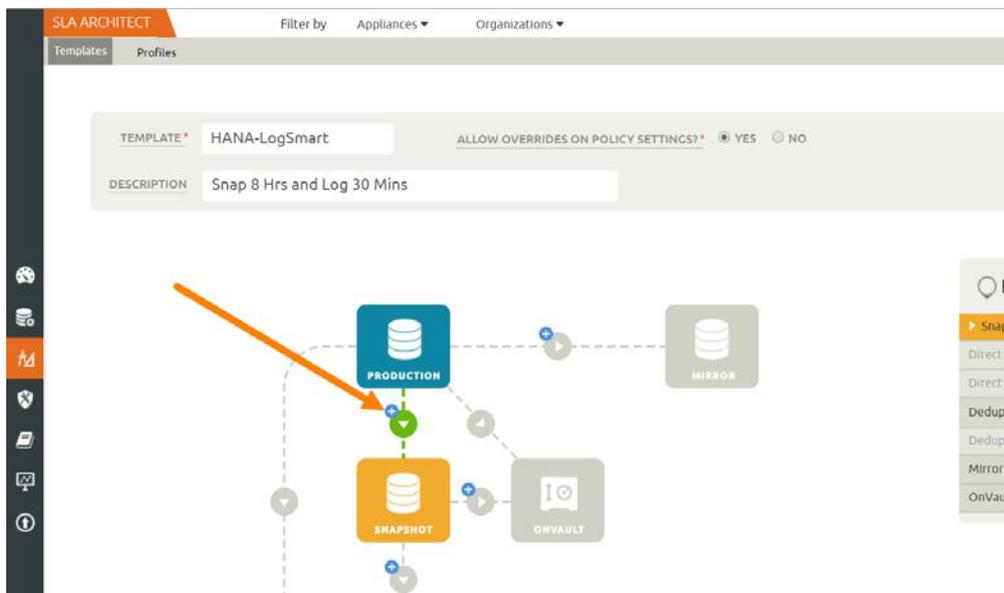
Setting up the Log Backup in Actifio AGM

To enable and set up the HANA database log backup:

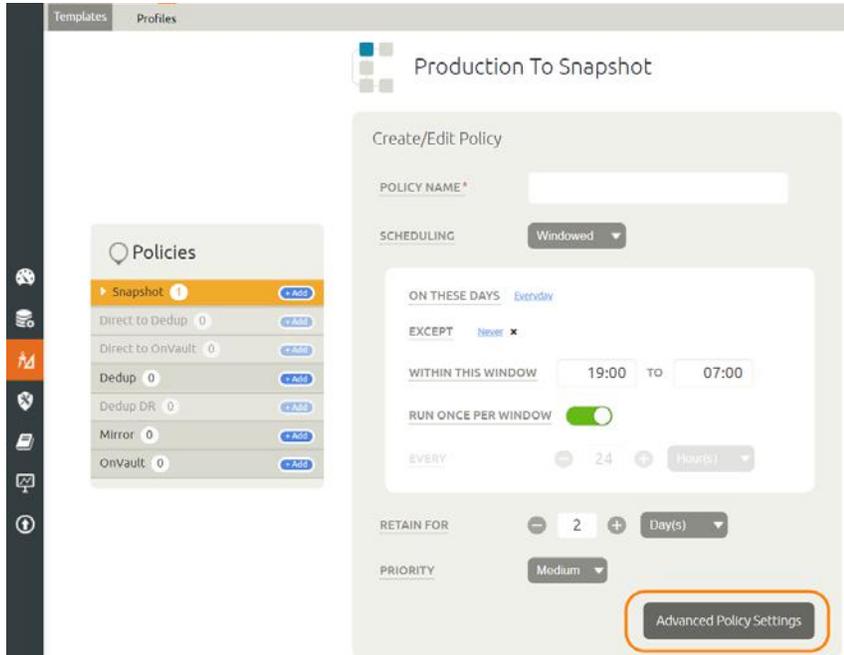
1. From the SLA Architect page, edit the template created for HANA database protection:



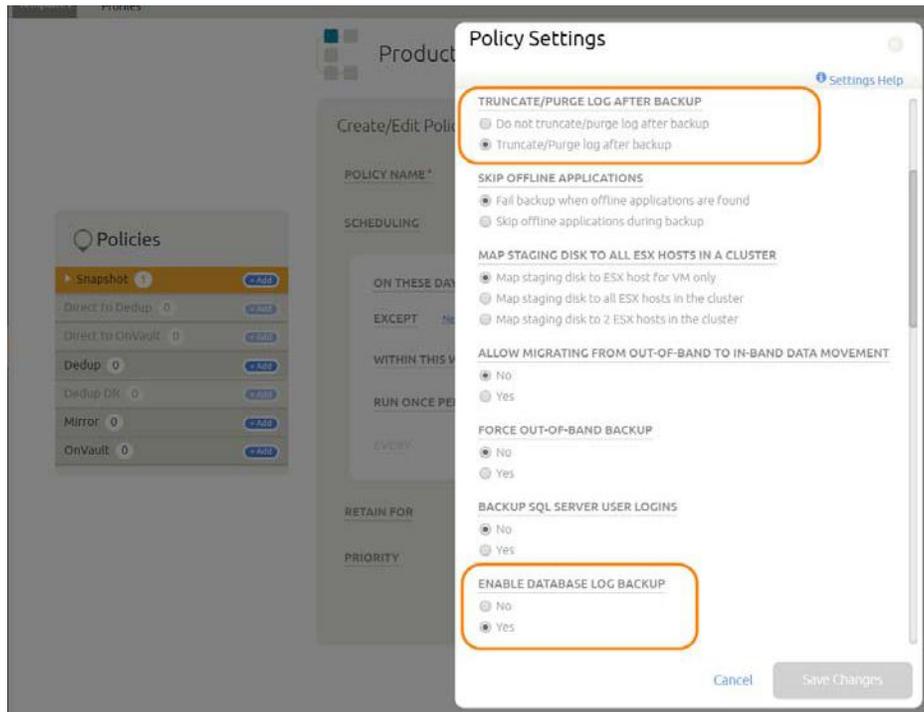
2. Click the Production to Snapshot "+".



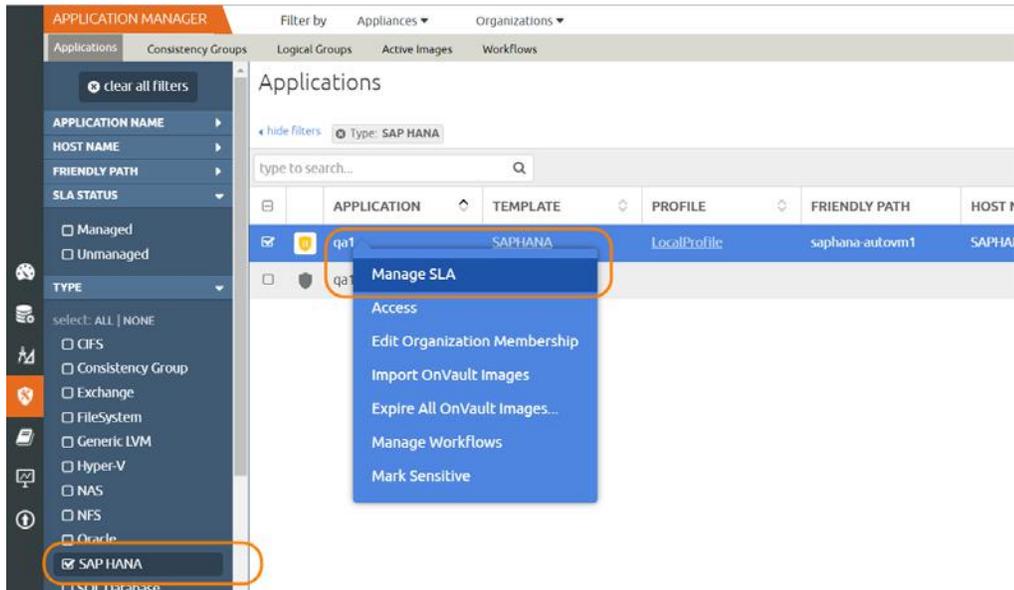
3. Select **Advanced Policy Settings**.



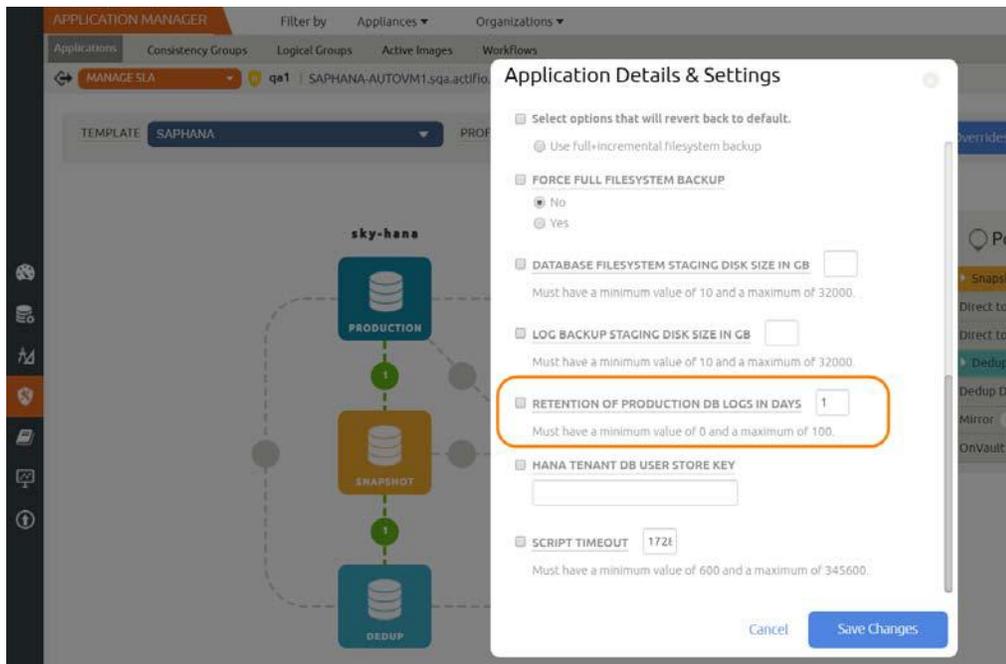
4. Set the log policy options (you will have to scroll to see them all):
 - o Truncate/Purge Log After Backup: Select this.
 - o Enable Database Log Backup: Select this.
 - o RPO (Minutes): Enter the desired frequency of log backup
 - o Log Backup Retention Period (in Days): the SLA to retain the backup 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 logs older than the data backup id will be purged. The default value is 0 days. With the default value, all logs prior to last data backup are purged.

8 Restoring, Accessing, or 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 39](#)

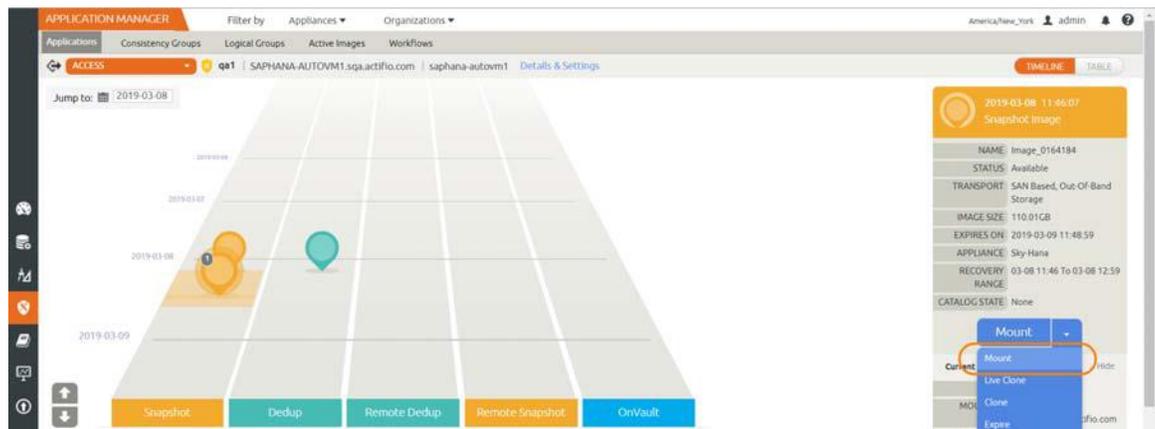
[Workflow to Automate Mount and Refresh from Block-Based LVM Snapshot with CBT to a Target SAP HANA Database as a Virtual Application on page 41](#)

[Restoring and Recovering an SAP HANA Database on page 43](#)

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 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 new target.
 - o For Manage New Application, if you want to reprotect the database, click and 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.

ACCESS | ppt | saphana6 | saphana6 | Details & Settings

2019-03-20 06:30:55
Snapshot Image

NAME	Image_0213479
STATUS	Available
TRANSPORT	SAN Based, Out-Of-Band Storage
IMAGE SIZE	110.01GB
EXPIRES ON	2019-03-22 06:33:46
APPLIANCE	Sky-Hana
RECOVERY RANGE	03-20 06:30 To 03-20 08:31
CATALOG STATE	None

Mount +

Mount

TARGET * LABEL

saphana6

▼ Application Options

CREATE NEW VIRTUAL APPLICATION

ROLL FORWARD TIME HOST TIME USER TIME

TARGET DATABASE SID *

SAP_DB_USER_STORE-KEY *

MANAGE NEW APPLICATION

▼ Mapping Options

STORAGE POOL *

MOUNT LOCATION *

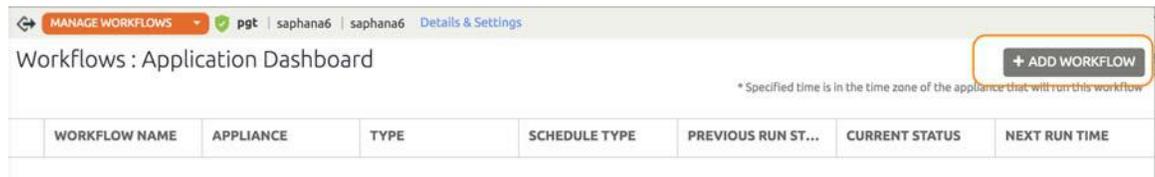
Cancel Submit

4. Click **Submit**.

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

You can use a workflow to automate the process of mounting and refreshing a HANA database from a snapshot:

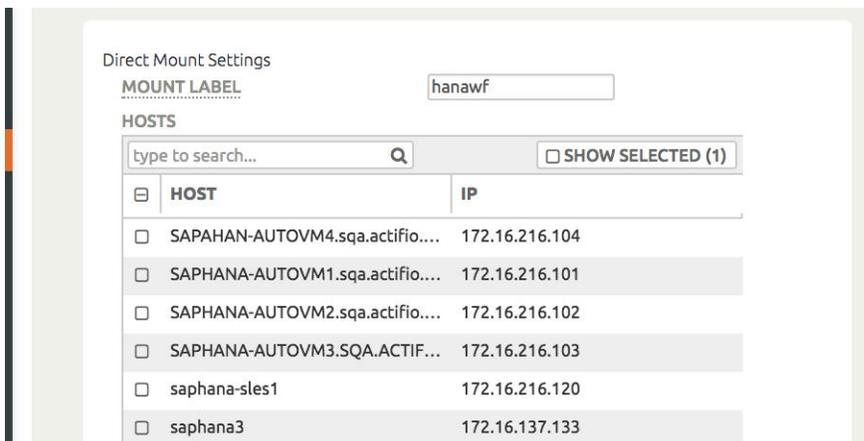
1. From the AGM Application Manager, right-click the HANA database and select **Manage Workflows**.
2. In the upper right corner of the Workflows: Application Dashboard page, click **+ Add Workflow**.



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 based on your requirement. 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 or hosts where the virtual HANA database copy will be created.



- o Mount Location: Specify a mount point to mount 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 **Connecting Hosts to Actifio Appliances**.
- 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).

The screenshot shows a configuration form with the following fields and values:

- Mapping Options**
 - MOUNT LOCATION**: /halmnt
- Script Options**
 - PRE-SCRIPT**: [Empty field] **TIME OUT (SECONDS)**: [Empty checkbox]
 - POST-SCRIPT**: [Empty field] **TIME OUT (SECONDS)**: [Empty checkbox]
- CREATE NEW VIRTUAL APPLICATION**: [Checked toggle]
- TARGET DATABASE SID ***: hal
- SAP DB USER STORE-KEY ***: ACTBACKUP

Optional, if you want to re-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.

Restoring and Recovering an SAP HANA Database

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

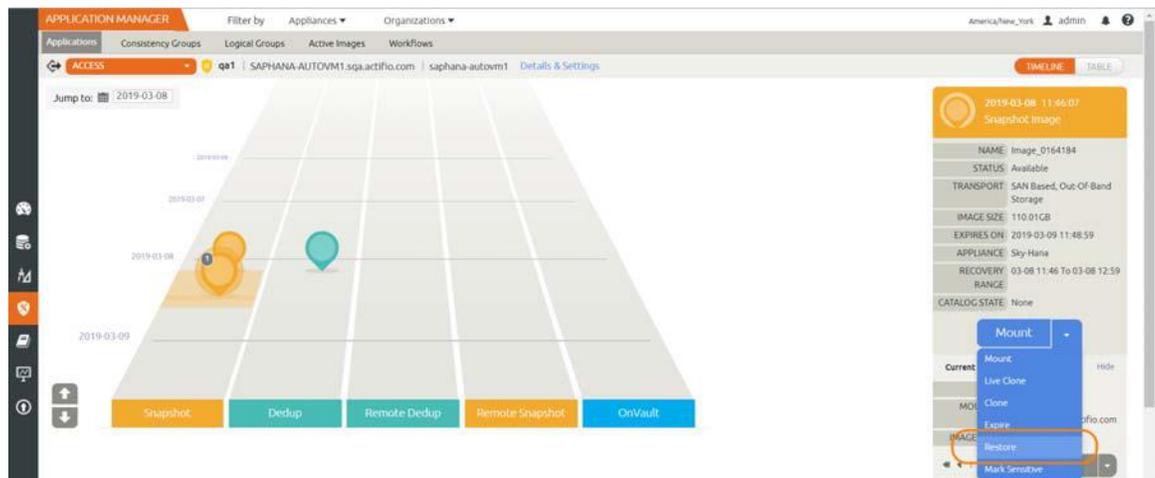
- Recovering from Block-Based LVM Snapshot with CBT on page 43
- Instant Recovery from Block-Based LVM Snapshot on page 44
- Recovering a Single Tenant Database using LVM Snapshot on page 46
- Recovering from a File-Based Backup Back to the Source on page 48
- Recovering from a File-Based Backup to a New Target on page 49

Recovering from Block-Based LVM Snapshot with CBT

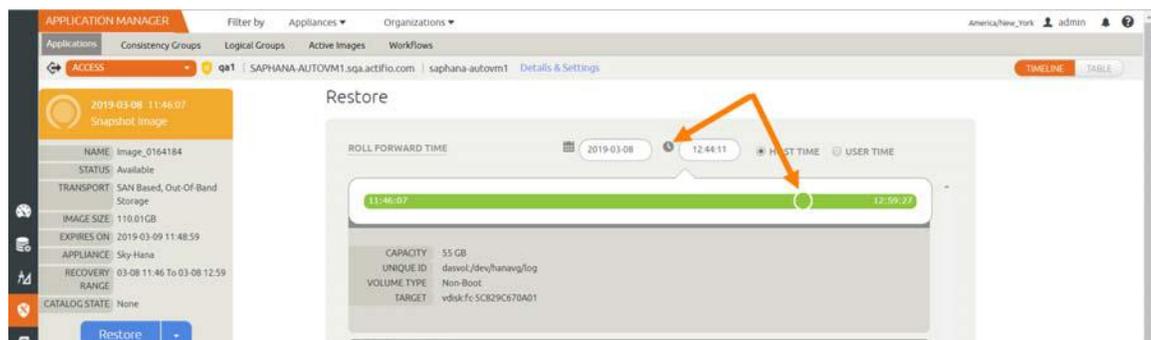
Use this procedure to restore and recover the source HANA database. This procedure uses physical recovery of the source data area.

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 a point in time on the slider bar for a database protected with logs to recover to the desired point in time.



3. Click **Submit**.

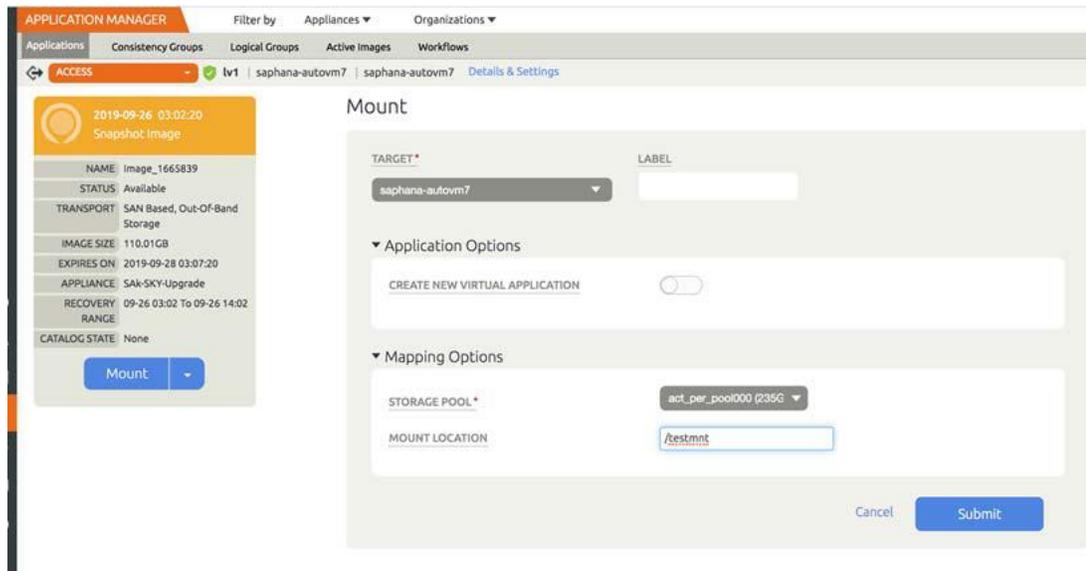
Instant Recovery from Block-Based LVM Snapshot

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

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

To recover from block-based LVM snapshot:

1. Login to AGM, select the application and perform simple mount as shown in the screenshot below.



2. Once the Simple 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.

Where:

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>

For 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.

- 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>
```

Note: Comment *PROD DATA VOLUME* and *PROD LOG VOLUME* from */etc/fstab* until migration is complete.

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
```

- Once this script completes, go to AGM and perform an unmount + delete as shown below.

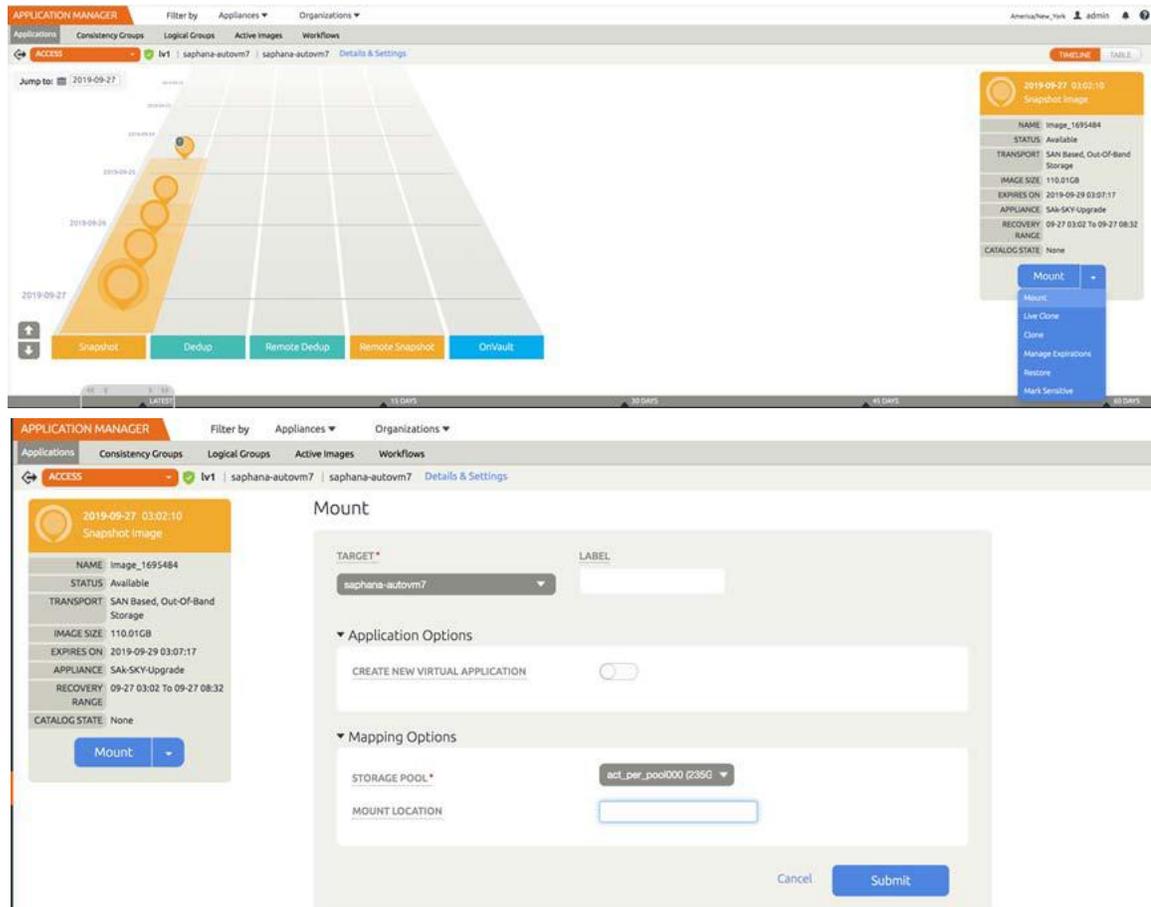
The screenshot shows the Application Manager interface. The top navigation bar includes 'APPLICATION MANAGER', 'Filter by', 'Appliances', and 'Organizations'. The main area displays a timeline of snapshots for 'saphana-autovm7' from 2019-09-24 to 2019-09-26. A modal dialog titled 'Unmount & Delete' is open, asking for confirmation to unmount and delete 'Image_1626906' mounted on 'saphana-autovm7'. The dialog includes a 'Force Unmount?' checkbox and a 'Submit' button.

- This completes the migration process. Uncomment the */etc/fstab* entries for DATA and LOG Volumes.

Recovering a Single Tenant Database using LVM Snapshot

To recover a single-tenant database:

1. Mount the SAP HANA database snapshot image by providing the mountpoint location using AGM (Actifio Global Manager).
2. From Application Manager > Protected Application > Access, select the latest snapshot to recover, and choose **Mount**.



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

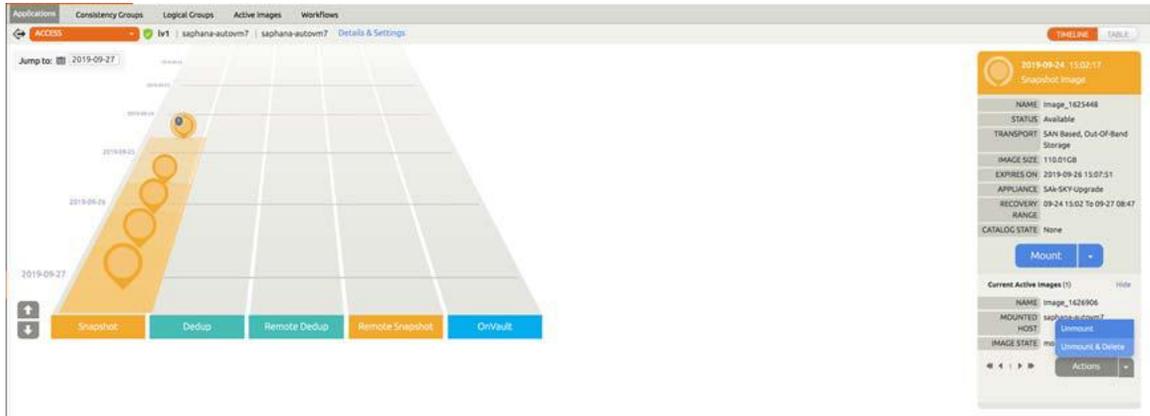
```
cd /act/custom_apps/saphana/restore
```
4. Execute the script for recovery:

```
./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'`
5. Once the script has completed, the Tenant DB is ready for validation.
6. Unmount and delete the backup image.

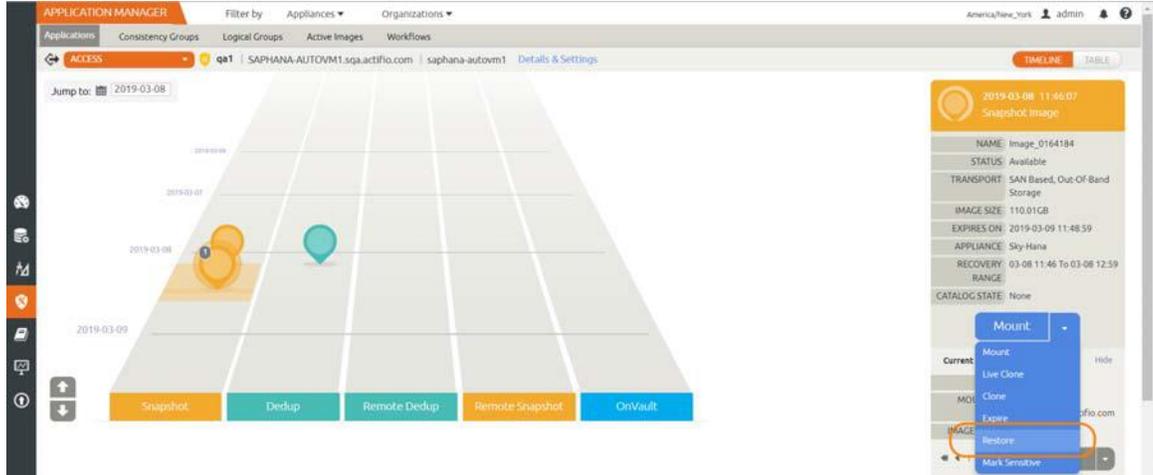


Unmounting and Deleting the Mounted Image

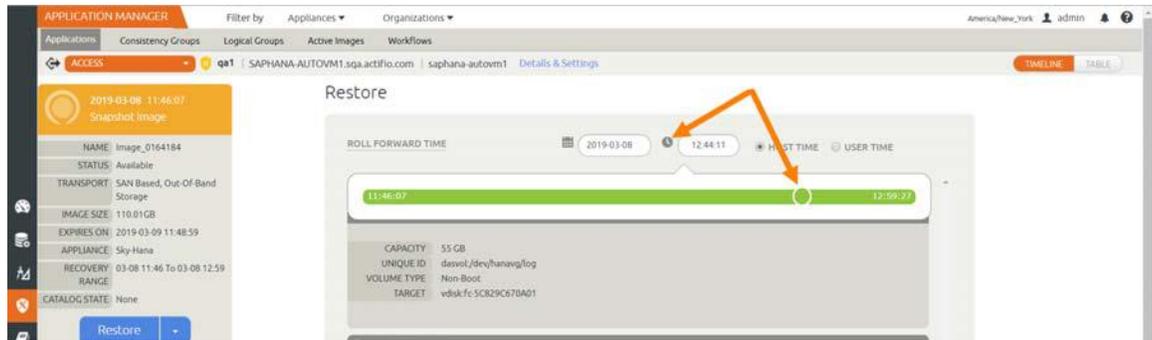
Recovering from a File-Based Backup Back to the Source

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

1. From Application Manager > Protected Application > Access.
2. Select the latest snapshot to recover, and choose **Restore**.

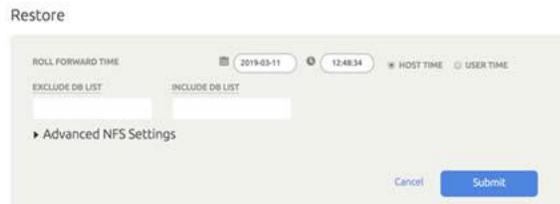


3. For a database protected with logs, on the Restore page, choose a date and a point in time on the slider.



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.



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

Recovering from a File-Based Backup to a New Target

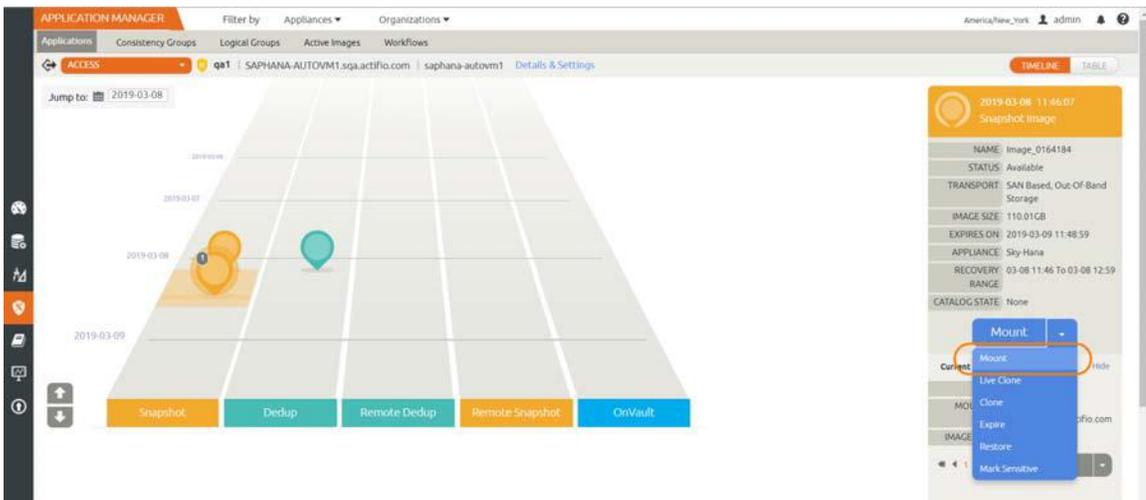
Use this procedure to restore and recover to a new target server.

Before you begin:

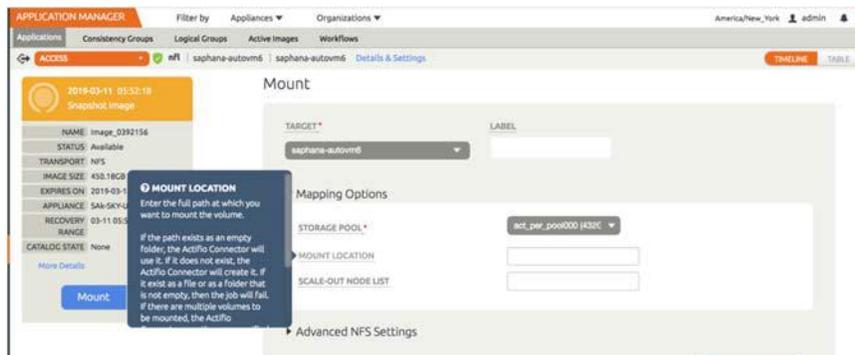
- Make sure the target HANA server is set up the same as the source HANA server (OS version, CPU and memory, HANA version).
- Make sure the HANA database on the target server is configured same as the source, *i.e.* global.ini, nameserver.ini.

To recover to a new target server:

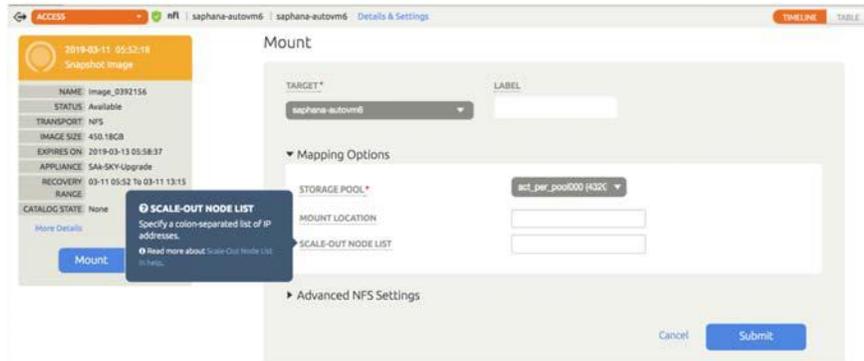
1. From Application Manager > Protected Application > Access, select the latest snapshot to recover, and choose **Mount**.



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 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`

9 HANA Database Management Using actHANADBM

DBAs and developers can use actHANADBM.pl to perform database access tasks using the command line interface. ActHANADBM 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. ActHANADBM.pl is installed on the database server automatically along with the Actifio Connector.

This section includes:

[Installing and Configuring actHANADBM.pl](#) on page 50

[actHANADBM Commands](#) on page 51.

[agmconfig](#) on page 51

[createTemplate](#) on page 52

[hostDiscovery](#) on page 53

[protectApp](#) on page 54

[backup](#) on page 55

[listImageDetails](#) on page 56

[mount](#) on page 57

[unmountdelete](#) on page 58

[restore](#) on page 59

[runwf](#) on page 60

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 50

[Enabling and Verifying Port 443](#) on page 50

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 **Connecting Hosts to Actifio Appliances** in your Actifio Documentation Library.

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 CDS 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>
<runwfm>
```

actHANADBМ Commands

The actHANADBМ commands are:

- [agmconfig](#) on page 51
- [createTemplate](#) on page 52
- [hostDiscovery](#) on page 53
- [protectApp](#) on page 54
- [backup](#) on page 55
- [listImageDetails](#) on page 56
- [mount](#) on page 57
- [unmountdelete](#) on page 58
- [restore](#) on page 59
- [runwф](#) on page 60

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

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 actHANADBMP.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>
```

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

hostdiscovery Parameters

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

protectApp

To protect the application, use -type protectApp.

Example

```
perl actHANADB.M.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>
```

protectApp Parameters

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

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

backup Parameters

Parameters	Use
<code>--appname</code>	Name of the application. This is a required parameter.
<code>--hostname</code>	Name of the application host. If not specified, host where script is running will be used.
<code>--backuptype</code>	Type of backup operation. This is optional parameter. If not specified, the default type is database backup (db)
<code>--jobpriority</code>	The priority for the job. This is optional parameter. Valid inputs are low, medium or high.
<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 value is yes.

listImageDetails

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

Example

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

listImageDetails Parameters

Parameters	Use
<code>--appname</code>	Name of the application. This is a required parameter.
<code>--hostname</code>	Name of the source host. This is a required parameter.
<code>--AGM</code>	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>
[--scaleoutodelist <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 <'yyyymmdd hh24:mi:ss'>]
--AGM <AGM name|ip>
[--wait <yes|no>]
```

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>--scaleoutodelist</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 'yyyymmdd 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>]
```

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>]
```

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>]
```

directmount 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.