
An Oracle DBA's Guide to Actifio GO

Last updated on August 29, 2023

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

Actifio GO

Copyright, Trademarks, and other Legal Matter

Copyright © 2022 Google LLC. All rights reserved.

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

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

Contents

Preface	vii
The ActifioNOW Customer Portal	vii
Actifio Support Centers.....	vii
Chapter 1 – Introduction to Actifio Copy Data Management	1
Capturing Oracle Data.....	2
Replicating Data	3
Accessing Data.....	4
Required Oracle Patches.....	6
Oracle Recommended Patches for dNFS	8
Chapter 2 – Best Practices for Using Actifio GO with Other Backup Products	9
Backup Schedule Clashing	9
Purging the Archivelog.....	10
RMAN Metadata Making Actifio Backups Obsolete	10
Chapter 3 – Actifio Prerequisites for Protecting an Oracle Database	11
Chapter 4 – Oracle Authentication	13
Using Oracle with OS Authentication.....	13
Enabling Database Authentication for an Oracle Server	14
Chapter 5 – Data Capture under File System and under ASM Disk Group	17
Protecting an Oracle Database Under a File System as a File System.....	18
Protecting an Oracle Database Under an ASM Disk Group as an ASM Disk Group	18
ASM Scalability and Limits (from Oracle Doc ID 370921.1).....	18
Protecting an Oracle Database Under a File System as an ASM Disk Group	20
Protecting an Oracle Database Under an ASM Disk Group as a File System	21
Chapter 6 – Preparing Oracle Databases for Protection	23
Preparing Oracle Databases in a Linux Environment Using OS Authentication.....	24
Preparing to Capture a Database from Oracle ASM to Oracle ASM.....	25
Preparing to Capture a Database from Oracle ASM to Filesystem	25
Preparing Oracle Database Authentication in a Linux Environment	25
Enable Database Block Change Tracking (optional).....	27

Protecting from an Oracle Data Guard Node.....	27
Configuring RAC Transparent Failover of Actifio RMAN Backup to Other Nodes.....	28
Oracle Archive Logs Compression.....	31
Manually Calculating Log Staging Disk Size (optional).....	31
Configuring Oracle Database Services for Load Balancing across Multiple Nodes.....	32
Chapter 7 – Details and Settings for Oracle Databases	35
Application Details & Settings for Oracle Databases.....	35
Policy Overrides for Oracle Databases.....	38
Chapter 8 – Configuring dNFS for Protecting and Mounting Virtual Oracle Databases	41
Before You Begin.....	41
Configuring AGM for Protecting and Mounting Virtual Oracle Databases over dNFS.....	42
Actions to be Performed on the Host for dNFS to Work.....	43
For a Virtual Database Mount.....	43
Troubleshooting dNFS: Database Issues.....	44
Alert Log.....	44
Database Trace Files.....	44
Database Hang.....	44
dNFS Views.....	44
The Oracle dNFS Monitor Package.....	48
Chapter 9 – Virtualizing an Oracle Database for Data Protection and Agility	51
Chapter 10 – Accessing, Recovering, or Restoring an Oracle Database	53
Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access.....	54
Mounting an Oracle Database Image Protected Under a File System for Data Access.....	56
Mounting an Oracle Database as a Virtual Application.....	58
Bringing Actifio-Protected ASM Diskgroups Back Online after Reboot of a Target DB Server.....	63
Restoring a Database, Overwriting the Production Database.....	63
Chapter 11 – Recovering an Oracle Database Manually Using RMAN	65
Recovering a Non-RAC Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio.....	66
Recovering a RAC ASM Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio.....	67
Recovering a Non-RAC Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog.....	69
Recovering a RAC ASM Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog.....	71
Recovering an Oracle Database to a Scheduled Backup Point if the archivelog is not Protected through Actifio.....	73
Recovering an Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog.....	75
Chapter 12 – Recovering Tablespace and Data Files	77
Recovering a Single Tablespace of a Production Database on an ASM Disk Group.....	77
Recovering a Corrupt Database Block.....	78
Recovering Lost Control Files.....	79

Recovering an Oracle Pluggable Database.....	80
Chapter 13 – Instant Oracle Database Recovery or Migration Using ASM Switch and Rebalance	81
Chapter 14 – Performing an Oracle ASM Switch and Rebalance	85
Chapter 15 – Protecting and Recovering Oracle Databases in a Windows Environment	91
Preparing Oracle Protection in a Windows Environment	91
Identifying Database Instances On Windows.....	92
Backing Up an Oracle Database in a Windows Environment.....	93
Watch Script to Watch for Database Volumes Being Mounted.....	94
Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point if the archivelog is Not Protected through Actifio.....	95
Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog.....	96
Chapter 16 – Using Actifio VDP with Oracle Exadata Database or Oracle ExaCC	97
Using Actifio VDP with Oracle Exadata in an iSCSI Environment.....	97
Using Oracle Exadata with Actifio VDP in a dNFS Environment.....	98
Chapter 17 – Protecting SAP ECC/BW with an Oracle Database	99
Protecting the Oracle Database.....	99
Refreshing the Database.....	99
After the Refresh, on the Target Oracle Database.....	99
After the Refresh, on the Target SAP Application Server.....	100
Chapter 18 – Oracle RMAN Logs	101
Oracle Protection Logs on Linux.....	101
Oracle Protection Logs on Windows.....	102
Retrieving a Specific Oracle DB Archive Log Sequence Number from an Actifio Backup Image.....	103
Chapter 19 – Introduction to Provisioning Environments With Workflows	105
Workflow Benefits.....	106
Workflow Data Access Methods.....	106
Workflow Access Control.....	107
Configuring Roles, Organizations, and Users.....	107
Example Role for Limited Workflow Access.....	108
Example Organization for Limited Workflow Access.....	108
Example User for Limited Workflow Access.....	109
Chapter 20 – Useful Workflows	111
Direct Mounting Application Data or a Virtual Application.....	112
Creating Scrubbed Data or a Scrubbed Virtual Application with a LiveClone and Masking Tools.....	115
Using an Actifio Workflow to Refresh Oracle Database Schemas.....	119

Before You Begin	120
Creating the Workflow.....	120
Running the Workflow.....	122
Unmounting Mounted Images.....	122
Presenting an Oracle 12c Database PDB as a Virtual PDB to an Existing Database Container on a Target.....	123
Before You Begin	123
Creating a Workflow to Perform the PDB Clone Job.....	123
Running the Workflow.....	125
Unmounting Mounted Images.....	125
Chapter 21 - Workflow Pre and Post Scripts	127
Environment Variables.....	127
Example Script	128
Chapter 22 - Oracle Database Management Using actDBM	129
Installing and Configuring actDBM.pl.....	130
ActDBM Commands.....	132
listImageDetails.....	133
listApplication.....	134
listDiscoveredHost	134
backup.....	135
restore	136
clone.....	137
mount.....	142
cleanup (Unmount and Delete an Image)	143
runwf.....	144
createliveclone.....	146
refreshliveclone.....	147
restoreASMswitch (Instant Oracle Database Recovery).....	148
restoreASMrebalance	149
actDBM.pl Script Template	150
Perl Examples of actDBM Usage and Results.....	154
RESTful API Examples of actDBM Usage and Results	155
Chapter 23 - Best Practices for Application Details & Settings	157
Staging disk size calculation.....	158

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 Oracle databases.

The ActifioNOW Customer Portal

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

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

To log into the ActifioNOW customer portal:

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

Actifio Support Centers

To contact an Actifio support representative, you can:

- Send email to: support@actifio.com
- 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

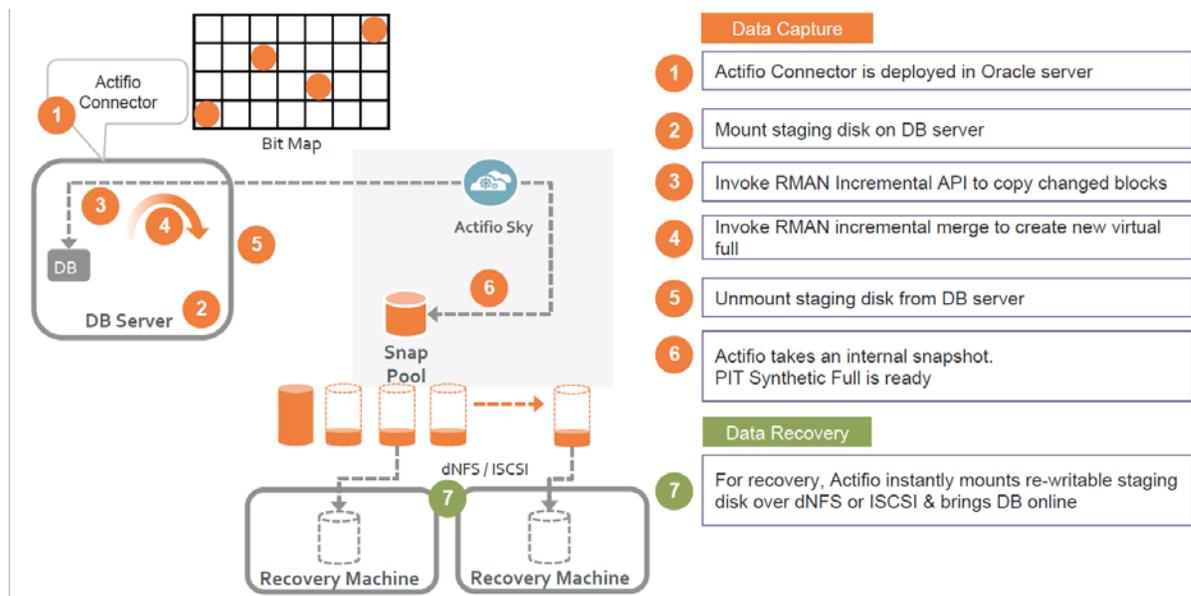
1 Introduction to Actifio Copy Data Management

This chapter provides a high-level overview of basic Actifio concepts and procedures used to capture and access Oracle databases. It includes:

- [Capturing Oracle Data](#) on page 2
- [Replicating Data](#) on page 3
- [Accessing Data](#) on page 4
- [Required Oracle Patches](#) 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 or physical copies of the data whenever and wherever they are needed.



Capture, Manage and Access Application Data

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.

Capturing Oracle Data

Capturing Oracle data consists of four steps:

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

The Actifio Connector

The Actifio Connector is used to capture selected Oracle databases. The Actifio Connector is a small-footprint, lightweight service that can be installed on either virtual or physical servers. The Actifio Connector makes use of Oracle RMAN for capture and access operations.

Specifically, the Actifio Connector:

- Discovers Oracle databases.
- Uses RMAN image copy and incremental merge API to capture data at block level in incremental forever fashion.
- Identifies changes to database data for Actifio's incremental forever capture strategy.
- Captures and manages archivelog:
 - o Captures Oracle database(s) and logs with one SLA.
 - o Purges Oracle database archivelog.
 - o Rolls forward Oracle database archivelog for point-in-time recovery when accessing virtual copies.

Oracle Database Block Change Tracking (BCT)

Oracle tracking enables fast database backups by identifying which blocks have changed. Only changed blocks are included in the backup operation.

- Actifio incremental-forever supports both databases running with BCT enabled and databases running with BCT disabled.
- Change Block Tracking is enabled at database level.
- Oracle records the changed blocks in each data file in a tracking file (small binary file stored in the database area).
- With tracking enabled, RMAN uses the BCT file to get the changed blocks for incremental backup.
- RMAN scans each block in a data file for all data files in the database during incremental backup when Change Block Tracking on the database is not enabled.
- With BCT not enabled the incremental backup time will increase.

Protecting Oracle Databases in an Actifio Consistency Group

In out-of-band configurations (most Sky Appliance and CDX Appliance configurations), a consistency group can contain a single Oracle database application and any number of file system applications from the Oracle server. A consistency group is a good choice for Oracle databases in test/dev and other business agility use cases.

Oracle Databases with TDE

Actifio supports a variety of capture and presentation methods for Oracle databases under various configurations. This includes backup, recovery, and Application Aware mount operations of Oracle database with Transparent Data Encryption (TDE) configured.

For Oracle databases with TDE, wallet files from the source backup host must be available to the target host of any Application Aware mounts. This can be accomplished in several ways.

- The wallet files can be copied from the backup source server to the target mount server and Oracle configured to access them.
- If the Oracle wallet files are stored on a central, shared device on the network, the Appaware mount target Oracle instance should be configured to access them.
- If the Oracle wallet files were captured during backup by setting the Oracle Configuration File Location advanced setting, the wallet files can be retrieved with the following steps:
 1. Do a simple mount of the database to the target host.
 2. Copy the wallet files from the simple database mount to the target host and configure Oracle to use them.
 3. Unmount the simple database from the target host.
 4. Perform an Application Aware mount of the database to the target host.

Note: Application Aware mounts for TDE enabled databases require the wallet to be copied to the appropriate location on the target host and the wallet must be configured and open.

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 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.
- 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 OnVault 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 Oracle database server.

An Actifio Appliance provides two ways to mount an Oracle database:

- The standard mount presents and makes a captured Oracle database backup image copy available to a target server as a file system or as an ASM Disk group depending on the capture method. This is useful for any tablespace/datafile recovery on source or to make a physical copy on target using RMAN duplicate.
- The Application Aware mount presents and makes the captured Oracle data available to a target server as a virtual Oracle database. This allows you to address the unique challenges associated with creating and managing copies of production databases for non-production use. Application Aware mounts are performed from the Actifio Appliance and do not require manual intervention by database, server, or storage administrators. Application Aware mounts can be used for such things as database reporting, analytics, integrity testing, and test and development. Application Aware mounts are described in [Mounting an Oracle Database as a Virtual Application](#) on page 58.

LiveClones

The LiveClone is an independent copy of Oracle 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.

Workflows

Workflows are built with captured Oracle data. Workflows can present data as either a direct mount or as a LiveClone:

- Direct mounts (standard or application aware) work well for Oracle 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 Oracle data without actually moving the data.
- A LiveClone is a copy of your production Oracle data that can be updated manually or on a scheduled basis. You can mask sensitive Oracle data in a LiveClone prior to making it available to users.

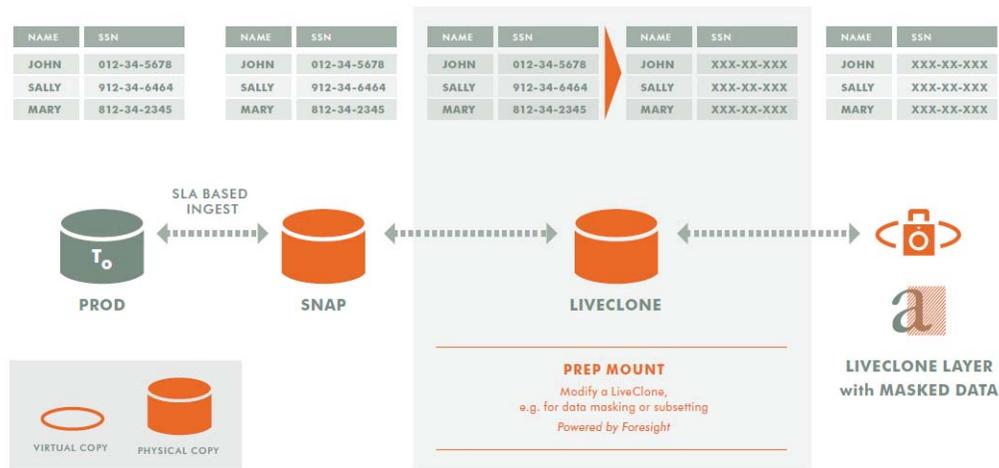
Combining Actifio's automated Oracle data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now users can provision their own environments almost instantly.

For example, an Actifio administrator can create an SLA Template Policy that captures Oracle data according to a specified schedule. Optionally, the administrator can mark the captured production Oracle 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 Oracle data available as a LiveClone or as a direct mount
- Updates the LiveClone or mountable Oracle data on a scheduled or on-demand basis
- (Optional) Automatically applies scripts to the LiveClone's Oracle data after each update. This is useful for masking sensitive Oracle data.

Once the Workflow completes, users with proper access can provision their environments with the LiveClone or mountable Oracle data via AGM.



Workflow With Masked Social Security Data

Workflows are introduced in [Chapter 19, Introduction to Provisioning Environments With Workflows](#), and detailed in the following chapters.

Required Oracle Patches

Table 1: Actifio VDP - Required Oracle Patches

Oracle Version	Needs Patch	Due to these Bugs	Notes
10.2.0.1 to 11.2.0.1	8579188	RMAN RESTORE COMMAND FAILED WITH ORA-1861 (RMAN Recovery Session Fails with ORA-1861 (Doc ID 852723.1))	Actifio Oracle backup may fail Fixed in: 11.2.0.2+
11.1.0.6 to 11.1.0.7	13037524	ORA-01455 Querying v\$asm_disk from database instance raises ORA-01455: converting column overflows integer datatype (Doc ID 1473647), caused by Oracle unpublished Bug 13037524	Fixed in: 11.2.0.1 +. Actifio Oracle backup may fail.
11.2.0.1 to 12.1.0.2.	19621704	ORA-00600 [723] [memory Leak] Error With Leaked Memory For "mbr node memory"	Actifio Application Aware mounts may fail
11.2.0.3	13366202	DBNEWID [nid] does not allow TARGET=/ (NID-106) (Doc ID 13366202.8)	Actifio Application Aware mount may fail
11.2.0.4	28019962	An Oracle database 11.2.0.4 gets ORA-01455 when running a query on V\$ASM_DISK when using 12.2.0.1 ASM Instance (Doc ID 2508802.1)	Patch 28019962 is mandatory for this issue.
12.1.0.2	22709877	ORA-00600: INTERNAL ERROR CODE, ARGUMENTS: [KCVFDB_PDB_SET_CLEAN_SCN: CLEANCKPT]	Actifio Application Aware mount may fail Fixed in: 12.2.0.1+
12.1.0.2 to 19c	26978857	CONTROLFILE BACKUP TO TRACE REFERS PDB DATAFILE OFFLINE IN CDB\$ROOT CONTEXT	Actifio Application Aware mounts may fail.
12.1.0.2	23019710	RMAN LIST BACKUP SUMMARY FAILS With any of: ORA-01507: database not mounted ORA-01219: database or pluggable database not open: queries allowed on fixed ORA-00972: identifier is too long ORA-01589: must use RESETLOGS or NORESETLOGS option for database open ORA-01426: numeric overflow ORA-01152: file 3 was not restored from a sufficiently old backup ORA-01110: data file 3: ORA-01109: database not open ORA-01034: ORACLE not available	Actifio Oracle backups may fail
12.1.0.2	18845653	ORA-600 from PDB close if PDB renamed in another session	Actifio Application Aware mounts may fail.
12.1.0.2	19075256	ORA-600 [kcfmis_internal: enq] from PDB RENAME	Fixed in 12.2.0.1+

Table 1: Actifio VDP – Required Oracle Patches

Oracle Version	Needs Patch	Due to these Bugs	Notes
12.1.0.2	22809813 (Win) 19404068 (Linux)	ORA-1610 ON RECOVER DATABASE FOR CREATED CONTROLFILE	Actifio Application Aware mounts may fail. Oracle Patch release notes: https://updates.oracle.com/Orion/Services/download?type=readme&aru=20122528#BABCGCAB
12.2 to 19c	30045273	PDB: ORA-00600 [KDSGRPI] / BLOCK INCONSISTENCIES AFTER DATABASE RECOVER FROM RMAN IMAGE COPIES USING NEW CONTROLFILE	Actifio Application Aware mounts may fail.
12.2.0.2	31718134	ORA-00304: requested INSTANCE_NUMBER is busy	Restarting new database instance for the first time in no mount state may fail
18c & 19c	3041950	ORA-65250: invalid path specified for file - /act/mnt/Staging_631487/datafile/24v0gutb_1_1	Actifio Oracle database backups may fail. Workaround until patch is installed: Unset parameter CREATE_FILE_DESTINATION for pluggable database.
18.1.0 and later	Check with Oracle for interim patches for earlier versions.	Bug 29056767 - STANDBY: Datafiles Checkpoint not Updated at Standby Database when Media Recover is running (Doc ID 29056767.8)	Actifio Application Aware mounts may fail. Fixed in: 19.4.0.0.190716 Jul 2019 DB RU 18.8.0.0.191015 Oct 2019 DBRU 20.1.0.

Create controlfile with set database does not work with datafile added to PDB Sev 1 SR

This affects an Application Aware mount back to source host, when there were datafile(s) added to the parent database and then captured by archive log backup. An Application Aware mount will attempt to restore the newly-added file back to its source location, corrupting the source datafile. After backup, subsequent Application Aware mounts fail, and the source database will have issues if any of those datafiles are accessed.

This affects both Standalone ASM and RAC, for PDB and for the container itself, on Oracle 12.1.0.2. We suspect filesystem as well, but need confirmation. Actifio tracking number in release notes is 82465.

No patch available; Oracle SR 3-15240183821

Workaround: Take a new database backup with the newly-added file, and then make an Application Aware mount of the new image to a new location. Do not mount back to the source.

Queries on DBA_FREE_SPACE are Slow (Doc ID 271169.1)

During Actifio backup, the Actifio Connector queries the **dba_free_space** Oracle metadata table to determine the database allocated and free space. Sometimes the sql queries to dba_free_space become costly. This is Oracle known issue Doc ID 271169.1 (Queries on DBA_FREE_SPACE are slow). This can be observed on a hung system by running `ps -ef | grep -i dbFreeSize.sql` from the command line.

Oracle recommends to purge the recycle bin from the database:

1. Login to the database as sysdba:

```
sqlplus / as sysdba
SQL>purge dba_recyclebin;
SQL>exit;
```
2. It is a good idea to run the statistics on fixed objects. This can take a few minutes.
Login to the database as sysdba:

```
sqlplus / as sysdba
SQL> exec dbms_stats.GATHER_FIXED_OBJECTS_STATS
SQL>exit;
```

Product Oracle Database - Enterprise Edition, Release Oracle Database 10.1.0.2 and later.
See also [ASM Scalability and Limits \(from Oracle Doc ID 370921.1\)](#) on page 18.

ORA-01157: cannot identify/lock data file % when trying to start a mounted Oracle Database

This is Oracle known issue Doc ID 2183663.1.

See https://support.oracle.com/knowledge/Oracle%20Database%20Products/2183663_1.html

Oracle Recommended Patches for dNFS

Oracle maintains a list of required/recommended patches in the My Oracle Support document "Recommended Patches for Direct NFS Client (Doc ID 1495104.1)".

Unable to open the PDBs (App-mounted) on node-2 after the server reboot.

2 Best Practices for Using Actifio GO with Other Backup Products

Actifio can coexist with legacy products capturing data from production databases, but be sure to follow these best practices:

[Backup Schedule Clashing](#) on page 9

[Purging the Archivelog](#) on page 10

[RMAN Metadata Making Actifio Backups Obsolete](#) on page 10

Backup Schedule Clashing

Database logs are used to capture individual transactions in a database, enabling point-in-time recoveries. Typically, most agility use cases center around getting database snapshots on a periodic basis from production. Common frequency ranges from once a day to once a week or once in two weeks, depending on the use case. As a result, application developers do not commonly have the need to position their non-prod instance to a specific point-in-time from the source (production). This usually eliminates the need to capture and manage logs as a part of an Actifio agility solution.

Requirement	Do not schedule legacy backup software and Actifio to run jobs in a way that would allow any overlap in time..
Best Practice	Schedule Actifio Oracle jobs to begin at a time when the legacy backup software should be finished. Do not schedule the legacy backup software to run immediately after an Actifio job would normally complete.
Reason	If legacy backup jobs and Actifio jobs run concurrently, it may result in a serious performance impact on the database server leading to instability and possibly an outage. Additionally, for Oracle, this may result in invalid backup images for one or both solutions.

Purging the Archivelog

Oracle uses archive logs generated during a database backup to ensure the consistency and recoverability of that backup. As a result, if archive logs are purged during a database backup job, that backup copy will be unrecoverable.

Requirement	Do not allow Oracle archive logs to be purged during an Actifio job, and do not allow Actifio to purge archive logs during a legacy backup RMAN job.
Best Practice	Configure disable archive log purge jobs in the legacy backup software at the start of the Actifio backup job, and resume purge jobs at the end or retain archive log for a minimum of 24 hours before purging.
Reason	If archive logs are purged during an RMAN job, that database backup/image copy will be corrupt and unrecoverable.

RMAN Metadata Making Actifio Backups Obsolete

Actifio database backup is incremental forever. This is achieved by using RMAN image copy with RMAN incremental merge API.

The first RMAN backup is a full image copy of the database datafile on Actifio backup disk with internal snapshot of backup disk.

For the second and all subsequent backups, RMAN incremental backup runs with RMAN incremental merge on the Actifio backup disk, updating the last full with incremental changes before snapshot. However if any third party database backup or crosscheck of backup runs after the Actifio database backup, then all backup datafiles under the Actifio backup are marked obsolete under RMAN metadata.

Note: If the Actifio Application Details & Settings parameter *DO NOT UNCATALOG* is set to Yes, you may get Error: Failed to catalog image copies from staging device and a backup failure. If Actifio GO must co-exist with other backup products, then keep *DO NOT UNCATALOG* set to **NO**.

Requirement	Set Actifio Application "Details & Settings" parameter DO NOT UNCATALOG to NO.
Best Practice	Actifio database backup is incremental forever. This is achieved by using RMAN image copy with RMAN incremental merge API. The first RMAN backup is a full image copy of the database datafile on Actifio backup disk with internal snapshot of backup disk. Subsequent RMAN incremental backup runs with RMAN incremental merge on Actifio backup disk, updating the last full with incremental changes before snapshot. However if a third party database backup or crosscheck of backup runs after the Actifio database backup, then all backup datafiles under Actifio backup are marked obsolete under RMAN metadata. Actifio Application "Details & Setting" parameter "DO NOT UNCATALOG" set to "Yes" results in Error: "Failed to catalog image copies from staging device" and backup failure. Keep "DO NOT UNCATALOG" set to "NO" to co-exist with other legacy backup products.
Reason	By default the parameter DO NOT UNCATALOG in Actifio Application Details & Settings is set to NO. Setting this to YES interferes with other backup products.

3 Actifio Prerequisites for Protecting an Oracle Database

After the Actifio preparation and before you can virtualize and access Oracle databases,

1. Review the concepts in [Chapter 4, Oracle Authentication](#) and [Chapter 5, Data Capture under File System and under ASM Disk Group](#),
2. Prepare the database according to the steps in [Chapter 6, Preparing Oracle Databases for Protection](#).

Step	Where	What	These procedures are in:
1	The Database Server	Install/upgrade the Actifio Connector. Always use the most recent Actifio Connector.	Network Administrator's Guide to Actifio GO
2	AGM Manage > Hosts	The database server must be added as a host or as a VM.	AGM Online Help
3	AGM Oracle Databases Wizard	The database must be discovered as an application.	AGM Online Help
4	SLA Architect	You need one or more suitable SLA templates and resource profiles for the database.	Policy Overrides for Oracle Databases on page 38
5	App Manager	There are many Oracle-specific Application Details & Settings that must be set.	Application Details & Settings for Oracle Databases on page 35

4 Oracle Authentication

This section describes two forms of Oracle database user authentication:

[Using Oracle with OS Authentication](#)

[Enabling Database Authentication for an Oracle Server](#) on page 14

Note: Actifio RMAN backup runs as the Oracle binary owner. If the user running the database instance is not the Oracle OS owner and group access privileges for the users are not the same, then backup fails. In an SAP environment, sometimes the Oracle database instance gets started as sapadmin instead of as the Oracle OS owner account. The right configuration is to start and run the database instance as Oracle OS user who owns the Oracle binary. If the database instance must run as a different user such as sapadmin, then sapadmin and the Oracle OS user should have all the same group access privileges.

Using Oracle with OS Authentication

OS Authentication is the default setting in Linux environments. No database user account and no service name are needed. An Actifio backup uses “/ as sysdba” to connect to the database.

With OS authentication, the backup cannot be run in parallel from multiple nodes in a RAC environment using backup under ASM disk group.

From the AGM Application Details and settings, you can validate the authentication configuration, as shown at right. Application Details & Settings for an Oracle database when OS Authentication is configured include:

- **Number of Channels:** Specify the number of channels for RMAN based on the number of cores on the database server. Consider the number of channels allocated to the other database backup on this server to optimize the channel allocation. The default value is 1 RMAN channel.
- **Oracle Data Guard Primary Node Servicename:** This is required only when you are protecting data from the standby node of an Oracle Data Guard pair. See [Protecting from an Oracle Data Guard Node](#) on page 27. With Data Guard, you also need the database username and password to connect to primary to switch the archive log for consistent database copy during the backup.

The screenshot shows the 'Application Details & Settings' window for an Oracle database. The window title is 'Application Details & Settings'. Below the title, there is a checkbox 'Select options that will revert back to default.' and a 'Settings Help' link. The application name is 'BOSTON'. The configuration details are as follows:

APPLICATION TYPE	Oracle
HOST	Dgvm5.sqa.actifio.com
HOST IP ADDRESS	172.16.15.76
PATH	Dgvm5.sqa.actifio.com
OPERATING SYSTEM	Linux
APPLIANCE	Mastiff
APPLIANCE IP ADDRESS	Turner.sqa.actifio.com

The 'Authentication' section is highlighted with an orange border and contains the following fields and buttons:

- USERNAME: username
- PASSWORD: *****
- Validate Configuration button

At the bottom of the window, there are 'Settings', 'Cancel', and 'Save Changes' buttons.

Note: OS Authentication is not supported in Windows environments.

Enabling Database Authentication for an Oracle Server

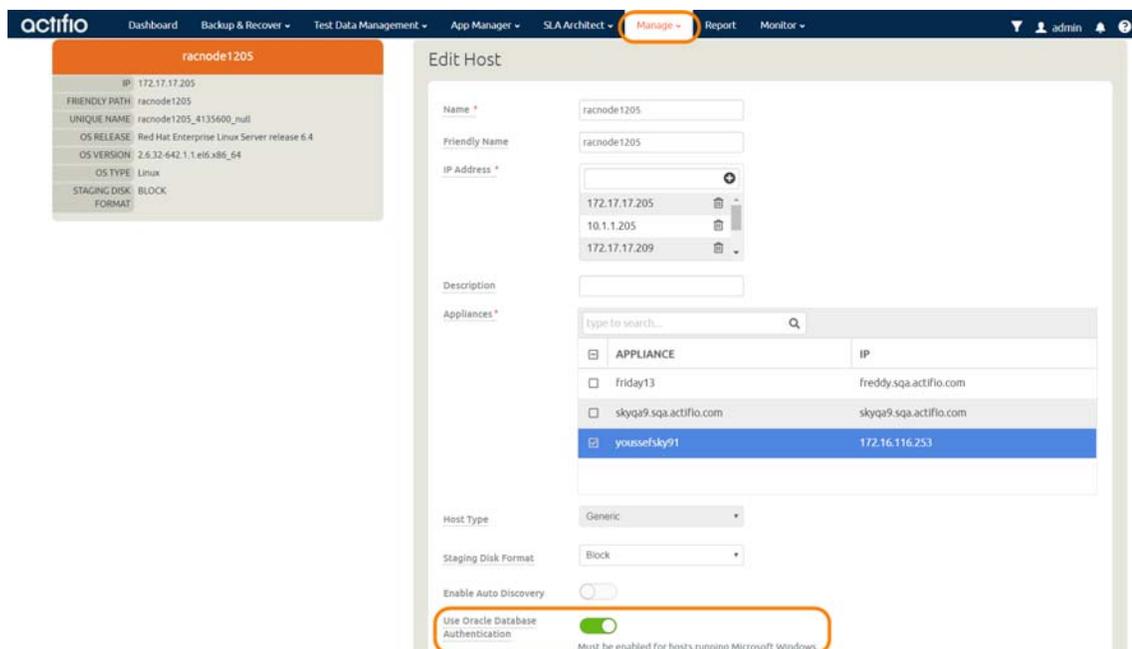
Oracle database authentication uses Oracle database credentials. With Oracle Database Authentication, you must provide database credentials to connect to the database with sysdba privilege (or sysbackup for Oracle 12c).

See Oracle Metalink note: Doc ID 469777.1 for sysdba privilege requirement for RMAN backup.

Enabling Oracle Database Authentication

To enable Oracle Database Authentication:

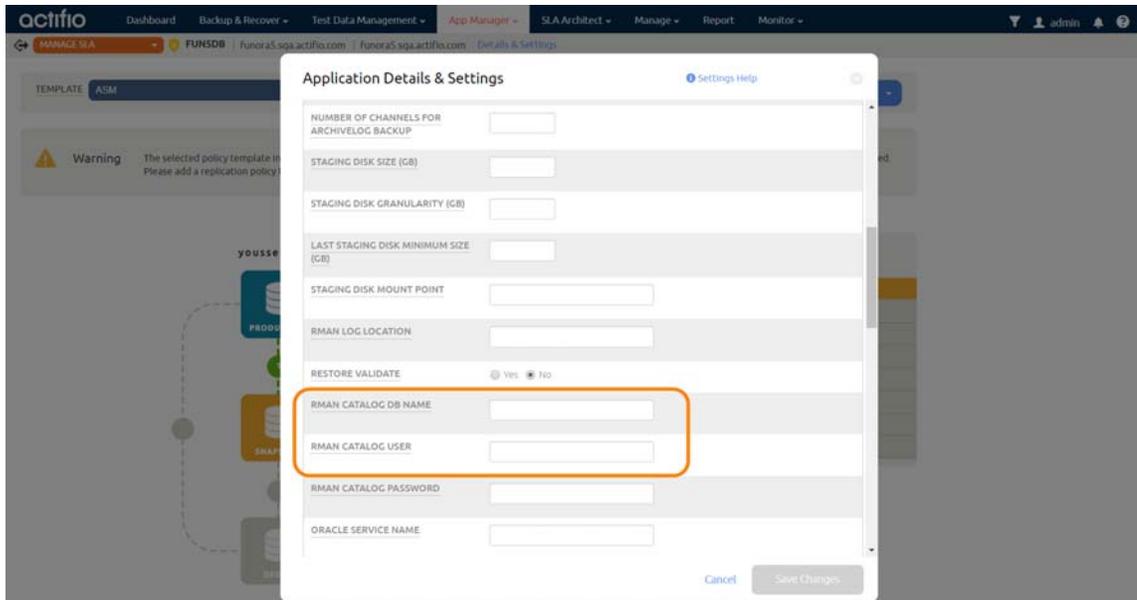
1. Open AGM to Manage > Hosts.
2. Right-click the database server to authenticate against and click **Edit**.
3. On the Edit Host page, slide the **Use Oracle Database Authentication** button to the right. Backup jobs will use Database Authentication for all databases on that database server. This requires a database user account under Application Details & Settings for the backup to succeed.



4. Go to the App Manager and right-click a database on the host that you just enabled for database authentication. Select **Manage SLA**.
5. At the top of the page, click **Details and Settings**.



6. Scroll down to **RMAN Catalog User** and **Password**. Enter the credentials and fill in other settings as required.



7. Repeat [Step 4](#) through [Step 6](#) for each database that will be managed from this database server.

Note: By default the user to connect to the database is sysdba. In an Oracle 12c environment you can choose sysbackup if the database user is granted sysbackup instead of sysdba.

5 Data Capture under File System and under ASM Disk Group

Oracle database capture has different properties depending on whether the images are protected under:

File System: For all source database configurations other than Oracle ASM, the backup is under file system. See [Protecting an Oracle Database Under a File System as a File System](#) on page 18.

ASM Disk Group: If a source database configuration is using ASM, the backup is under an ASM disk group. The Application Details & Settings include ASM configuration options that enable the database to be mounted back to an ASM Disk Group. For more information on protecting an Oracle database under an Oracle ASM disk group, see [Protecting an Oracle Database Under an ASM Disk Group as an ASM Disk Group](#) on page 18.

Note: Oracle backup to ASM is not supported on VMware VMs when the transport method is via NFS datastore (i.e., via the ESX Server). Use RDM directly to the VM.

During the capture, you can convert the database from one format to the other:

From File System to ASM Disk Group: Oracle databases can be protected under ASM Disk Group even if the database being protected is actually on a file system. For more information on protecting an Oracle database from a file system under an Oracle ASM disk group, see [Protecting an Oracle Database Under a File System as an ASM Disk Group](#) on page 20.

From ASM to File System: Oracle databases can be protected as a file system even if the database being protected is actually on an Oracle ASM Disk Group. To protect an ASM database to a file system format, see [Protecting an Oracle Database Under an ASM Disk Group as a File System](#) on page 21.

Table 1: Supported Data Capture and Data Presentation

Database Configuration	Data can be Captured Under	Data can be Presented as an Application Aware Mount as
Database data files under file system or Raw Devices Database data files under file system	File System ASM Disk Group	Standalone File System Standalone ASM or ASM RAC (one or more nodes)
Database data files under RAC or Standalone ASM	File System ASM Disk Group	Standalone File System Standalone ASM or ASM RAC (one or more nodes)

Protecting an Oracle Database Under a File System as a File System

When you capture an Oracle database image under a file system, an Actifio staging disk is mapped to the Oracle server (protected node). A new file system based on file system on the OS is created on an Actifio staging disk (for example, if the source database is on Linux ext4, an ext4 file system will be created).

RMAN image copies of all data files for the entire database will be captured on an Actifio presented file system. A snapshot of the staging disk will be taken.

dNFS is supported, see [Chapter 8, Configuring dNFS for Protecting and Mounting Virtual Oracle Databases](#).

Protecting an Oracle Database Under an ASM Disk Group as an ASM Disk Group

When you capture an Oracle database image under an Oracle backup ASM disk group, an Actifio staging disk is mapped to the Oracle database server and presented to the Oracle ASM layer. An ASM backup disk group is created under ASM using a mapped disk.

The RMAN image copy of all data files for the entire database is captured on an Actifio-presented ASM disk group retaining the ASM header information. A snapshot of the staging disk with ASM header information is taken.

To run backup from	and mount the staging disk to	add RAC member node to
protected node only	protected node only	public IP of protected node
protected node only	more than one node	public IP of protected node first and then public IP of each other node
more than one node	more than one node	public IP of protected node first and then public IP of each other node

To run the backup from more than one node configure tnsnames as described in [Configuring Parallel RMAN Image Copy from Multiple Nodes](#) on page 32.

The Application Details & Settings required for managing databases from an Oracle ASM Disk Group are:

- **Auto Discover RAC Members**
- **RAC Member Nodes** (If auto discovery is selected then RAC Member Nodes is not required. All RAC member nodes will participate.)
- **AU_SIZE**

These are detailed in [Application Details & Settings for Oracle Databases](#) on page 35.

For best results, pay attention to [ASM Scalability and Limits \(from Oracle Doc ID 370921.1\)](#) on page 18.

ASM Scalability and Limits (from Oracle Doc ID 370921.1)

This depends on:

[Oracle Database, Enterprise Edition](#) on page 19

[Oracle Database12c](#) on page 19

[With Oracle Exadata Storage](#) on page 19

[Without Exadata Storage, COMPATIBLE.ASM or COMPATIBLE.RDBMS disk group attribute < 12.1](#) on page 19

[Without Exadata Storage, COMPATIBLE.ASM and COMPATIBLE.RDBMS disk group attributes > 12.1](#) on page 19

Oracle Database, Enterprise Edition

For Oracle Database, Enterprise Edition, Versions 10.1.0.2 to 11.1.0.7 and 11.2.0.3, ASM imposes the following limits:

- 63 disk groups in a storage system
- 10,000 ASM disks in a storage system
- 2 terabyte maximum storage for each ASM disk (the Bug 6453944 allowed larger sizes, but that led to problems, see Note 736891.1 "ORA-15196 WITH ASM DISKS LARGER THAN 2TB")
- 40 exabyte maximum storage for each storage system
- 1 million files for each disk group
- 2.4 terabyte maximum storage for each file

Oracle Database12c

For Oracle Database12c, ASM imposes the following limits:

- 511 disk groups in a storage system for Oracle Database 12c Release 1 or later
- 10,000 Oracle ASM disks in a storage system
- 1 million files for each disk group

With Oracle Exadata Storage

With all Oracle Exadata Storage, Oracle ASM has the following storage limits:

- 4 PB maximum storage for each Oracle ASM disk with the AU size equal to 1 MB
- 8 PB maximum storage for each Oracle ASM disk with the AU size equal to 2 MB
- 16 PB maximum storage for each Oracle ASM disk with the AU size equal to 4 MB
- 32 PB maximum storage for each Oracle ASM disk with the AU size equal to 8 MB
- 320 EB maximum for the storage system

Without Exadata Storage, COMPATIBLE.ASM or COMPATIBLE.RDBMS disk group attribute < 12.1

Without any Oracle Exadata Storage, Oracle ASM has the following storage limits if the COMPATIBLE.ASM or COMPATIBLE.RDBMS disk group attribute is set to less than 12.1:

- 2 terabytes (TB) maximum storage for each Oracle ASM disk
- 20 petabytes (PB) maximum for the storage system

Without Exadata Storage, COMPATIBLE.ASM and COMPATIBLE.RDBMS disk group attributes > 12.1

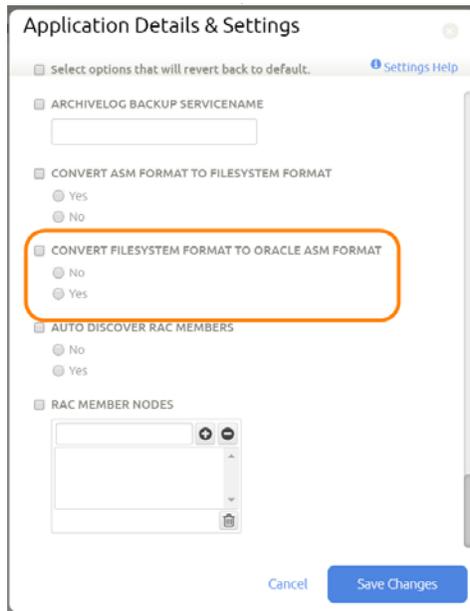
Without any Oracle Exadata Storage, Oracle ASM has the following storage limits if the COMPATIBLE.ASM and COMPATIBLE.RDBMS disk group attributes are set to 12.1 or greater:

- 4 PB maximum storage for each Oracle ASM disk with the allocation unit (AU) size equal to 1 MB
- 8 PB maximum storage for each Oracle ASM disk with the AU size equal to 2 MB
- 16 PB maximum storage for each Oracle ASM disk with the AU size equal to 4 MB
- 32 PB maximum storage for each Oracle ASM disk with the AU size equal to 8 MB
- 320 exabytes (EB) maximum for the storage system

Protecting an Oracle Database Under a File System as an ASM Disk Group

When you capture a file system Oracle database image under an Oracle backup ASM disk group, an Actifio staging disk is mapped to the Oracle database server and presented to the Oracle ASM layer. An ASM backup disk group is created under ASM using a mapped disk.

To protect a specific file system database to ASM Disk group format, in the Application Details & Settings check the check box **Convert Filesystem Format to Oracle ASM Format**. This requires ASM to be installed and running on the protected database node.



The screenshot shows a dialog box titled "Application Details & Settings". At the top, there is a checkbox "Select options that will revert back to default." and a "Settings Help" link. Below this, there are several settings sections:

- ARCHIVELOG BACKUP SERVICENAME**: A text input field.
- CONVERT ASM FORMAT TO FILESYSTEM FORMAT**: Two radio buttons, "Yes" and "No", both unselected.
- CONVERT FILESYSTEM FORMAT TO ORACLE ASM FORMAT**: A checked checkbox. Below it, two radio buttons, "No" and "Yes", with "Yes" selected. This section is highlighted with an orange rectangle.
- AUTO DISCOVER RAC MEMBERS**: Two radio buttons, "No" and "Yes", both unselected.
- RAC MEMBER NODES**: A list box with a plus icon, a minus icon, and a trash icon.

At the bottom right, there are "Cancel" and "Save Changes" buttons.

Application Details & Settings to Capture an Oracle Database under File System to ASM Disk Group

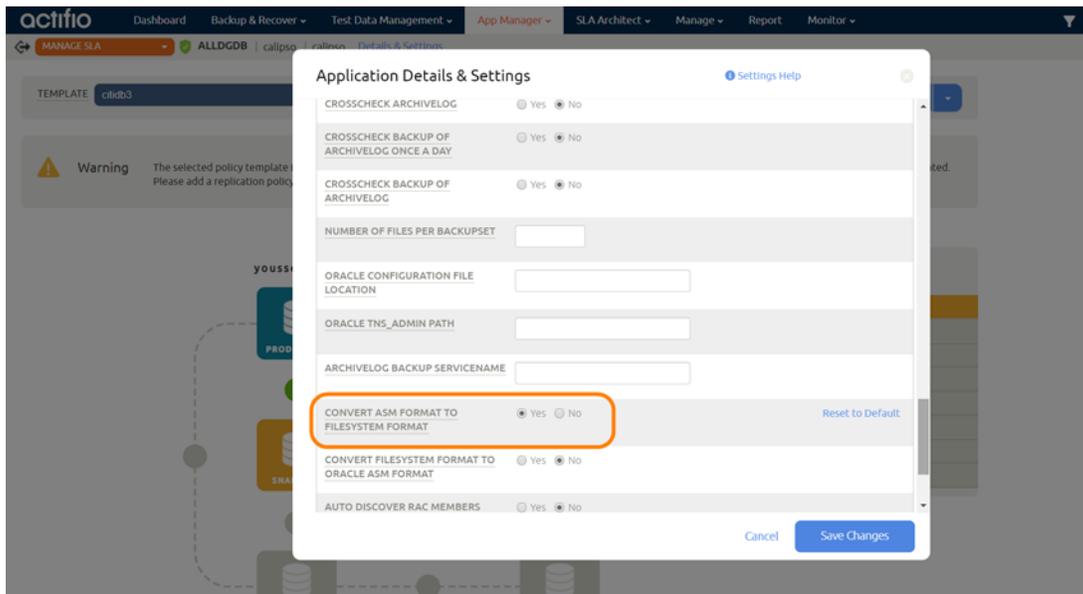
Note: Oracle backup to ASM is not supported on VMware VMs when the transport method is via NFS datastore (i.e., via the ESX Server). Use RDM directly to the VM.

Protecting an Oracle Database Under an ASM Disk Group as a File System

To protect an ASM database to a file system format, in the Application Details & Settings select **Yes** under **Convert ASM Format to Filesystem Format**. File system backup will be used for all source databases including ASM.

Note: Oracle backup to ASM is not supported on VMware VMs when the transport method is via NFS datastore (i.e., via the ESX Server). Use RDM directly to the VM.

If you are protecting an Oracle ASM database to a filesystem, then check that the **Force out-of-band backup** application advanced setting is enabled. Application Details & Settings are detailed in [Application Details & Settings for Oracle Databases](#) on page 35.



Application Details & Settings to Capture an Oracle Database under ASM Disk Group to File System

6 Preparing Oracle Databases for Protection

Before Actifio Appliances can manage Oracle databases, these preparation steps must be performed by a DBA.

Table 1: Preparation Procedures for Oracle Databases in Linux Environments

Step	Preparation Procedure
1	Patching Oracle 12c on page 58 Preparing Oracle Databases in a Linux Environment Using OS Authentication on page 24
2	Preparing to Capture a Database from Oracle ASM to Oracle ASM on page 25 (This is needed only for RAC or Standalone ASM configurations.) Preparing to Capture a Database from Oracle ASM to Filesystem on page 25.
3	Preparing Oracle Database Authentication in a Linux Environment on page 25 Enable Database Block Change Tracking (optional) on page 27 Protecting from an Oracle Data Guard Node on page 27 Configuring RAC Transparent Failover of Actifio RMAN Backup to Other Nodes on page 28 Oracle Archive Logs Compression on page 31 Manually Calculating Log Staging Disk Size (optional) on page 31 Configuring Oracle Database Services for Load Balancing across Multiple Nodes on page 32

Preparing Oracle Databases in a Linux Environment Using OS Authentication

Before protecting an Oracle database, or if database protection jobs fail, make sure that the following settings are correct on the Oracle database server. If you plan to use Oracle Database Authentication, perform these steps first and then go to [Preparing Oracle Database Authentication in a Linux Environment](#) on page 25.

Each Oracle Database to be Protected Must be Running

Each Oracle database to be protected must be up and running. For example:

```
database: actdb
#ps -ef | grep pmon | grep -i actdb
oracle  27688      1  0  2015 ?          00:26:24 ora_pmon_actdb
```

The Database Must Be Running in Archive Log Mode

To verify that the database is running in archive log mode, log into the database server as Oracle OS user and set the database environment variable:

```
export ORACLE_HOME=<oracle home path>
(get this from /etc/oratab)
export ORACLE_SID=<database instance name> (you can get this through ps -ef | grep pmon)
export PATH=$ORACLE_HOME/bin:$PATH
```

Login to sqlplus:

```
#sqlplus / as sysdba
#SQL> archive log list;
Database log mode      Archive Mode
Automatic archival    Enabled
Archive destination    +FRA
Oldest online log sequence  569
Next log sequence to archive  570
Current log sequence    570
#SQL>
```

Note: If archive log mode is not enabled then get archive mode enabled before proceeding.

The Database Should be Using spfile

To verify that the database is running with spfile:

```
#sqlplus / as sysdba
SQL> show parameter spfile
NAME          TYPE VALUE
-----
spfile        string +DATA/ctdb/spfilectdb.ora
```

Note: If the value is **null** then get the spfile set. Actifio supports backing up using pfile as well. pfile should be available in default location. For example, a Linux pfile should be located under \$ORACLE_HOME/dbs.

For RAC under ASM, the Snapshot Control File Must Be Located Under Shared Disks

For an Oracle RAC database under ASM, the snapshot control file must be located under shared disks. To check this, connect to RMAN and run the show all command. Configure it if necessary:

```
RMAN target /
RMAN> show all
```

RMAN configuration parameters for database with db_unique_name CTDB are:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/DATA1/ctdb/snapcf_ctdb.f';
```

Configure it if necessary. For example, the above example is set to Local. To make it shared, use:

```
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '+<DG name>/snap_<DB name>.f';
```

Preparing to Capture a Database from Oracle ASM to Oracle ASM

The ASM diskstring Parameter Must Be Set

If you are using Oracle ASM protection out-of-band, then check that the ASM diskstring parameter is not null. Log into the database server as ASM OS user and set the ASM environment variable:

```
# export ORACLE_HOME=<oracle ASM home path> (get this from /etc/oratab)
#export ORACLE_SID=<ASM instance name> (you can get this through ps -ef | grep pmon)
#export PATH=$ORACLE_HOME/bin:$PATH
```

Connect to sqlplus:

```
#sqlplus / as sysasm
#sql> show parameter asm_diskstring
NAME          TYPE  VALUE
-----
asm_diskstring string ORCL:*, /dev/sdt1, /dev/sdu1
```

If the result of value is null, then get the correct ASM disk string value for existing ASM disks before proceeding with Actifio protection. The Actifio backup will add its diskstring path (/dev/actifio/asm/*) for its backup staging disk to map to ASM.

Note: For Oracle 10g, make sure the kfed utility is configured in the grid home. If it is not configured, configure kfed tool using Oracle Metalink Document ID 1346190.1.

Preparing to Capture a Database from Oracle ASM to Filesystem

The Force Out-Of-Band Setting Must Be Enabled

If you are protecting an Oracle ASM database to a filesystem, check these Application Details & Settings (see [Application Details & Settings for Oracle Databases](#) on page 35):

- **Force out-of-band backup** application advanced setting is enabled.
- **Convert ASM to File System** is set to Yes.

Preparing Oracle Database Authentication in a Linux Environment

These additional preparation steps are required only if you will use database authentication. Oracle database authentication is described in [Chapter 4, Oracle Authentication](#).

1. Follow the steps in [Preparing Oracle Databases in a Linux Environment Using OS Authentication](#) on page 24.

2. Create a database user account for Actifio backup (if not provided):

```
sql> create user act_rman_user identified by <password>;
```
3. Grant sysdba access to all RAC nodes by logging into sqlplus to all nodes and running:

```
sql> grant create session, resource, sysdba to act_rman_user;
```

For Oracle 12c this role can be sysbackup instead of sysdba, and the database user name starts with #.
4. Verify that the sysdba role has been granted on all nodes in the RAC environment:

```
#sqlplus / as sysasm
# sql> select * from gv$pwfile_users;
INST_ID USERNAME SYSDB SYSOP SYSAS
-----
```

INST_ID	USERNAME	SYSDB	SYSOP	SYSAS
1	SYS	TRUE	TRUE	FALSE
2	SYS	TRUE	TRUE	FALSE
1	ACT_RMAN_USER	TRUE	TRUE	FALSE
2	ACT_RMAN_USER	TRUE	TRUE	FALSE
5. Test the service name as described in:
[Creating and Verifying the Oracle Servicename in a non-RAC Environment](#) on page 26
[Creating and Verifying the Oracle Servicename in a RAC Environment](#) on page 26

Creating and Verifying the Oracle Servicename in a non-RAC Environment

The Oracle Servicename is used for database authentication only. It is not needed for OS authentication.

Example: Database name: dbstd, Instance Name: dbstd

1. If the Oracle Servicename is not listed, then create the service name entry in the `tnsnames.ora` file at `$ORACLE_HOME/network/admin` or at `$GRID_HOME/network/admin` by adding the entry:

```
act_svc_dbstd =
(DESCRIPTION =
(AADDRESS = (PROTOCOL = TCP)(HOST = <IP of the database server>)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dbstd)
) )
```

If the `tnsnames.ora` file is in a non-standard location, then provide the absolute path to it in the Application Details & Settings described in [Application Details & Settings for Oracle Databases](#) on page 35.
2. Test that the service name entry for the database is configured:
Login as Oracle OS user and set the Oracle environment:

```
TNS_ADMIN=<tnsnames.ora file location>
tnsping act_svc_dbstd
```
3. Check the database user account to be sure the Actifio backup can connect:

```
sqlplus act_rman_user/act_rman_user@act_svc_dbstd as sysdba
```
4. Provide the service name created (`act_svc_dbstd`) under the Oracle Service Name setting in Application Details & Settings described in [Application Details & Settings for Oracle Databases](#) on page 35.

Creating and Verifying the Oracle Servicename in a RAC Environment

The Oracle Servicename is used for database authentication only. It is not needed for OS authentication.

Example three-node RAC:

Database name: dbrac, Instance1 name: dbrac1, Instance2 name: dbrac2, Instance3 name: dbrac3 with database protection being set from Node3 (Instance name dbrac3):

1. Create a Servicename Entry in tnsnames.ora file at \$ORACLE_HOME/network/admin or at \$GRID_HOME/network/admin by adding the entry:

```
act_svc_dbrac3 =
(DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP)(HOST = <IP of the database server>)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(INSTANCE_NAME = dbrac3)
(SERVICE_NAME = dbrac)
) )
```

Where:

HOST = This can be SCAN IP in a RAC environment or VIP or IP of the node 3 database server.

SERVICE_NAME = database name

INSTANCE_NAME = database instance name on node3

2. Test the service name entry created above:

Login as Oracle OS user and set the Oracle environment:

```
TNS_ADMIN=<tnsnames.ora file location>
```

```
tnsping act_svc_dbrac3
```

3. Check the database user account to be sure the Actifio backup can connect:

```
sqlplus act_rman_user/act_rman_user@act_svc_dbrac3 as sysdba
```

4. Provide the service name created (act_svc_dbrac3) under the Oracle Service Name setting in Application Details & Settings ([Application Details & Settings for Oracle Databases](#) on page 35).

If the tnsnames.ora file is in a non-standard location, then provide the absolute path to the tnsnames.ora file under the Oracle TNS_Admin Path setting in the Application Details & Settings described in [Application Details & Settings for Oracle Databases](#) on page 35.

Enable Database Block Change Tracking (optional)

To check if database block change tracking is enabled:

```
#sqlplus / as sysdba
#sql>select * from v$block_change_tracking;
```

```
STATUS  FILENAME          BYTES
-----
DISABLED
```

Note: Tracking is optional. Oracle Standard Edition and Oracle Express Edition do not support tracking. Tracking is described in [Oracle Database Block Change Tracking \(BCT\)](#) on page 2.

If tracking is not enabled, then enable database block change tracking from sqlplus:

Using ASM Disk Group

```
sql>alter database enable block change tracking using file '<ASM Disk Group Name>/<database name>/<dbname>.bct';
```

Using File System

```
sql>alter database enable block change tracking using file '$ORACLE_HOME/dbs/<dbname>.bct';
```

Protecting from an Oracle Data Guard Node

You can protect an Oracle database from primary database nodes or from Oracle Data Guard standby nodes. If protection is set from an Oracle Data Guard node, then make sure to set the primary node credentials in Application Details & Settings.

For Database Authentication

Username/Password: The database user account credentials. In order for this user account to be available on the Data Guard node with sysdba access, this user must be created with sysdba privilege at the Primary node (see creating backup user account with sysdba access). Then the password file (under \$ORACLE_HOME/dbs/) from the primary node must be copied over to the Data Guard node.

For OS Authentication

Username/Password: Under OS Authentication, sysdba privilege is not required. This database user account needs "connect, alter system privilege" In order for this user account to be available on the Data Guard node, this user must be created at the primary node.

Grant "connect, alter system" access:

```
sql> grant connect, alter system to act_rman_user;
```

If the user does not have the sysdba role, then the user also needs:

```
> grant select on dba_tablespaces to act_rman_user;
```

This is to allow gathering info of READONLY tablespaces during backup.

Oracle Data Guard Primary Node Servicename: This is the servicename in the tnsnames.ora file configured on the Data Guard node to connect to the primary node from the standby node.

For full details on all Details & Settings, see [Application Details & Settings for Oracle Databases](#) on page 35.

Table 2: Authentication to Data Guard Primary and Secondary Nodes

Database Node	OS Authentication	Database Authentication
Primary	No database credentials are needed.	Database credentials are needed. If no role is selected, then sysdba is used
Standby	Database credentials are needed even for OS Auth (to connect to primary to switch log). The database credentials do not need to have a sysdba/sysbackup role. If a sysdba/sysbackup account is used, then set the user role in User Role in the Database in the Application Details & Settings.	Database credentials are needed. The database credentials must be for either the sysdba or sysbackup role, and User Role in the Database must be set to sysdba or sysbackup in the Application Details & Settings. To set up database authentication, see Enabling Database Authentication for an Oracle Server on page 14.

Configuring RAC Transparent Failover of Actifio RMAN Backup to Other Nodes

The Actifio Connector must be installed and running on all nodes that will be part of the backup failover configuration. The protection is set up from one node only.

Note: In an Oracle One Node environment, both nodes must be discovered and protected with the same template and profile.

In Details & Settings, Cluster Nodes, specify the failover node choice in a Oracle RAC environment:

```
<Failover choice>:<Node IP>:<Servicename>:<Role>
```

Where:

Failover Choice: the order of node in which to fail over.

Node IP: the IP address of the node where you want the backup to run

Servicename: the name of the service created and specified in the tnsnames.ora for Actifio RMAN backup. This can be a new dedicated service created for Actifio backup or the SID name (instance name) of the database on that node.

Role: F, indicating it is a failover node

To create a new servicename on failover node under tnsnames.ora file (\$ORACLE_HOME/network/admin/tnsnames.ora or at \$GRID_HOME/network/admin/tnsnames.ora)

Example in an Oracle One Node Environment

RAC One Node consists of two nodes:

172.15.157.200

172.15.157.201

It has one database OneN running only at one of the nodes at any given time. OneN is protected from 172.15.157.200, with Cluster Node settings specified as 1:172.15.157.201:OneN:F

If OneN fails over to 172.15.157.201, Actifio backup follows it and starts the next backup job from 172.15.157.201 instead of 200. If failover occurs in the middle of a backup job, then the job fails and the next job uses the failover node to start new backup.

Example in an Oracle Environment, NOT One Node

- 2 node RAC (dbrac1, dbrac2)
- Protection is set using database name "dbrac" from dbrac1 and failover is to be set to dbrac2
- Service name on node2: act_svc_dbrac2
- Node2 IP or scan IP: 172.1.1.0

act_svc_node2 =

```
(DESCRIPTION =  
(ADDRESS = (PROTOCOL = TCP)(HOST = 172.1.1.0)(PORT = 1521))  
(CONNECT_DATA =  
(SERVER = DEDICATED)  
(INSTANCE_NAME = dbrac2)  
(SERVICE_NAME = dbrac)  
) )
```

In Application Details & Settings, the Cluster Nodes entry will be:

Failoverchoice:NodeIP:Servicename:Role

1:172.1.1.1:act_svc_node2:F

The Behavior of Cluster Node Entries F and M

F: Failover node, only participates when protecting node is not able to perform the backup.

M: Maintenance node, replaces protecting node if validated to be able to perform the backup.

Table 3: Examples of Cluster Node Configuration and Behavior

Cluster Nodes Entry	Behavior of This Configuration
<p>1: <Node IP Address><ServiceName>: F 2: <Node IP Address><ServiceName>: F</p>	<p>If only F (failover) nodes are present in the list, the first-column order of 1,2,3 will be followed when checking the next backup node.</p> <p>For example, when the primary backup node is unavailable (host down, connector not running, db down, service not running etc), then cluster node validation will go to the F nodes in the order in the first column; when a node is validated that can be used for backup, it is chosen and backup starts from that node. Validation follows the numerical order until one node is validated.</p>
<p>1: <Node IP Address><ServiceName>: M 2: <Node IP Address><ServiceName>: F <i>or</i> 1: <Node IP Address><ServiceName>: F 2: <Node IP Address><ServiceName>: M</p>	<p>When an M (maintenance) node is specified in the cluster node list, the backup will run on the M node (even if the protecting node is able to take the backup).</p> <p>If the M node can not be validated, normal validation is performed on the protection node and on the failover node. If the protection node is validated, then it is used for backup, otherwise the failover node goes through the validation process and acts as the backup node.</p>
<p>1: <Node IP Address><ServiceName>: M 2: <Node IP Address><ServiceName>: M</p>	<p>The first M entry is used to replace the protecting node if validated. The second entry is ignored (there should be only one M entry in the cluster node list, if that's the intention, as it replaces the protecting node for backup).</p>

Oracle Archive Logs Compression

Actifio archivelog backup supports Oracle log backupset compression. The type of compression you select depends on these RMAN configuration settings. Select an option based on your use case.

- Lower compression ratios create the least impact on backup throughput. These are best suited for environments where CPU resources are the limiting factor.
- Medium compression is recommended for most environments. These provide a good combination of compression ratios and speed.
- High compression ratios are resource intensive and are best suited for backups over slower networks where the limiting factor is network speed.

The default setting is BASIC. BASIC does not require Oracle Advanced Compression. RMAN 11G offers a wider range of compression levels with the Advanced Compression Option (ACO).

To check the type of compression set in the environment, run "SHOW ALL" from an RMAN prompt:
rman> show all

Table 4: Selecting an Oracle Compression Algorithm

Algorithm Name	Use This For	Oracle Versions
BASIC	good compression ratio	9.2.0.0 and later
BZIP2	good compression ratio	9.2.0.0 and later
LOW	maximum possible compression speed	11.2.0.0 and later
ZLIB	balance between speed and compression ratio	11.0.0.0 and later
MEDIUM	balance between speed and compression ratio	11.0.0.0 and later
HIGH	maximum possible compression ratio	11.2.0.0 and later

To configure the compression algorithm, use the Algorithm Name from the table above:

```
RMAN> CONFIGURE COMPRESSION ALGORITHM '<alg_name>';
```

RMAN compresses the backupset contents before writing to disk. No extra decompression steps are required during recovery for RMAN compressed backup.

Manually Calculating Log Staging Disk Size (optional)

The Actifio Connector calculates the log staging disk size based on the high water mark of last 60 days of archive generation. In case of specific behavior of archive generation rate, you can specify log staging disk size under the Application Details & Settings, detailed in [Application Details & Settings for Oracle Databases](#) on page 35.

To calculate the archive size and archive generation rate:

1. As Oracle OS user: set the database environment (ORACLE_HOME, ORACLE_SID, PATH).
2. To check the current total log size, connect as sysdba from sqlplus:

```
Sqlplus / as sysdba  
Sql> select sum(blocks*block_size)/(1024*1024*1024) from v$archived_log where deleted = 'NO';
```
3. Check the archive generation rate for (sixty) days:

```
Sqlplus / as sysdba
```

```

Sql>col Day format a10
Sql>col NB_SWITCHS format 9999999999999999
Sql>col TOTAL_SIZE_GB format 999999999999999999
Sql>col AVG_SWITCHS_PER_HOUR format a22
Sql>set pagesize 1000
Sql>SELECT trunc(first_time) DAY,
count(*) NB_SWITCHS,
trunc(count(*)*log_size/1024)/(1024*1024) TOTAL_SIZE_GB,
to_char(count(*)/24, '9999.9') AVG_SWITCHS_PER_HOUR
FROM v$loghist,
(select avg(bytes) log_size from v$log)
where first_time > sysdate - 60
GROUP BY trunc(first_time),log_size
order by 1;

```

Configuring Oracle Database Services for Load Balancing across Multiple Nodes

This procedure applies only to Oracle ASM databases protected out-of-band. In this example, assume a four-node RAC environment; nodes 3 and 4 are to be load-balanced for backup use.

See:

- [Configuring Parallel RMAN Image Copy from Multiple Nodes](#)
- [Configuring Oracle Database Services for Load Balancing across Multiple Nodes](#)

Configuring Parallel RMAN Image Copy from Multiple Nodes

In a RAC environment, you can configure backup to run in parallel from multiple nodes.

1. Install the Actifio Connector on all nodes.
2. Setup the ASM disk group mapping to node 3 and node 4 using Application Details & Settings.
3. Create a database service using srvctl to run from node 3 and node 4.
4. Use this service to specify under Application Details & Settings. Choose Number of channels under Advance Settings (# of Channels). RMAN will distribute the channels between node 3 and node 4.
5. Also set Oracle Servicename and RAC Member Nodes.

Configuring Oracle Database Services for Load Balancing across Multiple Nodes

1. Configure in Application Details & Settings, RAC Member Nodes: IP of node3 and IP of node 4.
2. Create a database service for the maintenance node to be used by Actifio for backup:

```

srvctl add service -d <dbname> -s act_service_<dbname> -r <dbinstance3>,<dbinstance4>
srvctl start service -d <dbname> -s act_service_<dbname>

```
3. Add the tns entry for the Oracle service name created on backup nodes (dbinstance3 and dbinstance4 node in this example) under tnsnames.ora file (\$ORACLE_HOME/network/admin/tnsnames.ora or at \$GRID_HOME/network/admin/tnsnames.ora)

```

act_service_<dbname> =
(DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP)(HOST = <SCAN IP>)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = <DATABASE NAME>)
) )

```
4. Test the servicename created above: `tnsping act_service_<dbname>`
5. Test the service name and user credentials:

```

sqlplus act_rman_user/act_rman_user@act_service_<dbname> as sysdba

```

6. Specify this servicename under Application Details & Settings Oracle Servicename.
7. Create a dedicated Archivelog Backup service on a protected node (e.g. node 3) to be used for backup:

```
srvctl add service -d <dbname> -s act_arc_service_<dbname> -r <dbinstance3>  
srvctl start service -d <dbname> -s act_arc_service_<dbname>
```
8. Add the tns entry for the Archivelog Backup service name created under tnsnames.ora file (\$ORACLE_HOME/network/admin/tnsnames.ora or at \$GRID_HOME/network/admin/tnsnames.ora)

```
act_arc_service_<dbname> =  
(DESCRIPTION =  
(ADDRESS = (PROTOCOL = TCP)(HOST = <SCAN IP>)(PORT = 1521))  
(CONNECT_DATA =  
(SERVER = DEDICATED)  
(INSTANCE_NAME = <node 3 instance>)  
(SERVICE_NAME = act_arc_service_<dbname>)  
) )
```
9. Test the servicename created above: `tnsping act_arch_service_<dbname>`
10. Specify this servicename under Application Details & Settings Archivelog Backup Servicename.

7 Details and Settings for Oracle Databases

There are two kinds of advanced settings:

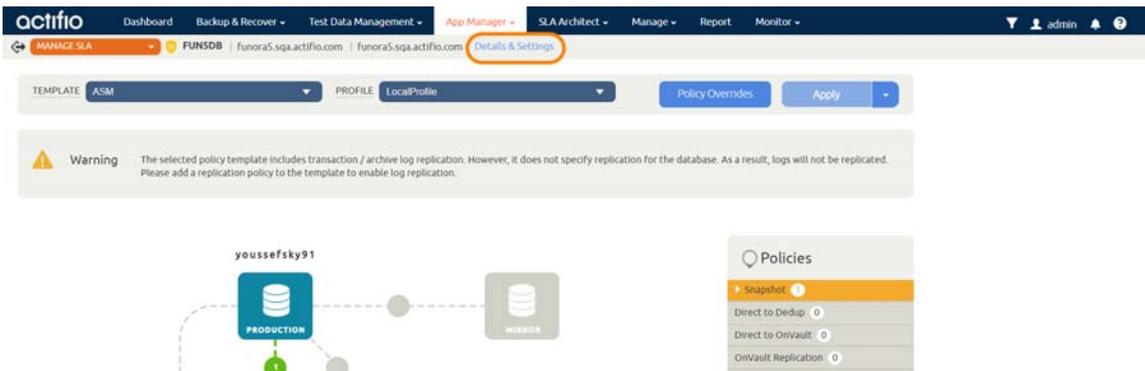
[Application Details & Settings for Oracle Databases](#) on page 35

[Policy Overrides for Oracle Databases](#) on page 38

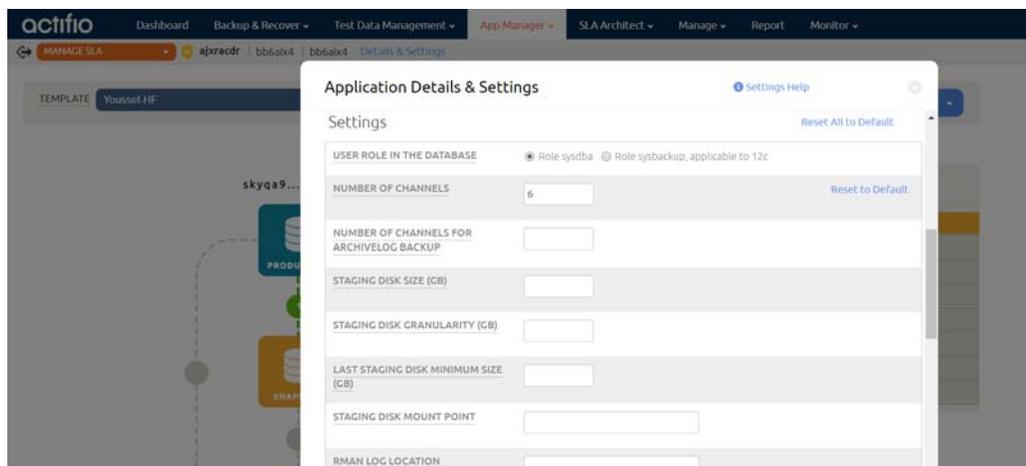
Application Details & Settings for Oracle Databases

To set database-specific details:

1. Go to the App Manager > Applications and right-click a database. Select **Manage SLA**.
2. At the top of the page, click **Details & Settings**.



This opens the Policy Settings page, where you can enter all of the application-specific details and settings below.



The Application Details & Settings for an Oracle database are:

- **Username and Password** (in the Authentication section): When OS Authentication is not or cannot be employed, enter an Oracle user `act_rman_user` username and password for database authentication. Make sure the database user account has the proper role granted based on the **User Role in the Database** below.
- **User Role In The Database**: In all cases except an Oracle Data Guard standby node, the default value is `sysdba`, but select `sysbackup` for an Oracle 12c database. Standby has no default value.
- **Number Of Channels**: Enter the number of RMAN channels based on the host computing power. Number of channels should be configured based on # of cores available on the server, taking into account other database backups configured to run in parallel. The default number of channels is one.
For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Number Of Channels for Archivelog Backup**: Specify the number of RMAN channels (parallel copy processes) to use during archivelog backup.
- **Staging Disk Size**: By default, the connector calculates the size as 1.5 times the maximum size of the database. To specify a value manually, allocate a staging disk to allow for two years future growth of the database. Do not confuse this entry with Log Staging Disk Size, below.
For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Use Staging Disk Granularity as Minimum Staging Disk Size**: Use this for applications that are under the size of the granularity setting and that tend to periodically grow. This option is useful to avoid frequent costly full backups. Because the staging disk is thin provisioned, there is no initial cost to use a staging disk that is larger than required for immediate use.
The default values are 0 for No and the Staging Disk Granularity setting for Yes.
- **Staging Disk Granularity**: Maximum size of each staging disk when multiple staging disks are used for an application. The default value is 1000GB.
- **Last Staging Disk Minimum Size**: Minimum size of the last staging disk created for an application with multiple staging disks. This value is also used for additional disks allocated to accommodate growth. The default value is 250GB.
- **Staging Disk Mount Point**: Allows you mount the staging disk to a specific location.
- **RMAN Log Location**: By default the rman log location is `/var/act/log/rman<db name>.log`. This entry allows you to change the RMAN log file location. enter the full path, with RMAN filename.
- **Restore Validate**: RMAN provides restore validation for the backups. When this box is checked, the connector will invoke RMAN restore validation for each backup. This validation will add time to the backup.
- **RMAN Catalog DB Name**: Optional: This is the CATALOG database SID name. This is for the user environment where RMAN CATALOG DATABASE is set up for RMAN backup. The CATALOG database SID name must have an entry in the `tnsnames.ora` file for Actifio to connect.
- **RMAN Catalog User** and **RMAN Catalog Password**: Catalog database user name/password for RMAN.
- **Oracle Service Name**: Provides the ability to specify a new service name in `tnsnames.ora` file to be used by Actifio backup, as described in [Creating and Verifying the Oracle Servicename in a non-RAC Environment](#) on page 26 and in [Creating and Verifying the Oracle Servicename in a RAC Environment](#) on page 26. If not specified, then by default Actifio will use the Oracle SID name (instance name) as the service name. Either the new service name or the default SID name must have an entry in the `tnsnames.ora` file for Actifio to connect. The Oracle Servicename is used only with database authentication.
- **Oracle Data Guard Primary Node Service Name**: This is the service name in the `tnsnames.ora` file configured on the Data Guard node to connect to the primary node from the standby node. This is required only when you are protecting data from Oracle Data Guard. For more information, see [Protecting from an Oracle Data Guard Node](#) on page 27.

- **Cluster Nodes:** Specify a failover node choice in format:
Failover choice:Node IP:service:role.
This is used for RAC only, see [Configuring RAC Transparent Failover of Actifio RMAN Backup to Other Nodes](#) on page 28.
Example: 1:172.16.16.21:svc_orarac2_act:F
role should be **F** (failover). role can also be **M** (maintenance). When an appliance member role is M, then the Actifio Appliance uses this as the backup node instead of using the original protected node.
- **Connector Options:** Use this only under the direction of Actifio Support.
- **Log Purging Retention Period:** In the space provided, enter the number of hours to retain archive logs in the primary log destination. For example, if this is set to 4, then archive logs older than four hours will be purged from the database primary archive destination. The default value is 24 hours. For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).

Note: If you set **Log Purging Retention Period** to 0, then the log will be purged immediately after the backup job is finished. If you do this, set **Successful Log Backups Before Purge** to at least 1.

- **Successful Log Backups Before Purge:** By default, archive purging does not check for the number of successful log backups. Enter a number of successful log backups after which to run the archive purge. For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Maxcorrupt Parameter Setting:** RMAN backup will continue with backup, skipping this number of corrupted data blocks in each datafile. By default this value is 0 and backup will fail if there is any corrupt data block in any data file.
- **AU_SIZE:** AU_SIZE: Parameter to configure ASM Diskgroup AU size, in MB, default is 4MB. This only takes effect during diskgroup creation, which is during level 0 job. Set this before the first snapshot, or select **Force new level 0** to recreate the disk group (be sure to have enough free space when using this option).
For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Section Size Override:** Section size for multisection image copy backup for 12c or higher, in MB. Default section size is 16000. Enter a new size (1-200000) if you wish to override the default value..
- **Log Staging Disk Size:** Used if log backup policy is set. By default Actifio uses the 30-day high-water mark to determine the staging disk size for archive backup staging disk. To specify a value, refer to [Manually Calculating Log Staging Disk Size \(optional\)](#) on page 31 for more information on determining this value.
- **Do Not Uncatalog:** To keep RMAN datafile backup cataloged after each backup job. By default, Actifio datafile backup will be cataloged at the start of backup and then be uncataloged at the end of the backup. Archivelog is not cataloged.
For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Force New Level 0 Backup:** If for any reason a full level 0 backup is required, overwriting the Actifio incremental backup, then check this box for a single backup job. Be sure to **uncheck** it after the full level 0 backup is complete, or else this will force each backup to be a new level 0 Oracle RMAN out-of-band backup. This has impact on snapshot pool storage.
- **Crosscheck Archivelog:** Select this to run crosscheck and delete expired archivelogs on archive backup. For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#)
- **Crosscheck Backup of Archivelog Once a Day:** Select this to crosscheck the backup of archivelog once a day during log backup after database backup. This option will override crosscheck during each archivelog backup if both are selected.
- **Crosscheck Backup of Archivelog:** Select this to run crosscheck on the current backed up archivelog before the new logs are backed up, and delete expired archivelogs.
For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#)

- **Number of Files per Backupset:** Specify the number of archivelogs to include in a backupset during archivelog backup. For additional information, see [Chapter 23, Best Practices for Application Details & Settings](#).
- **Oracle Configuration File Location:** Use this when backing up Oracle configuration files with an Oracle OOB backup such as wallet for encryption support. Requires a full path name. If a folder name is specified, all files under that folder are backed up. If a file name is specified then only the specified file is backed up. Keys are not backed up with the database backup.

Note: For Oracle databases with TDE, the wallet for TDE can be captured by setting the Oracle Configuration File location advanced setting for the Oracle application. Application aware mounts for TDE enabled databases require the wallet to be copied to the appropriate location on the mount host and the wallet must be configured and open.

- **Oracle TNS_Admin Path:** If tnsnames.ora is in a nonstandard location, then provide the full path of the directory where it is located. The Oracle TNS_Admin Path is used only with database authentication.
- **Archivelog Backup Servicename:** Provide a dedicated Oracle database service name for the archive log backup in RAC environment when Oracle service name is set to run from more than one node. The Archivelog Backup Servicename is used only with database authentication.
- **Convert ASM Format to Filesystem Format:** By default, the database is captured in its native format, either ASM or file system. The backup destination is ASM to ASM and non-ASM to file system. Set this to Yes if the source database is ASM and backup destination must be set to file system.
- **Convert Filesystem Format to Oracle ASM Format:** By default, the database is captured in its native format, either ASM or file system. The backup destination is ASM to ASM and non-ASM to file system. Set this to Yes if the source database is under file system and backup destination must be set to Oracle ASM. This requires ASM to be installed on the Oracle server.
- **Auto Discover RAC Members:** Check this to autodiscover all members of the RAC databases in an ASM disk group out-of-band configuration. This enables mapping the staging disk to all nodes. Auto-discovery will not work if the hostname does not have a FQDN. In that case add the nodes manually.
- **RAC Member Nodes:** If you choose not to autodiscover RAC members, then provide a node list for mapping the staging disk as a shared volume for backup. List the protected nodes first. Use this only for protecting Oracle databases in an ASM disk group.
- **Prefer LVM for Single Staging Disk:** Select this to create an LVM even when a single staging disk is enough for backup.

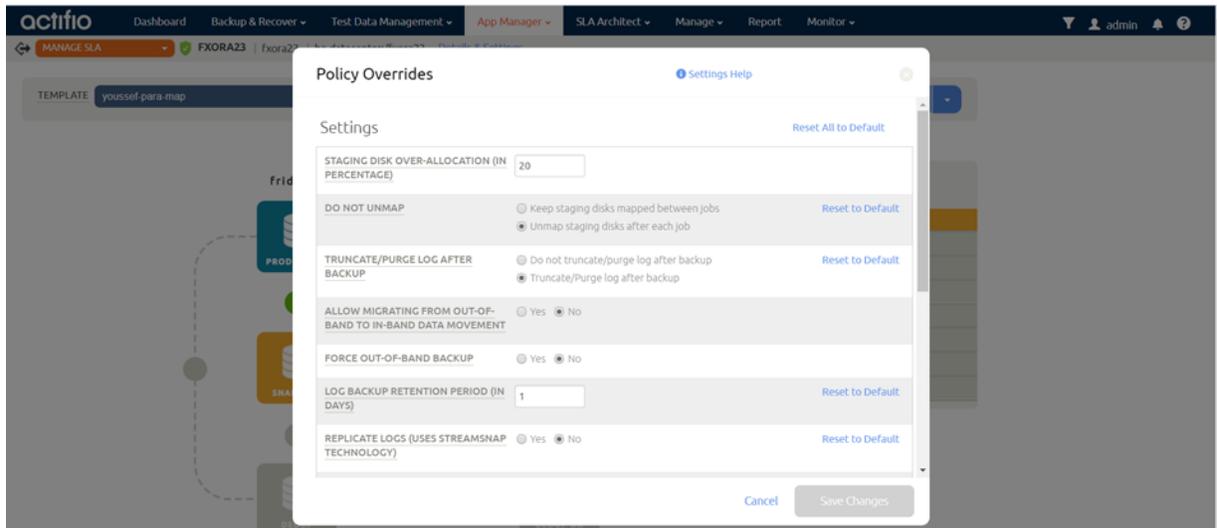
Policy Overrides for Oracle Databases

SLA policy overrides allow you to customize an SLA policy for a specific application. To set policy overrides:

1. Go to the App Manager > Applications and right-click a database. Select **Manage SLA**.
2. At the top of the page, click **Policy Overrides**.



This opens the Policy Overrides page, where you can enter the policy override details below.



Staging Disk Over-Allocation: This parameter determines the extra space allocated for staging disk to accommodate growth of the application. The range is 0 to 1000%. For more information, see [Chapter 23, Application Details & Settings Recommended Settings](#).

Do Not Unmap: Select whether to keep staging disks mapped between jobs.

Truncate/Purge Log After Backup: To manage log purging, select this. The default is Do Not Truncate. If a policy with Enable Database Log Backup is set to No, and if Truncate Logs After Backup is Yes, then archive log purging runs at the end of each database backup, purging all the logs.

Allow Migrating from Out-of-Band to In-Band Data Movement: Backup an in-band application using in-band mode even when there are existing out-of-band backups.

Force Out-of-Band Backup: Select this to enable the RMAN archive backup to run in compress mode.

Log Backup Retention Period: The archive log backup under Actifio staging disk will be retained to the value set here. Backup log retention can be different from snapshot retention.

Replicate Logs: Select if you want the database logs to be replicated, enabling for point-in-time recoveries at the remote location. Requires the database to also be replicated.

Log Staging Disk Growth Size: Set a percentage by which to grow the staging disk when needed.

Estimated Change Rate: Estimate the percentage by which the database data changes daily.

Compress Database Log Backup: Use this to enable RMAN archive backup to run in compress mode.

Script <> Timeout: Timeout values for each script type: Init, Freeze, Unfreeze, Finish, Post Replication.

Table 1: Archive Log Purge Behavior if SLA Advanced Policy Settings are Overridden by Application Policy

TEMPLATE SLA Advanced Policy Settings	APPLICATION Policy Overrides IF Allow Override is set to Yes	Behavior
Truncate/Purge log	Truncate/Purge log	Purge log will run at the end of each archive backup with a retention of 24 hours (delete archivelog older than sysdate -1)
Do not Truncate/Purge	Truncate/Purge log	Purge log will run at the end of each archive backup with a retention of 24 hours (delete archivelog older than sysdate -1)
Truncate/Purge log	Do not Truncate/Purge	Archive log will not be purged
Do not Truncate/Purge	Do not Truncate/Purge	Archive log will not be purged

8 Configuring dNFS for Protecting and Mounting Virtual Oracle Databases

This chapter includes:

[Before You Begin](#) on page 41

[Configuring AGM for Protecting and Mounting Virtual Oracle Databases over dNFS](#) on page 42

[Actions to be Performed on the Host for dNFS to Work](#) on page 43

[Troubleshooting dNFS: Database Issues](#) on page 44

Actifio can present staging disks to the source or target hosts of an Oracle database via NFS. Historically NFS has been too slow for use with databases.

Starting with Oracle 11.2.0.4, Oracle has built in the ability to use Direct NFS (dNFS) which makes NFS a viable option. dNFS bypasses several O/S layers of the NFS stack (the root cause for most bottlenecks) and allows each process to establish direct IO communications to the storage provider.

Before You Begin

In order to use dNFS with an Actifio appliance, the following requirements must be met:

- Actifio VDP 10.0.2 or higher
- ARC Cache increased to minimum of 8GB on Actifio Appliance
- Sufficient network bandwidth between database server and Actifio Appliance (minimum 10Gb, recommended 25Gb)
- Use all Oracle required/recommended patches. Oracle maintains a list of required/recommended patches in the Oracle Support documentation.

Configuring AGM for Protecting and Mounting Virtual Oracle Databases over dNFS

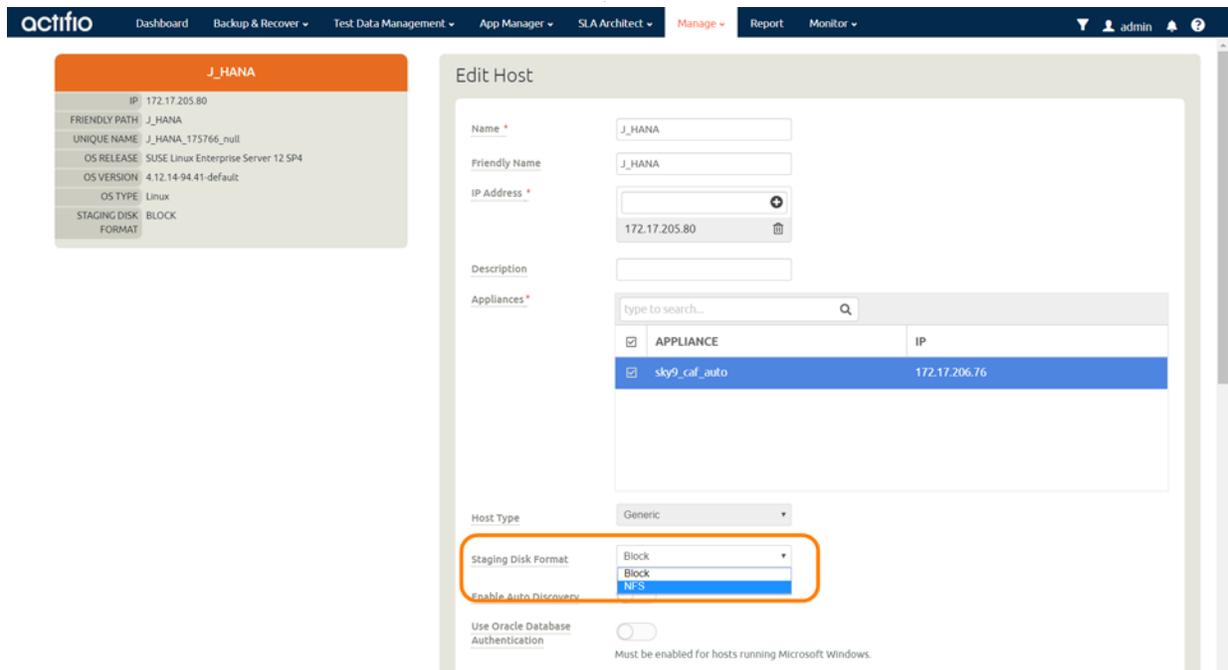
To perform Oracle Direct NFS (dNFS) based backup, you must set the Actifio Appliance staging disk format (disk preference) to NFS. You can set the staging disk format on the database host in two ways:

- [Setting the Staging Disk Format from AGM on page 42](#)
- [Setting the Staging Disk Format from the CLI on page 42](#)

Setting the Staging Disk Format from AGM

To set the Staging Disk Format (disk preference) to NFS:

1. Go to **Manage, Hosts**.
2. Right-click the host and select **Edit**.
3. At Staging Disk Format, select **NFS** and then click **Save**.



Setting the Staging Disk Format from the CLI

To set the Staging Disk Format (disk preference) to "NFS" from CLI for the specific backup host, use `udstask chhost`:

```
udstask chhost -diskpref NFS <hostid>
```

Example:

```
[07:05:21] localhost:~ # udstask chhost -diskpref NFS 404373
404373
[07:05:44] localhost:~ # udsinfo lshost 404373 |grep -i disk
diskpref NFS
[07:06:10] localhost:~ #
```

Actions to be Performed on the Host for dNFS to Work

Perform these actions to be sure that dNFS is configured correctly:

1. Check for this message under DB Alert.log to confirm that dNFS is enabled:
Oracle instance running with ODM: Oracle Direct NFS ODM Library Version 3.01.
2. If dNFS is not enabled, then enable it:
 - NFS Client packages must exist on the database host for protection jobs, and on any Oracle host on which you might mount a captured Oracle database via dNFS.
For example, for Linux, the nfs-util package should exist on the host. To check:

```
rpm -qa |grep nfs-util
```
 - Enable dNFS on the Oracle host:

```
cd $ORACLE_HOME/rdbms/lib  
make -f ins_rdbms.mk dnfs_on
```
 - Restart the databases running on that ORACLE_HOME, then check for this message under DB Alert.log to confirm that dNFS is enabled:
Oracle instance running with ODM: Oracle Direct NFS ODM Library Version 3.0
3. Trigger the Oracle-protected DB backup job. During the backup job, run this query to check dNFS usage:

```
select * from gv$dnfs_servers;
```

You can see the NFS READ/WRITE stats for the happening i/o:

```
select inst_id, PNUM, NFS_READ, NFS_WRITE, NFS_COMMIT, NFS_MOUNT from gv$dnfs_stats  
where NFS_READ>0 or NFS_WRITE>0 order by inst_id, PNUM;
```

We can see the dnfs channel process information.

```
select c.inst_id, program, pid,pname, local, path from gv$process p,  
gv$dnfs_channels c where p.inst_id = c.inst_id and c.pnum = p.pid;
```
4. Test dNFS usage with a new application aware mount to the same host and to a different host.

For a Virtual Database Mount

These are the operations to be performed for the Application Aware mount.

There is no need to set the staging disk format. When you select the NFS/dNFS based backup image and perform an application aware mount, Backup and DR automatically uses NFS export irrespective of the "Staging Disk Format" (disk preference) set for the host.

1. On the mounted target host, check for this message in DB Alert.log to confirm dNFS is enabled:
Oracle instance running with ODM: Oracle Direct NFS ODM Library Version 3.01.
If dNFS is not enabled, then enable it:
 - NFS Client packages must exist on the database host for protection jobs, and on any Oracle host on which you might mount a captured Oracle database via dNFS.
For example, for Linux, the nfs-util package should exist on the host. To check:

```
rpm -qa |grep nfs-util
```
 - Enable dNFS on the Oracle host:

```
cd $ORACLE_HOME/rdbms/lib  
make -f ins_rdbms.mk dnfs_on
```
 - Restart the databases running on that ORACLE_HOME, then check for this message under DB Alert.log to confirm that dNFS is enabled:
Oracle instance running with ODM: Oracle Direct NFS ODM Library Version 3.0
2. Perform the Application Aware mount by selecting the backup image. It will use the dNFS on the target host as we already enabled.

Troubleshooting dNFS: Database Issues

This includes:

- [Alert Log](#) on page 44
- [Database Trace Files](#) on page 44
- [Database Hang](#) on page 44
- [dNFS Views](#) on page 44
- [The Oracle dNFS Monitor Package](#) on page 48

Alert Log

The first stop for any debug operation is to check the alert log for dNFS related messages. A common issue observed on databases with dNFS is with the socket buffer size being limited. Oracle will try to adjust the size, but this can be limited by the O/S. In this case, an error like this one will be found in the alert log:

```
Direct NFS: Failed to set socket buffer size.wtmax=[1048576] rtmax=[1048576], errno=-1
```

Other items to look for in the alert log include if the correct network cards are being used to communicate with the filer. This can be determined by looking for a message similar to the following:

```
Direct NFS: channel id [0] path [192.168.56.3] to filer [192.168.56.3] via local [] is UP
```

Database Trace Files

If I/O issues are occurring, the following events can be set in the database to capture additional logging information. Set these events, wait for the incident to occur, then review to trace files.

```
ALTER SYSTEM SET MAX_DUMP_FILE_SIZE =UNLIMITED;
ALTER SYSTEM SET EVENTS '10298 trace name context forever, level 1'; # KSFD I/O tracing
ALTER SYSTEM SET EVENTS '19392 trace name context forever, level 8'; # kgnfs tracing
ALTER SYSTEM SET EVENTS '19394 trace name context forever, level 8'; # skgnfs tracing
ALTER SYSTEM SET EVENTS '19396 trace name context forever, level 6'; # kgodm tracing
ALTER SYSTEM SET EVENTS '19398 trace name context forever, level 128'; # mount tracing errors
```

Database Hang

If a database running on dNFS is hanging, then login as SYSDBA via sqlplus and perform a hang analysis/dump:

```
SQL> oradebug setmypid
SQL> oradebug unlimited
SQL> oradebug hanganalyze 3
SQL> oradebug dump systemstate 266
```

If database is a RAC database, then add a -g option to the last two oradebug commands.

dNFS Views

The dNFS client is actually in the database kernel. Therefore, several v\$ views exist within the database to monitor and check the health of dNFS from within the database. Oracle provides a package that can be used to quickly monitor dNFS performance. This package is in [The Oracle dNFS Monitor Package](#) on page 48.

Once deployed, a DBA can perform the following to get information (parameters: dnfs_monitor(<sleep time>), dnfs_itermonitor(<sleep time>,<nbr of times to check>), sleep time is in seconds):

```
SQL> set serveroutput on
SQL> set lines 200
SQL> exec dnfs_monitor(60);
Started at 01/18/2017 10:09:46 AM
Finished at 01/18/2017 10:10:46 AM
READ IOPS:          2
WRITE IOPS:         3
```

```

TOTAL IOPS:          5
READ Throughput:    0 MB/s
WRITE Throughput:   0 MB/s
TOTAL Throughput:   0 MB/s
SQL> exec dnfs_itermonitor(2,10)
Started at 01/18/2017 10:20:18 AM
TIMESTAMP          READ IOPS  WRITE IOPS  TOTAL IOPS  READ (MB/s)  WRITE (MB/s)  TOTAL (MB/s)
01/18/2017 10:20:20 AM  15         7           22          0            0             0
01/18/2017 10:20:22 AM   2         3            5          0            0             0
01/18/2017 10:20:24 AM   0         3            3          0            0             0
01/18/2017 10:20:26 AM   2         2            4          0            0             0
01/18/2017 10:20:28 AM   0         3            3          0            0             0
01/18/2017 10:20:30 AM   2         3            5          0            0             0
01/18/2017 10:20:32 AM   4         3            7          0            0             0
01/18/2017 10:20:34 AM   0         3            3          0            0             0
01/18/2017 10:20:36 AM   2         3            5          0            0             0
01/18/2017 10:20:38 AM   2         3            5          0            0             0
Finished at 01/18/2017 10:20:38 AM

```

The V\$ Views are:

- **V\$DNFS_SERVER**: Shows information for all NFS server connections (one for each NFS server). View is useful to verify connectivity and TCP socket settings.
- **V\$DNFS_CHANNELS**: Shows information for all network paths created to the NFS servers. Each dNFS client creates one channel per process per network path. If multiple paths exists (multiple NICs), the dNFS client load balances over all channels. Data reflects activity since last select.
- **V\$DNFS_FILES**: Shows files that are currently open via dNFS client.
- **V\$DNFS_STAT**: Performance metrics for dNFS client.

Table 1: V\$DNFS_SERVER

Column	Description
SRVNAME	NFS Server Name
DIRNAME	Volume exported by NFS Server
MNTPORT	Local Mount Port
NFSPORT	NFS Server Port
WTMAX	Max write size for NFS Server
RTMAX	Max read size for NFS Server

Table 2: V\$DNFS_CHANNELS

Column	Description
PNUM	Oracle Process Number (link to PID in v\$process)
SVRNAME	NFS Server Name
PATH	Network path to server

Table 2: V\$DNFS_CHANNELS

Column	Description
CH_ID	dNFS Channel ID
SVR_ID	dNFS Server ID
SENDS	Send operations over channel since last select.
RECVS	Receive operations over channel since last select.
PINGS	Ping operations over channel since last select.

Table 3: V\$DNFS_FILES

Column	Description
FILENAME	Name of file.
FILESIZE	Size of file.
PNUM	Process ID (link to PID in v\$process)
SRV_ID	NFS Server ID

Table 4: V\$DNFS_STAT

Column	Description
PNUM	Oracle Process Number (link to PID in v\$process)
NFS_NULL	Null operations
NFS_GETATTR	Get attribute operations
NFS_SETATTR	Set attribute operations
NFS_LOOKUP	Lookup operations
NFS_ACCESS	Access operations
NFS_READLINK	Read link operations
NFS_READ	Read operations
NFS_WRITE	Write operations

Table 4: V\$DNFS_STAT

Column	Description
NFS_CREATE	Create operations
NFS_MKDIR	Make directory operations
NFS_MKNOD	Make node operations
NFS_SYMLINK	Symbolic link operations
NFS_REMOVE	Remove operations
NFS_RMDIR	Remove directory operations
NFS_RENAME	Rename operations
NFS_LINK	Link operations
NFS_READDIR	Read directory operations
NFS_READDIRPLUS	Read directory plus operations
NFS_FSSTAT	File system status operation
NFS_FSINFO	File system information operations
NFS_PATHCONF	Path configuration operations
NFS_COMMIT	Commit operations
NFS_MOUNT	Mount operations

The Oracle dNFS Monitor Package

```
CREATE OR REPLACE PROCEDURE dnfs_monitor
  (sleepSecs IN NUMBER)
IS
  startTime      DATE;
  startReadIOPS  NUMBER;
  startWriteIOPS NUMBER;
  startReadBytes NUMBER;
  startWriteBytes NUMBER;
  endTime        DATE;
  endReadIOPS    NUMBER;
  endWriteIOPS   NUMBER;
  endReadBytes   NUMBER;
  endWriteBytes  NUMBER;
  readThr        NUMBER;
  writeThr        NUMBER;
  readIOPS        NUMBER;
  writeIOPS       NUMBER;
  elapsedTime    NUMBER;
BEGIN

  SELECT sysdate, SUM(stats.nfs_readbytes), SUM(stats.nfs_writebytes), SUM(stats.nfs_read),
SUM(stats.nfs_write)
  INTO startTime, startReadBytes, startWriteBytes, startReadIOPS, startWriteIOPS
  FROM dual, v$dnfs_stats stats;

  DBMS_OUTPUT.PUT_LINE('Started at ' || TO_CHAR(startTime, 'MM/DD/YYYY HH:MI:SS AM'));

  DBMS_LOCK.SLEEP(sleepSecs);

  SELECT sysdate, SUM(stats.nfs_readbytes), SUM(stats.nfs_writebytes), SUM(stats.nfs_read),
SUM(stats.nfs_write)
  INTO endTime, endReadBytes, endWriteBytes, endReadIOPS, endWriteIOPS
  FROM dual, v$dnfs_stats stats;

  DBMS_OUTPUT.PUT_LINE('Finished at ' || to_char(endTime, 'MM/DD/YYYY HH:MI:SS AM'));

  elapsedTime := (endTime - startTime) * 86400;
  readThr := (endReadBytes - startReadBytes)/(1024 * 1024 * elapsedTime);
  writeThr := (endWriteBytes - startWriteBytes)/(1024 * 1024 * elapsedTime);
  readIOPS := (endReadIOPS - startReadIOPS)/elapsedTime;
  writeIOPS := (endWriteIOPS - startWriteIOPS)/elapsedTime;

  DBMS_OUTPUT.PUT_LINE('READ IOPS:          ' || LPAD(TO_CHAR(readIOPS, '999999999'), 10, ' '));
  DBMS_OUTPUT.PUT_LINE('WRITE IOPS:         ' || LPAD(TO_CHAR(writeIOPS, '999999999'), 10, ' '));
  DBMS_OUTPUT.PUT_LINE('TOTAL IOPS:         ' || LPAD(TO_CHAR(readIOPS + writeIOPS, '999999999'),
10, ' '));
  DBMS_OUTPUT.PUT_LINE('READ Throughput:   ' || LPAD(TO_CHAR(readThr, '999999999'), 10, ' ') || '
MB/s');
  DBMS_OUTPUT.PUT_LINE('WRITE Throughput:  ' || LPAD(TO_CHAR(writeThr, '999999999'), 10, ' ') || '
' MB/s');
  DBMS_OUTPUT.PUT_LINE('TOTAL Throughput:  ' || LPAD(TO_CHAR(readThr + writeThr, '999999999'),
10, ' ') || ' MB/s');
END;
/

CREATE OR REPLACE PROCEDURE dnfs_itermonitor
  (sleepSecs IN NUMBER,
  iter      IN NUMBER)
IS
  startTime      DATE;
  startReadIOPS  NUMBER;
  startWriteIOPS NUMBER;
  startReadBytes NUMBER;
  startWriteBytes NUMBER;
  endTime        DATE;
```

```

endReadIOPS    NUMBER;
endWriteIOPS   NUMBER;
endReadBytes   NUMBER;
endWriteBytes  NUMBER;
readThr        NUMBER;
writeThr       NUMBER;
readIOPS       NUMBER;
writeIOPS      NUMBER;
i              NUMBER;
elapsedTime    NUMBER;
BEGIN

    DBMS_OUTPUT.PUT_LINE('Started at ' || TO_CHAR(SYSDATE, 'MM/DD/YYYY HH:MI:SS AM'));

    DBMS_OUTPUT.PUT_LINE(
        LPAD('TIMESTAMP', 15, ' ')||
        LPAD('READ IOPS', 33, ' ')||
        LPAD('WRITE IOPS', 15, ' ')||
        LPAD('TOTAL IOPS', 15, ' ')||
        LPAD('READ (MB/s)', 15, ' ')||
        LPAD('WRITE (MB/s)', 15, ' ')||
        LPAD('TOTAL (MB/s)', 15, ' '));

    FOR i IN 1..iter
    LOOP
        SELECT sysdate, SUM(stats.nfs_readbytes), SUM(stats.nfs_writebytes), SUM(stats.nfs_read),
        SUM(stats.nfs_write)
        INTO startTime, startReadBytes, startWriteBytes, startReadIOPS, startWriteIOPS
        FROM dual, v$dtrfs_stats stats;

        DBMS_LOCK.SLEEP(sleepSecs);

        SELECT sysdate, SUM(stats.nfs_readbytes), SUM(stats.nfs_writebytes), SUM(stats.nfs_read),
        SUM(stats.nfs_write)
        INTO endTime, endReadBytes, endWriteBytes, endReadIOPS, endWriteIOPS
        FROM dual, v$dtrfs_stats stats;

        elapsedTime := (endTime - startTime) * 86400;
        readThr := (endReadBytes-startReadBytes)/(1024 * 1024 * elapsedTime);
        writeThr := (endWriteBytes-startWriteBytes)/(1024 * 1024 * elapsedTime);
        readIOPS := (endReadIOPS - startReadIOPS)/elapsedTime;
        writeIOPS := (endWriteIOPS - startWriteIOPS)/elapsedTime;

        DBMS_OUTPUT.PUT_LINE(
            TO_CHAR(endTime, 'MM/DD/YYYY HH:MI:SS AM') ||
            LPAD(TO_CHAR(readIOPS, '999999999'), 15, ' ') ||
            LPAD(TO_CHAR(writeIOPS, '999999999'), 15, ' ') ||
            LPAD(TO_CHAR(readIOPS + writeIOPS, '999999999'), 15, ' ') ||
            LPAD(TO_CHAR(readThr, '999999999'), 15, ' ') ||
            LPAD(TO_CHAR(writeThr, '999999999'), 15, ' ') ||
            LPAD(TO_CHAR(readThr + writeThr, '999999999'), 15, ' '));
    END LOOP;
    DBMS_OUTPUT.PUT_LINE('Finished at ' || to_char(endTime, 'MM/DD/YYYY HH:MI:SS AM'));

END;
```

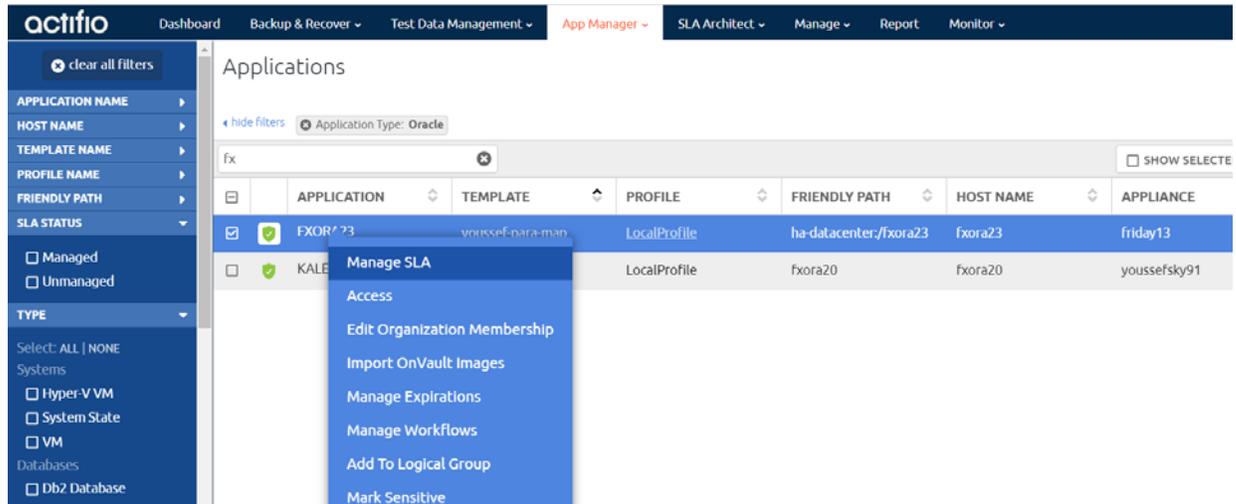

9 Virtualizing an Oracle Database for Data Protection and Agility

Virtualizing an Oracle database allows you to maintain up-to-date copies of it and to mount them for different business resiliency and agility purposes such as data protection and test/dev work. Before you can virtualize and protect Oracle databases, see:

- Chapter 3, Actifio Prerequisites for Protecting an Oracle Database
- Chapter 4, Oracle Authentication
- Chapter 5, Data Capture under File System and under ASM Disk Group
- Chapter 6, Preparing Oracle Databases for Protection.

To capture an Oracle database and its logs:

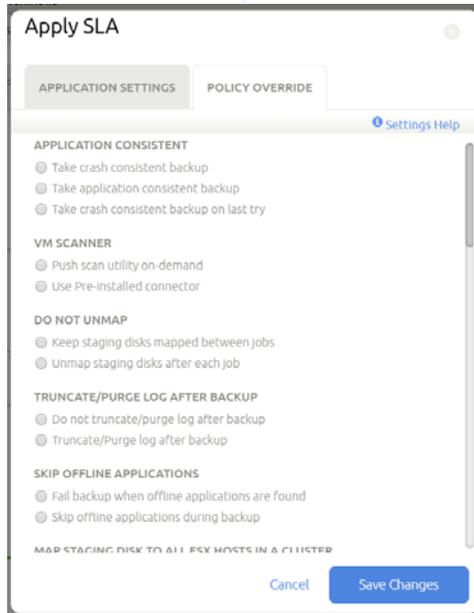
1. Open the AGM to the **App Manager > Applications** and enter the database application name or use the filters to make it easier to get to the database that you need.
2. Right-click the application and select **Manage SLA**.



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



4. The Apply SLA dialog opens. Assign policy overrides and change application settings as needed.



Applying SLA Policy Overrides

5. You can wait for the job to run during the period scheduled in the SLA, or you can run the job at the next opportunity by clicking on the desired job and selecting **Run SLA**.

Note: To take Oracle ASM to FS backups over NFS, you must enable **Convert ASM to FS format** in the database Application Details & Settings.



Running the SLA (Optional: If You Do Not Want to Wait for the Scheduled Time)

6. The job runs as soon as the scheduler has an opening, often immediately. You can go to the Jobs Monitor to view the progress and details of the job.

Note: If the template will capture logs, and if you have software that purges logs through RMAN, be sure to disable it. If that purge runs during an Actifio backup job, the backup may have incomplete log information.

10 Accessing, Recovering, or Restoring an Oracle Database

Actifio offers several ways to access data, including mounting and restoring. The most common ways to access an Oracle database:

The **standard mount** 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. Standard mount methods include:

- o [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access on page 54](#)
- o [Mounting an Oracle Database Image Protected Under a File System for Data Access on page 56](#)

The **Application Aware mount** presents and makes the captured Oracle database available to a target server as a virtual Oracle database. This allows you to address the unique challenges associated with creating and managing copies of production databases for non-production use. Application Aware mounts are performed from the Actifio Appliance and do not require manual intervention by database, server, or storage administrators. Application Aware mounts can be used for such things as database reporting, analytics, integrity testing, and test and development. Application Aware mounts are described in [Mounting an Oracle Database as a Virtual Application on page 58](#).

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. To restore a database, see [Restoring a Database, Overwriting the Production Database on page 63](#).

You can also clone and LiveClone Oracle databases following the general procedures in the AGM online help; there are no Oracle-specific procedures for those methods of data access.

Oracle-specific workflows in the AGM are introduced in [Chapter 19, Introduction to Provisioning Environments With Workflows](#) and detailed in the chapters that follow.

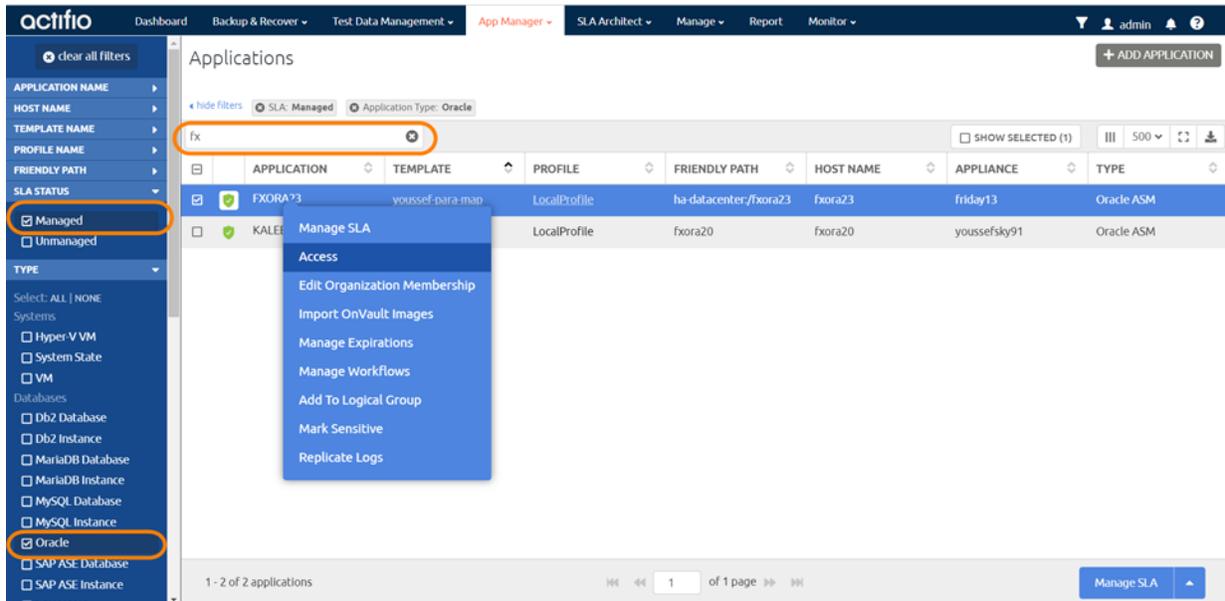
After any database server reboot where an Actifio image is mounted, or if Actifio backups are in progress for the database at the time of reboot/crash, please see [Bringing Actifio-Protected ASM Diskgroups Back Online after Reboot of a Target DB Server on page 63](#).

Note: With Oracle full database restore, the database incarnation will change and the log sequence will get reset. This requires a full level 0 backup as the previous backups become obsolete for backing up the database with the new incarnation. Actifio backup keeps track of incarnation and will check the incarnation for any change before each database backup job. If it detects an incarnation change it will automatically trigger a FULL LEVEL 0 BACKUP. This will consume the additional space (based on the size of the database) for full backup in snapshot, and the job takes longer to complete.

Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access

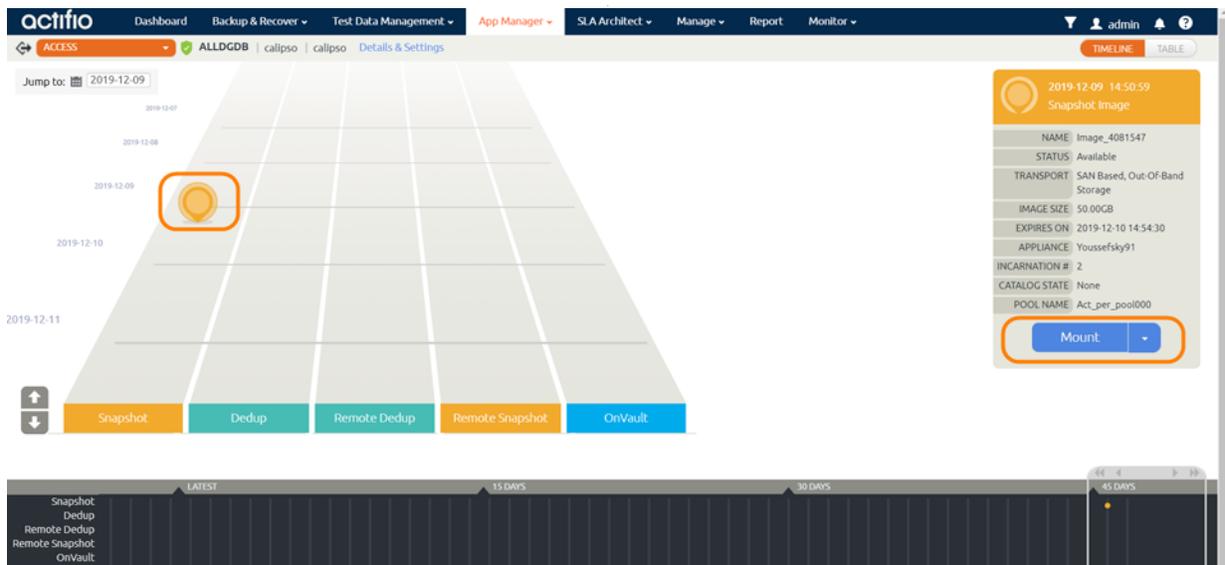
To mount an Oracle database image for data access:

1. Open the AGM to the **App Manager > Applications** and enter the database application name or use the filters to make it easier to get to the database image that you need.
2. Right-click the application and select **Access**.



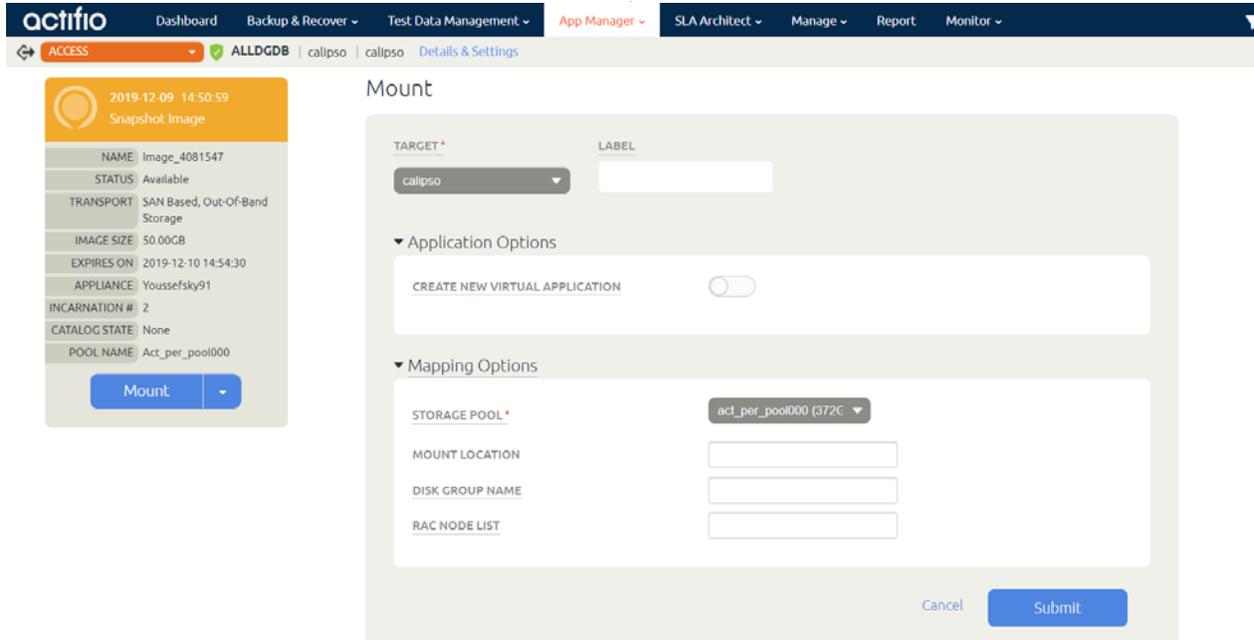
Selecting an Oracle Database

3. On the Access page, select the desired image and click **Mount** under the Mount menu.



Selecting a Managed Oracle Database Image

4. On the Mount page, fill in the required information.



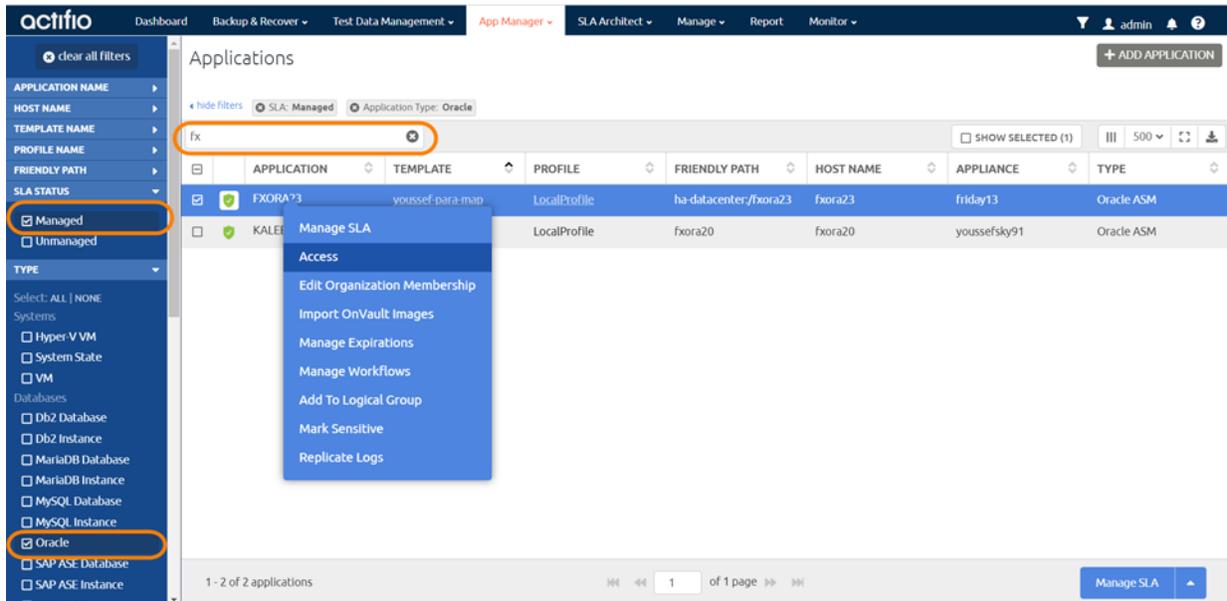
Mount Details for an Oracle Database

- o Under **Target**, select a host to mount the new database on, and provide a **Label** as needed. This is optional.
 - o Under **Application Options**, deselect **Create New Virtual Application**. If you want to create an application aware mount, then see [Mounting an Oracle Database as a Virtual Application](#) on page 58.
 - o Open the **Mapping Options** by clicking on the arrow icon to the left of the title.
 - o If necessary, change the default storage pool from the **Storage Pool** drop-down list. The available free space in the pool is indicated in parentheses.
 - o **Mount Location:** Enter the drive letter or the full path where to mount the new database: If the path exists as an empty folder, the Actifio Connector will use it. If it does not exist, the Connector will create it. If it exists as a file or as a folder that is not empty, then the job will fail. If there are multiple volumes to be mounted, the Connector uses the mount location that you specify for one of the volumes, and for the remaining volumes it appends an underscore followed by a number, i.e., <specified>_#
 - o Enter a diskgroup name for the mounted image copy at **Disk Group Name**.
 - o Enter the **RAC Node List:** To provision a RAC database on the target RAC cluster, specify the IP address of all nodes for the target RAC cluster separated by a colon (:) in the order of RAC nodes 1...n. The first IP address in RAC Node list *must* be the selected host's IP address. To provision a single node RAC database on a target RAC cluster or a standalone database under ASM on a non-RAC ASM target, provide the IP address of the target node.
 - o The Oracle **database image** will be mounted to ASM with a disk group name specified under Disk Group Name.
 - o If logs are Actifio-protected, then the **logs image** will be mounted to /act/mnt/<jobid>_Log, and subsequent logs images to /act/mnt/<jobid>_Log_1, /act/mnt/<jobid>_Log_2, and so on.
5. Click **Submit**. The job runs as soon as the scheduler has an opening, often immediately. You can go to the Job Monitor to view the progress and details of the job.

Mounting an Oracle Database Image Protected Under a File System for Data Access

To mount an Oracle database image for data access:

1. Open the AGM to the **App Manager > Applications** page and enter the database application name or use the filters to make it easier to get to the database image that you need.
2. Right-click the application and select **Access**.



Selecting an Oracle Database

3. On the Access page, select the desired image and click **Mount** under the Mount menu.



Selecting a Managed Oracle Database Image

4. On the Mount page, fill in the required information.
 - o Under **Target**, select the host for the newly mounted database. The requested information changes depending on whether you select a physical host, a VM as a vRDM, or a VM as a pRDM.

- o Provide a label as needed. This is optional.
- o Deselect **Create New Virtual Application**. If you want to make an application aware mount, then follow the procedure in [Mounting an Oracle Database as a Virtual Application](#) on page 58.

Table 1: Mapping Options

Host Type	Mapping Option
VM	Map to All ESX Hosts: Mapping staging disks to more than one ESX host in cluster ensures that VM can fail over to another ESX host in the event of an ESX host failure.
VM	Mount Mode: physical compatibility RDM (pRDM) or virtual compatibility RDMs (vRDM)
VM (vRDM)	Mark Dependent: A vRDM can be dependent or independent.
All	If necessary, change the default storage pool from the Storage Pool drop-down list.
All	Select a Mount Location : Enter the full path at which you want to mount the volume. <ul style="list-style-type: none"> • If the path exists as an empty folder, the Actifio Connector will use it. • If it does not exist, the Connector will create it. • If it exists as a file or as a folder that is not empty, then the job will fail.

If there are multiple volumes to be mounted, then the Connector uses:

Volume(s)	Mount location	No mount location
Database Image	/<mountpoint>	/act/mnt/<jobid>
Logs Image	/<mountpoint>_Log and subsequent logs images to <mountpoint>_Log_1, <mountpoint>_Log_2, and so on	/act/mnt/<jobid>_Log and subsequent logs images to /act/mnt/ <jobid>_Log_1, /act/mnt/<jobid>_Log_2, and so on

5. Check **Submit** to submit the job.

Mounting an Oracle Database as a Virtual Application

An Actifio Application Aware mount mounts a captured Oracle database as a virtual application. It allows you to quickly bring a database online without having to actually move the data and without having to manually configure a new instance of the database. An Application Aware mount addresses the challenges of creating and managing copies of production databases without manual intervention by database, server, and storage administrators.

Note: Oracle virtual applications consume a minimum 300 MB to account for redo logs and control file.

Note: Oracle virtual applications are not supported for databases installed with the virtual account.

Note: An SSH connection between RAC nodes is required for application aware mount to RAC nodes.

Note: The ASM Diskstring parameter must be set on any target server. See [The ASM diskstring Parameter Must Be Set](#) on page 25.

Patching Oracle 12c

Actifio Application Aware mounts may fail if your Oracle 12c installation does not include this patch, which can be downloaded from the Oracle support portal:

Oracle Database 12c Bug# 19404068 (ORA-1610 ON RECOVER DATABASE FOR CREATED CONTROLFILE)

- (Patch 19404068) Linux x86-64 for Oracle 12.1.0.2.0

To see if the patch is installed, run:

```
$cd $ORACLE_HOME/OPatch
$./opatch lsinventory -details
$./opatch lsinventory -details | grep 19404068
```

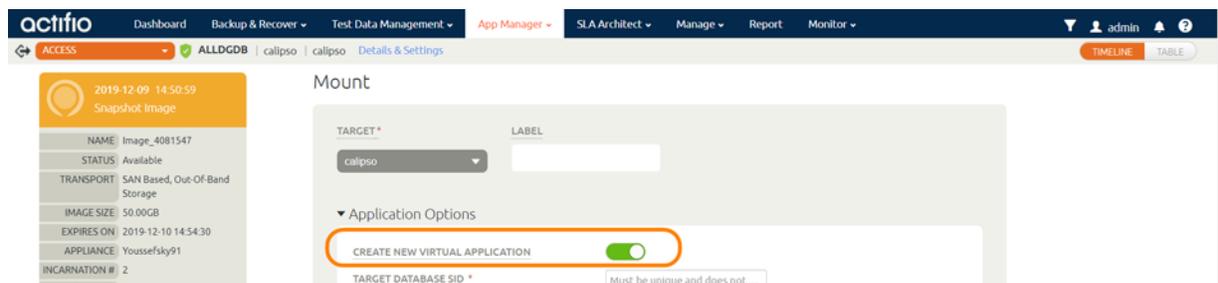
Pre-checks to Mount an Oracle Database as a Virtual Application

- Make sure that the database versions are matching between source host and target host:
 - Source: Where the source database is running
 - Target: Where you are trying to perform an Application Aware Mount
- Make sure that there are enough resources (like memory & CPU) on the target database server based on your performance requirements.

Mounting an Oracle Database as a Virtual Application

To mount an Oracle database as a virtual application:

1. Start detailed in [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56. Under Application Options, enable **Create New Virtual Application**.



Mounting and Confirming an Application Aware Oracle Database Mount

2. Fill in the form as needed for this virtual application. Fields marked with an asterisk (*) are required. You can click on each entry for additional helpful information.

Application Options

- o **Roll Forward Time:** Select the time of the image that you want to roll forward to.
- o **Target Database SID:** Specify the SID for the new Oracle database to be provisioned on the target. Follow standard Oracle naming conventions for this value. Make sure that the target host database version matches the version of the source host, and that there are enough resources on the target database server. Follow standard Oracle naming conventions for this value.
- o **User Name:** Specify Oracle Operating System user credentials on the target.
- o **Oracle Home Directory:** Specify the Oracle Home Directory (\$ORACLE_HOME) on the target database server.

Note: (Windows only) A user who runs an Application Aware mount must be the administrator user who owns the ORACLE_HOME, or else has READ/WRITE privileges on mount locations used during the App Aware mount. Otherwise a call to oradim will prompt for password and hang.

- o The application aware mount will be a new database. If you want to protect the new database, then enable **Manage New Application** to apply an SLA to the new database. New Template and Profile fields will appear where you can select any of your existing SLA templates and resource profiles.
The application aware mount will be a new database. You can have the new database protected by applying an SLA when you create the new database. The snapshots of the database are incremental unless you apply a policy template with Force Out-of-Band Backup checked.

Note: There is one exception: if the target server is a VMware VM, you must select "prdm" for the mount if you want the child database to have the efficient incremental snapshots. If you forget and leave the default of "vrDM", then the first snapshot job will be a full backup.

Advanced Options

- o **Password:** A password is required for Oracle 12c or above when ORACLE_HOME user is different from administrator, on Windows (only).
- o **TNS Admin Directory path:** Specify TNS_ADMIN Directory path (path of tnsnames.ora file) on the target database server.
- o **Database Memory Size in MB:** Database total memory size, in MB, for the database being provisioned on the target. See the table below for the expected behavior depending on how this and SGA% (next) are set.
- o **SGA%:** Parameter to configure SGA/PGA memory, as a percentage of total memory, for the database being provisioned. See the table below for the expected behavior depending on how this and Database Memory Size in MB (above) are set.

Database Memory Size in MB	SGA%	Behavior
not specified	N/A	Total database memory size and memory parameter on target will be the same as source database.
specified	not specified	MEMORY_TARGET parameter will be set for the database being provisioned on the target.

Database Memory Size in MB	SGA%	Behavior
specified	specified	Set SGA and PGA for the database provisioned on the target to: SGA_TARGET = Database Memory Size in MB x (SGA%/100) PGA_TARGET = Database Memory Size in MB x (100-SGA%/100) Note: Do not set SGA to 100. To avoid database slowness, be sure to reserve some memory space for PGA.

- o **REDO Size:** Parameter to configure REDO size, in MB, for the database being provisioned. If not specified, REDO size is be set to 1000 MB.
- o **Shared_Pool_Size in MB:** Parameter to configure shared pool size, in MB, for the database being provisioned. If not specified, shared_pool_size will not be used.
- o **DB_Cache_Size in MB:** Parameter to configure database cache size, in MB, for the database being provisioned. If not specified, db_cache_size will not be used.
- o **DB_Recovery_File_Dest_Size in MB:** Parameter to configure database recovery file destination size, in MB, for the database being provisioned. If not specified, db_recovery_file_dest_size will be set to 50000 MB.
- o **inmemory_size:** Parameter to configure database inmemory_size, in MB, for the database being provisioned. The minimum size can be set is 100MB. If not specified, inmemroy_size parameter will be skipped regardless of Oracle version.
- o **Diagnostic_Dest:** Parameter to configure diagnostic destination on the host. If not specified, diagnostic_dest will be set to ORACLE_HOME.

Note: Diagnostic_Dest is not supported for Oracle 10g. Simply leave it blank.

- o **Max number of processes:** Parameter to configure max number of system user processes that can simultaneously connect to Oracle, for the database being provisioned. If not specified, processes will be set to 500.
- o **Max number of open cursors:** Parameter to configure maximum number of open cursors that a session can have at once, for the database being provisioned. If not specified, number of open cursors will be using source database settings.
- o **TNS Listener IP:** Specify IP address for the TNS Listener. It can be one of SCAN IP, VIP, or Host IP. If not specified, Host IP will be used.
- o **TNS Listener port:** TNS Listener port to be used to create service name under tnsnames.ora for provisioned database on target. If not specified, port 1521 is used.
- o **TNS Domain Name:** Specify domain name to be used with service name under tnsnames.ora for provisioned database on target. This is needed when database service is using Domain Name.
- o **PDB Prefix:** Specify a prefix for renaming PDB during child database creation.
- o **User to be removed:** This is a comma separated list of users that to be removed as part of the mount operation.
- o **Do not change database DBID:** If selected, new database's DBID will not be changed.
- o **No Archive Mode:** If selected, new database will be running in no-archivelog mode. Reprotection of the new instance will not be available.
- o **Clear Archivelog:** If selected, clear archivelogs after masking is performed.

- o **Do not update tnsnames.ora:** If selected, an entry for the new database will not be added to tnsnames.ora. This may require manual intervention for connections to the new database, and in some cases snapshot jobs for the new database will fail without this manual intervention.
- o **Do not update oratab:** If selected, an entry for the new database will not be added to oratab if one exists.
- o **Add TNS Listener Entry:** Default is false. If this option is enabled, a listener entry will be added to the tnsnames.ora file.
- o **Number of Channels:** The number of RMAN channels.
- o **Clear OS_Authent_Prefix:** OS_Authent_Prefix is a prefix that Oracle uses to authenticate users connecting to the server. Oracle concatenates the value of this parameter to the beginning of the user's operating system account name and password.
- o **Restore with Recovery:** If selected, brings the newly created database online: the provisioned database on target will be open for read and write. This is the default selection.
- o **Stand Alone Non-RAC:** This is only applicable for databases where the source database is in a non-RAC configuration and Actifio stores the copy in ASM format. If selected, this performs an application aware mount to a standalone ASM non-RAC instance. Do not select this option if a RAC node list has been provided.
- o **Use existing Oracle password file:** If this option is selected, a virtual database mount uses the existing Oracle password file and does not delete it during cleanup on unmount
- o **Environment variable:** If you have any user-defined environment variables to be passed to pre / post scripts, you can enter one here.

▼ Application Options

CREATE NEW VIRTUAL APPLICATION

ROLL FORWARD TIME HOST TIME USER TIME

TARGET DATABASE SID *

USER NAME *

ORACLE HOME DIRECTORY *

▼ MANAGE NEW APPLICATION

TEMPLATE *

PROFILE *

▼ Advanced Options

PASSWORD

TNS ADMIN DIRECTORY PATH

DATABASE MEMORY SIZE IN MB

SGA %

REDO SIZE

SHARED_POOL_SIZE IN MB

DB_CACHE_SIZE IN MB

DB_RECOVERY_FILE_DEST_SIZE IN MB

DIAGNOSTIC_DEST

MAX NUMBER OF PROCESSES

MAX NUMBER OF OPEN CURSORS

TNS LISTENER IP

TNS LISTENER PORT

Application Aware Oracle Database Mount: Application Options and Advanced Options

3. Click **Submit** to submit the job.

Bringing Actifio-Protected ASM Diskgroups Back Online after Reboot of a Target DB Server

After any database server reboot where Actifio copy is mounted, or Actifio backups are in progress for the database at the time of reboot/crash, please follow these steps to get the Actifio disk group mount back:

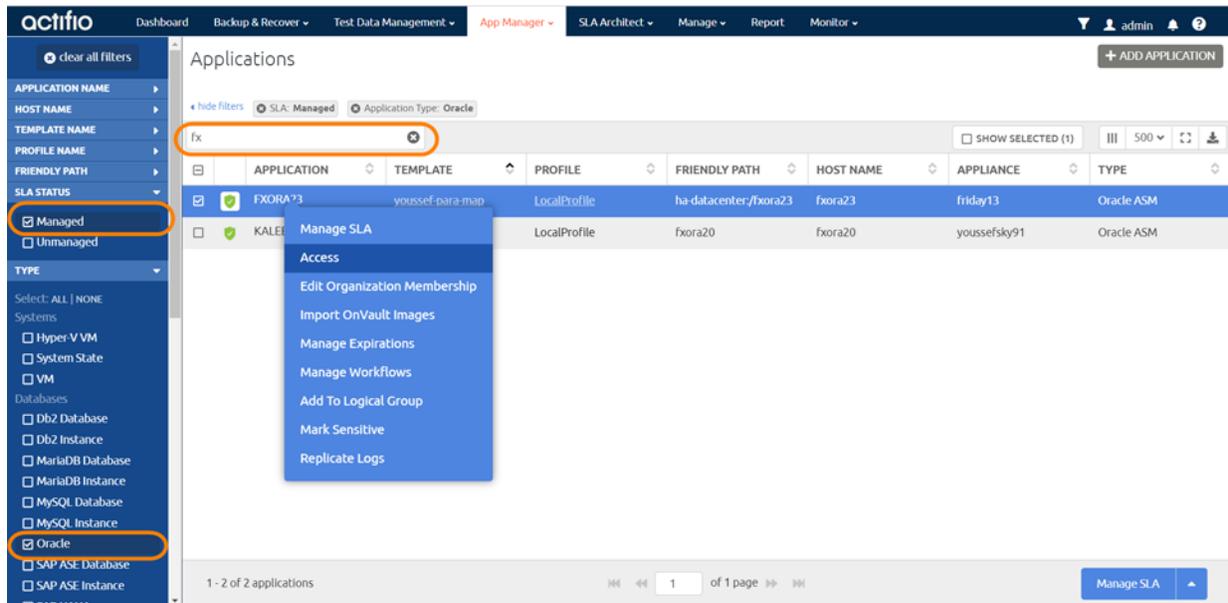
1. Check that the target database server is back up and that ASM and RAC system are also up.
2. Restart the Actifio Connector (from root).
3. Set ASM environment.
4. Login to ASM sqlplus and check the disk group status:

```
select name, state from v$asm_diskgroup where name = '<dg name>';
```
5. If unmounted, mount the disk group: `alter diskgroup <dg name> mount;`
6. Login to the Oracle OS and set the database environment, then start the database.

Restoring a Database, Overwriting the Production Database

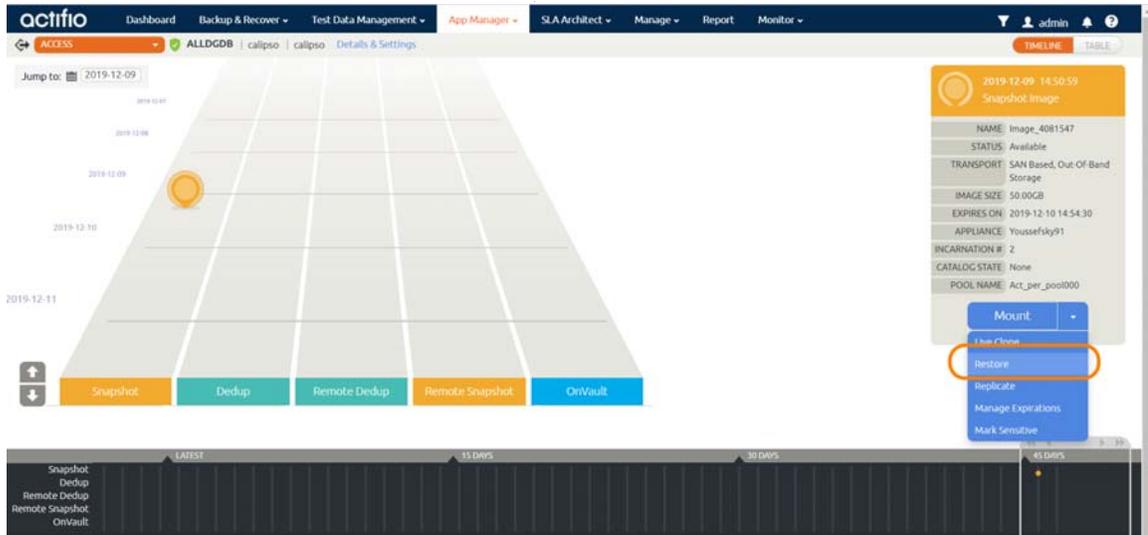
To restore an Oracle database out-of-band backup image, overwriting the original production database:

1. Open the AGM to the **App Manager** and enter the database application name or use the filters to make it easier to get to the database image that you need.
2. Right-click the application and select **Access**.



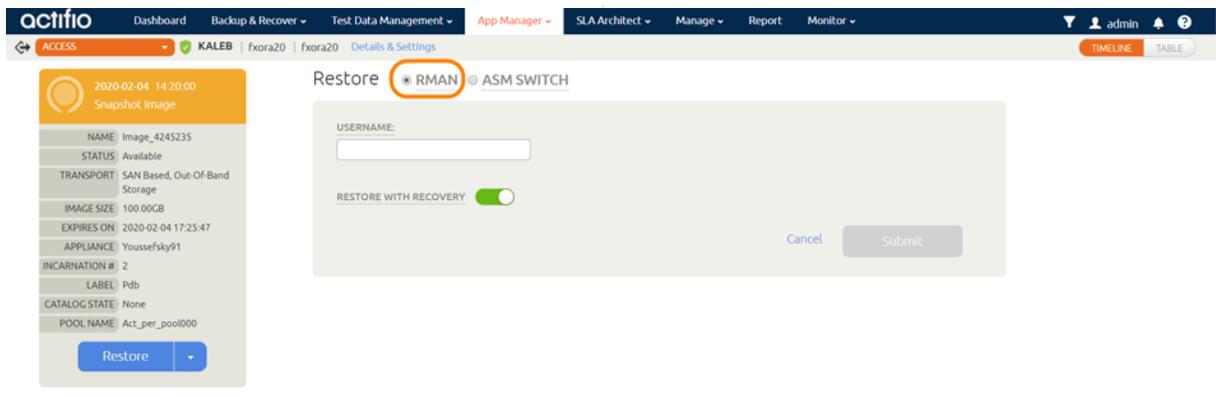
Selecting an Oracle Database to be Restored

3. On the Access page, select the desired image and click **Restore** under the Mount menu.



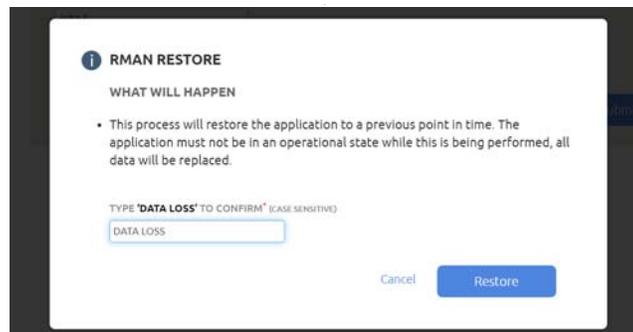
The Restore Option

- On the Restore page, select RMAN, then enter a username and click **Submit**.



Restoring an Oracle Database

- A warning dialog appears. Read it and enter **DATA LOSS** to confirm.



Confirmation

- The job is queued for the next available job slot. You can view progress from the Jobs Monitor.

Note: The SLA options (Run Schedule, Expire Data) of the database are disabled.

11 Recovering an Oracle Database Manually Using RMAN

The procedures to restore a database have subtle differences described below, but the basic procedure to recover to the point-in-time of the backup snapshot is:

1. Mount the latest database backup snapshot from Actifio back to the Oracle server.
2. Restore the parameter file and the control file.
3. Catalog the database backup snapshot to RMAN.
4. Restore and recover the database using an Actifio mounted backup.

The procedures vary depending upon whether the source database is RAC or non-RAC, whether the database is protected under a file system or under an ASM Disk Group, and whether the archivelog files are Actifio-protected or are not Actifio-protected.

Table 1: Six RMAN Procedures to Recover Databases

Source Database	Protected Under	Archivelog is	See
Non-RAC, Non-ASM	File System	Logs not Actifio-Protected	Recovering a Non-RAC Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio on page 66
RAC or Standalone ASM	File System	Logs not Actifio-Protected	Recovering a RAC ASM Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio on page 67
Non-RAC, Non-ASM	File System	Logs are Actifio-Protected	Recovering a Non-RAC Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog on page 69
RAC or Standalone ASM	File System	Logs are Actifio-Protected	Recovering a RAC ASM Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog on page 71
RAC or Standalone ASM	ASM Disk Group	Logs not Actifio-Protected	Recovering an Oracle Database to a Scheduled Backup Point if the archivelog is not Protected through Actifio on page 73
RAC or Standalone ASM	ASM Disk Group	Logs are Actifio-Protected	Recovering an Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog on page 75

Note: Actifio-Protected means database log protection is enabled in the SLA policy settings.

Note: With Oracle full database restore, the database incarnation will change and the log sequence will get reset. This requires a full level 0 backup as the previous backups become obsolete for backing up the database with the new incarnation. Actifio backup keeps track of incarnation and will check the incarnation for any change before each database backup job. If it detects an incarnation change it will automatically trigger a FULL LEVEL 0 BACKUP. This will consume the additional space (based on the size of the database) for full backup in snapshot, and the job takes longer to complete.

Recovering a Non-RAC Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio

Source Database	Protected Under	Archivelog is
Non-RAC, Non-ASM	File System	Not Actifio-Protected

To recover a non-RAC Oracle database for point-in-time recovery if the logs are not protected by Actifio:

1. Mount the image from the source database server to recover. In the Mount window under Mapping Options, provide a mount location for the image, for example: /acttestdb. For instructions on how to mount a database image, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.
2. Set the Oracle environment, and use sqlplus to shut down the database:


```
sqlplus / as sysdba
sql> shutdown immediate;
```

 Verify that the database is shut down. Kill any orphan process for the database.


```
ps -ef | grep <db name>
```
3. Create a new spfile from the existing pfile and restart the database:
 Start the database in nomount state using the parameter file from the mounted volume. The parameter file will be under a top mounted folder, for example: /acttestdb


```
sqlplus / as sysdba
sql> startup nomount pfile='/acttestdb/<database sid>__backup.ora';
```

 Create an spfile from the pfile:


```
sql> create spfile='$ORACLE_HOME/dbs/spfile<database sid>.ora' from pfile='/acttestdb/<database sid>__backup.ora';
```

 Restart the database in nomount state with the new spfile:


```
sql> shutdown immediate;
sql> startup nomount;
```
4. Use RMAN to restore the control file from the Actifio mounted volume:


```
rman target /
rman> restore controlfile from '/acttestdb/cf-D_<sid>-id_<id>.ctl' ;
```
5. Mount the database:


```
rman> alter database mount;
```
6. Catalog the datafile and archivelog folder from the Actifio mounted volume to RMAN:


```
rman> run { catalog start with '/acttestdb/datafile' noprompt;
catalog start with '/acttestdb/archivelog' noprompt; }
```
7. Restore and recover the database:


```
rman> run { restore database ; recover database ; }
```

Note: Ignore warning from RMAN looking for the next archivelog; this is a point-in-time recovery.

8. Open the database with the reset log option:


```
rman> alter database open resetlogs;
```

The database is available for read and write.

Recovering a RAC ASM Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through Actifio

Source Database	Protected Under	Archivelog is
RAC or Standalone ASM	File System	Not Actifio-Protected

To recover a standalone ASM or RAC Oracle database for point-in-time recovery if the logs are not Actifio-protected:

1. Mount the image from the source database server to recover. In the Mount window under Mapping Options, provide a mount location for the image, for example: /acttestdb. For instructions on how to mount a database image, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.
2. Shut down the Oracle database.
From node 1, su to the Oracle OS user:
su - oracle
Set the Oracle environment and use srvctl to stop the database across all nodes:
srvctl stop database -d <database name>
Verify that the database is shut down on all nodes. Kill any orphan process for the database.
ps -ef | grep <db name>
3. Start the database in nomount state using the parameter file from the mounted volume. The parameter file will be under a top mounted folder for example: /acttestdb
sqlplus / as sysdba
sql> startup nomount pfile='/acttestdb/<db name>__backup.ora';
4. Create a new spfile and restart the database:
To get the path of original spfile under disk group:
cat \$ORACLE_HOME/dbs/init<database sid>.ora
For example: spfile=+<preferred disk group>/<db name>/spfile<db name>.ora
sql> create spfile='+<preferred disk group>/<db name>/spfile<db name>.ora' from pfile='/acttestdb/<db name>__backup.ora';
Restart the database with spfile in nomount state:
sql> shutdown immediate;
sql> startup nomount;
5. Restore the control file using RMAN from the Actifio mounted volume.
rman target /
rman> restore controlfile from '/acttestdb/cf-D_<db name>-id_<id>.ctl' ;
6. Mount the database:
rman> alter database mount;
7. Catalog the datafile and archivelog folder from Actifio mounted volume to RMAN:
rman> run { catalog start with '/acttestdb/datafile' noprompt;
catalog start with '/acttestdb/archivelog' noprompt; }
8. Restore and recover the database:
rman> run { restore database ; recover database ; }

Note: Ignore any warning from RMAN looking for the next archivelog as this is a point-in-time recovery.

9. Open the database with the reset log option:

```
rman> alter database open resetlogs;
```

10. Shutdown the database on node 1 and start the database across all nodes.

Use sqlplus to shut down the database:

```
sqlplus / as sysdba  
sql> shutdown immediate;
```

Use srvctl to start database across all nodes:

```
srvctl start database -d <database name>
```

The database is available for read and write.

Recovering a Non-RAC Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog

Source Database	Protected Under	Archivelog is
Non-RAC, Non-ASM	File System	Actifio-Protected

To recover a non-RAC Oracle database for point-in-time recovery if the logs are Actifio-protected:

1. Mount the image from the source database server to recover. In the Mount window under Mapping Options, provide a mount location for the image, for example: /acttestdb. For instructions on how to mount a database image, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.

The database backup image will be mounted at:/acttestdb

The protected archive log will be mounted at: /acttestdb_Log

With high growth in archive generation, the protected archive image mount can be more than one mount, for example:

```
/acttestdb_Log
/acttestdb_Log_1
```

2. Set the Oracle environment and use sqlplus to shut down the database:

```
sqlplus / as sysdba
sql> shutdown immediate;
```

Verify the database is shut down. Kill any orphan process for the database.

```
ps -ef | grep <db name>
```

3. Start the database in nomount state using the backup parameter file from the mounted volume. The backup parameter file will be under top mounted folder, for example /acttestdb

4. Set the Oracle environment and use sqlplus to start the database:

```
sqlplus / as sysdba
sql> startup nomount pfile='/acttestdb/<database sid>__backup.ora';
```

5. Create a new spfile from the existing pfile and restart the database.

Create an spfile from the pfile:

```
sql> create spfile='$ORACLE_HOME/dbs/spfile<database sid>.ora' from pfile='/acttestdb/
<database sid>__backup.ora';
```

Restart the database with spfile in nomount state:

```
sql> shutdown immediate;
sql> startup nomount;
```

6. Restore the control file using RMAN from the Actifio mounted archive log image. Use the latest control file from Log mounted image, for example: /acttestdb_Log/cf-D_<sid>-id_<id>.ctl or if more than one log image: /acttestdb_Log_1/cf-D_<sid>-id_<id>.ctl

```
rman target /
rman> restore controlfile from '/acttestdb_Log_1/cf-D_<sid>-id_<id>.ctl' ;
```

7. Mount the database:

```
rman> alter database mount;
```

8. Catalog the datafile and archivelog folder from Actifio mounted database image and archive log image to RMAN

```
rman> run { catalog start with '/acttestdb/datafile' noprompt;
catalog start with '/acttestdb/archivelog' noprompt;
catalog start with '/acttestdb_Log' noprompt;}
```

9. Restore and recover the database:

```
rman> run { restore database ; recover database ; }
```

For a specific point in time recovery using the format `yyyymmddhh24mi`:

```
rman> run  
{  
restore database;  
recover database until time "to_date('<desired time stamp>', 'yyyymmddhh24mi')";  
}
```

10. Open the database with the reset log option:

```
rman> alter database open resetlogs;
```

The database is available for read and write.

Recovering a RAC ASM Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog

Source Database	Protected Under	Archivelog is
RAC or Standalone ASM	File System	Actifio-Protected

To recover a standalone ASM or RAC Oracle database for point-in-time recovery if the logs are Actifio-protected:

1. Mount the image from the source database server to recover. In the Mount window under Mapping Options, provide a mount location for the image, for example: /acttestdb. For instructions on how to mount a database image, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.

The database backup image will be mounted at the mount location: /acttestdb

The protected archive log will be mounted at:/acttestdb_Log

With high growth in archive generation, the protected archive image mount can be more than one mount, for example:

```
/acttestdb_Log
/acttestdb_Log_1
```

2. Shut down the Oracle database. From node 1, su to Oracle OS user:

```
su - oracle
```

Set the Oracle environment and use srvctl to stop the database across all nodes:

```
srvctl stop database -d <database name>
```

Verify the database is shutdown (all nodes). Kill any orphan process for the database.

```
ps -ef | grep <db name>
```

3. Start the database in no-mount state using the backup parameter file from the mounted volume. The backup parameter file will be under the top mounted folder, for example at /acttestdb

4. Set the Oracle environment and use sqlplus to start the database:

```
sqlplus / as sysdba
```

```
sql> startup nomount pfile='/acttestdb/<db name>__backup.ora';
```

5. Create a new spfile and restart the database.

To get the path of original spfile under disk group:

```
cat $ORACLE_HOME/dbs/init<database sid>.ora
```

For example: spfile=+<preferred disk group>/<db name>/spfile<db name>.ora

```
sql> create spfile='+<preferred disk group>/<db name>/spfile<db name>.ora' from
pfile='/acttestdb/<db name>__backup.ora';
```

Restart the database with spfile in nomount state:

```
sql> shutdown immediate;
```

```
sql> startup nomount;
```

6. Restore the control file using RMAN from the Actifio mounted archive log image. Use the latest control file from the Log mounted image (for example: /acttestdb_Log/cf-D_<db name>-id_<id>.ctl or if more than one log image: /acttestdb_Log_1/cf-D_<db name>-id_<id>.ctl

```
rman target /
```

```
rman> restore controlfile from '/acttestdb_Log_1/cf-D_<db name>-id_<id>.ctl' ;
```

7. Mount the database:

```
rman> alter database mount;
```

8. Catalog the datafile and archivelog folder from Actifio mounted database image and archive log image to RMAN:

```
rman> run { catalog start with '/acttestdb/datafile' noprompt;  
catalog start with '/acttestdb/archivelog' noprompt;  
catalog start with '/acttestdb_Log' noprompt;}
```

9. Restore and recover the database:

```
rman> run { restore database ; recover database ; }
```

For a specific point in time recovery using the format `yyyymmddhh24mi`:

```
rman> run  
{  
restore database;  
recover database until time "to_date('<desire time stamp>', 'yyyymmddhh24mi')";  
}
```

10. Open the database with the reset log option:

```
rman> alter database open resetlogs;
```

11. Shutdown the database on node 1 and start the database across all nodes.

Use sqlplus shut down the database:

```
sqlplus / as sysdba  
sql> shutdown immediate;
```

Use srvctl to start database across all nodes:

```
srvctl start database -d <database name>
```

The database is available for read and write.

Recovering an Oracle Database to a Scheduled Backup Point if the archivelog is not Protected through Actifio

Source Database	Protected Under	Archivelog is
RAC or Standalone ASM	ASM Disk Group	Not Actifio-Protected

To recover a standalone ASM or RAC Oracle database for point-in-time recovery if the logs are not protected by Actifio:

1. Mount the image from the source database server to recover. In the Mount window, provide a preferred disk group for the image mount under ASM on RAC Node 1. For details on how to mount a database image, see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access](#) on page 54.

For example, on the mount screen:

- o **Select Host:** RAC node 1 database server
- o **Preferred disk group:** acttestdg
- o **RAC node list:** IP of RAC node 1

The backup parameter file will be copied under /act/touch/<Preferred disk group>, for example:

```
/act/touch/acttestdg/<db name>__backup.ora
```

2. Shut down the Oracle database. From node 1, su to Oracle OS user:

```
su - oracle
```

Set the Oracle environment and use srvctl to stop the database across all nodes:

```
srvctl stop database -d <db name>
```

Verify database is shut down (all nodes in case of RAC). Kill any orphan process for the database.

```
ps -ef | grep <db name>
```

3. Start the database in nomount state using the backup parameter file from the mounted volume. The backup parameter file will be under the top mounted folder, for example at /act/touch/acttestdg/

4. Set the Oracle environment. Use sqlplus to start the database:

```
sqlplus / as sysdba
```

```
sql> startup nomount pfile='/act/touch/acttestdg/<db name>__backup.ora';
```

5. Create a new spfile and restart the database.

To get the path of original spfile under disk group:

```
cat $ORACLE_HOME/dbs/init<database sid>.ora
```

For example: spfile=+<preferred disk group>/<db name>/spfile<db name>.ora

```
sql> create spfile='+<preferred disk group>/<db name>/spfile<db name>.ora' from pfile='/act/touch/acttestdg/<db name>__backup.ora';
```

Restart the database with spfile in nomount state:

```
sql> shutdown immediate;
```

```
sql> startup nomount;
```

6. Restore control file using RMAN from the Actifio mounted volume.

```
rman target /
```

```
rman> restore controlfile from '+<preferred disk group>/<db name>/cf-D_<db name>-id_<id>.ctl' ;
```

7. Mount the database:
`rman> alter database mount;`
8. Catalog the datafile and archive log folder from Actifio mounted ASM disk group to RMAN
`rman> run { catalog start with '+acttestdg/<db name>/datafile' noprompt;
catalog start with '+acttestdg/<db name>/archive log' noprompt; }`
9. Restore and recover the database:
`rman> run { restore database ; recover database ; }`
10. Open the database with the reset log option:
`rman> alter database open resetlogs;`
11. Shutdown the database on node 1 and start the database across all nodes.
Use sqlplus to shut down the database:
`sqlplus / as sysdba
sql> shutdown immediate;`
Use srvctl to start database across all nodes:
`srvctl start database -d <database name>`

The database is available for read and write.

Recovering an Oracle Database to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog

Source Database	Protected Under	Archivelog is
RAC or Standalone ASM	ASM Disk Group	Actifio-Protected

To recover a standalone ASM or RAC Oracle database for point-in-time recovery if the logs are not protected by Actifio:

1. Mount the image from the source database server to recover. In the Mount window, provide a preferred disk group for the image mount under ASM on RAC Node 1. For details on how to mount a database image, see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access](#) on page 54.

For example, on the mount screen:

- o **Select Host:** RAC node 1 database server
- o **Preferred disk group:** acttestdg
- o **RAC node list:** IP of RAC node 1

The backup parameter file will be copied under /act/touch/<Preferred disk group>, for example:

```
/act/touch/acttestdg/<db name>__backup.ora
```

2. Shut down the Oracle database.

From node 1, su to Oracle OS user:

```
su - oracle
```

Set the Oracle environment. Use srvctl to stop database across all nodes:

```
srvctl stop database -d <db name>
```

Verify the database is shut down (on all nodes). Kill any orphan process for the database.

```
ps -ef | grep <db name>
```

3. Start the database in nomount state using the backup parameter file copied under /act/touch/acttestdg
4. Set the Oracle environment. Use sqlplus to start the database:

```
sqlplus / as sysdba
```

```
sql> startup nomount pfile='/act/touch/acttestdg/<db name>__backup.ora';
```

5. Create a new spfile and restart the database.

To get the path of original spfile under disk group:

```
cat $ORACLE_HOME/dbs/init<database sid>.ora
```

For example: spfile=+<preferred disk group>/<db name>/spfile<db name>.ora

```
sql> create spfile='+<preferred disk group>/<db name>/spfile<db name>.ora' from pfile='/act/touch/acttestdg/<db name>__backup.ora';
```

Restart the database with spfile in nomount state:

```
sql> shutdown immediate;
```

```
sql> startup nomount;
```

6. Restore control file using RMAN from the Actifio mounted archive log image. Use the latest control file from Log mounted image (for example: /acttestdb_Log/cf-D_<db name>-id_<id>.ctl or if more than one log image exists: /acttestdb_Log_1/cf-D_<db name>-id_<id>.ctl

```
rman target /
```

```
rman> restore controlfile from '/acttestdg_Log_1/cf-D_<db name>-id_<id>.ctl' ;
```

7. Mount the database:

```
rman> alter database mount;
```
8. Catalog the datafile and archivelog folder from Actifio mounted database image and archive log image to RMAN:

```
rman> run { catalog start with '+acttestdg/<db name>/datafile' noprompt;  
catalog start with '+acttestdg/<db name>/archivelog' noprompt;  
catalog start with '/acttestdg_Log' noprompt;}
```
9. Restore and recover the database:

```
rman> run { restore database ; recover database ; }
```

For a specific point in time recovery run the recover command as under:

```
rman> run  
{  
restore database;  
recover database until time "to_date('<desired time stamp>', 'yyyymmddhh24mi');  
}
```
10. Open the database with the reset log option:

```
rman> alter database open resetlogs;
```
11. Shutdown the database on node 1 and start the database across all nodes.

Use sqlplus to shut down the database:

```
sqlplus / as sysdba  
sql> shutdown immediate;
```

Use srvctl to start the database across all nodes:

```
srvctl start database -d <database name>
```

The database is available for read and write.

12 Recovering Tablespace and Data Files

To recover a single tablespace data file, for example, due to data corruption:

1. Mount the latest database snapshot from the Actifio Appliance back to the Oracle server.
2. Catalog the database backup snapshot to RMAN.
3. Restore and recover the tablespace using the backup snapshot as detailed below.

This section contains procedures for:

[Recovering a Single Tablespace of a Production Database on an ASM Disk Group](#) on page 77

[Recovering a Corrupt Database Block](#) on page 78

[Recovering Lost Control Files](#) on page 79

[Recovering an Oracle Pluggable Database](#) on page 80

Recovering a Single Tablespace of a Production Database on an ASM Disk Group

To recover a single tablespace of production database to the primary node:

1. Mount the database point-in-time snapshot as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
2. In the Mount window under Mapping Options, provide a mount location for the image. For example, for an image under ASM disk group provide a disk group name under Preferred Disk Group and for image under file system provide a mount location ex: /acttestdb.
 - o For instructions on how to mount a database image protected under file system, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.
 - o For details on how to mount a database image protected under ASM Disk Group: see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access](#) on page 54.

3. From the primary node, log into the database server as Oracle OS user.

4. Set the database environment and log into RMAN:

```
rman target /
```

5. At the RMAN prompt, catalog the backup data file and archive log folder:

Example: A database image protected under ASM Disk Group:
(Mounted ASM Disk Group name " + acttestdg"):

```
rman> catalog start with '+acttestdg/<db name>/datafile' noprompt;  
rman> catalog start with '+acttestdg/<db name>/archive' noprompt;
```

```
rman>catalog start with '<mountpoint_log>' noprompt; (If archive logs are protected by Actifio)
```

Example: A database image protected under file system (mounted file system name "/acttestdb"):

```
rman> catalog start with '/acttestdb/datafile' noprompt;  
rman> catalog start with '/acttestdb/archivelog' noprompt;  
rman> catalog start with '/acttestdb_log/archivelog' noprompt; (If archive logs are protected by Actifio)
```

Now you can run all RMAN recovery commands, such as:

- o [Recovering a Tablespace](#)
- o [Recovering a Datafile](#)
- o [Recovering a Corrupt Database Block](#)
- o [Recovering Lost Control Files](#)
- o [Recovering an Oracle Pluggable Database on page 80](#)

6. When finished, unmount and delete the image.

Recovering a Tablespace

To recover a tablespace:

```
rman> restore tablespace <tablespace name>;  
rman> recover tablespace <tablespace name>;
```

Recovering a Datafile

To recover a datafile

```
rman> restore datafile <file#>;  
rman> recover datafile <file#>;
```

Recovering a Corrupt Database Block

To recover a corrupt database block:

1. Mount the database point-in-time snapshot as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
 2. In the Mount window under Mapping Options, provide a mount location for the image. For example, for an image under ASM disk group provide a disk group name under Preferred Disk Group and for image under file system provide a mount location ex: /acttestdb.
 - o For instructions on how to mount a database image protected under file system, see [Mounting an Oracle Database Image Protected Under a File System for Data Access on page 56](#).
 - o For details on how to mount a database image protected under ASM Disk Group: see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access on page 54](#).
 3. From the primary node, log into the database server as Oracle OS user.
 4. Set the database environment and log into sqlplus, then query v\$database_block_corruption to check the corrupt blocks:

```
sqlplus / as sysdba  
sql> SELECT * FROM V$DATABASE_BLOCK_CORRUPTION;
```
 5. Login to RMAN to recover all corrupted blocks:

```
rman target /  
rman> RECOVER CORRUPTION LIST;
```

After the blocks are recovered, the database removes them from V\$DATABASE_BLOCK_CORRUPTION.
 6. To recover an individual corrupt block (ex: datafile 8 and block 13):
From RMAN prompt
RMAN> recover datafile 8 block 13;
-

Recovering Lost Control Files

To recover lost control files:

1. Mount the database point-in-time snapshot as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
2. In the Mount window under Mapping Options, provide a mount location for the image. For example, for an image under ASM disk group provide a disk group name under Preferred Disk Group and for image under file system provide a mount location ex: /acttestdb.
 - o For instructions on how to mount a database image protected under file system, see [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.
 - o For details on how to mount a database image protected under ASM Disk Group: see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access](#) on page 54.
3. From the primary node, log into the database server as Oracle OS user.
4. Set the database environment and log into sqlplus, then shut down the database and start in nomount state:

For standalone database:

```
sqlplus / as sysdba
sql> shutdown immediate;
sql> startup nomount;
```

For RAC database from the mounted image node shutdown the database across all nodes:

```
srvctl stop database -d <dbname>
sql> startup nomount;
```

5. Restore the control file from Actifio mounted image as above.
For example: /acttestdb (Filesystem) and +acttestdg (for ASM)

```
rman target /
rman> restore controlfile from '/acttestdb/cf-D_<db name>-id_<id>.ctl' ; (Filesystem mount)
rman> restore controlfile from '+acttestdb/cf-D_<db name>-id_<id>.ctl' ; (ASM mount)
```
6. Mount and open the database from sqlplus:

```
sqlplus / as sysdba
sql> alter database mount;
sql> recover database until cancel;
sql> alter database open resetlogs;
```

Recovering an Oracle Pluggable Database

To recover an Oracle 12c pluggable database (PDB):

1. Mount the image from the source database server to recover. In the Mount window, provide a preferred disk group for the image mount under ASM on RAC Node 1. For details on how to mount a database image, see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access](#) on page 54.
2. Close the pluggable database:
 - a. From Node 1, su to Oracle User
su - oracle
 - b. Set the Oracle environment. Connect to the Oracle database as "sysdba" user
sqlplus / as sysdba
SQL> alter pluggable database <Pluggable DB name> close;
3. Catalog the datafile and archivelog folder from Actifio mounted database image and archive log image to RMAN:

```
rman> run { catalog start with '+acttestdg/<db name>/datafile' noprompt;  
          catalog start with '+acttestdg/<db name>/archivelog' noprompt;  
          catalog start with '/acttestdg_Log' noprompt;  
        }
```
4. Restore and recover the pluggable database

```
rman> run  
    {  
      restore pluggable database <Pluggable DB name>;  
      recover pluggable database <Pluggable DB name> until time "to_date('<desired time  
stamp>', 'yyyymmddhh24mi')";  
    }
```

Note: These steps are applicable only for Oracle 12.1.0.2 version and above. If Oracle version is 12.1.0.1 or below 12.1.0.2, then recovery of all datafiles belonging to the Pluggable database must be performed after before executing Step 5.

For example: rman> recover datafile <PDB datafile number>;

5. Open the pluggable database:

```
rman> alter pluggable database <Pluggable DB name> open;
```

The Pluggable database is open for read and write.

13 Instant Oracle Database Recovery or Migration Using ASM Switch and Rebalance

You can also perform this procedure much more easily through the AGM; see [Chapter 14, Performing an Oracle ASM Switch and Rebalance](#).

Use this in case of storage failure or to migrate a database to new storage. There are five steps to this:

1. [Stop the Database to be Recovered on All Nodes](#) on page 81.
2. [Select and Mount an Image](#) on page 81.
3. [Switch the Database to the Newly Mounted Disk Group](#) on page 82.
At this point the database is up and running in its original configuration on Actifio storage. Now you can migrate the database back to production storage or to new storage.
4. [Migrate the Database Back to Production Storage \(Rebalance Operation\)](#) on page 84.
After the database migration, clean up the no-longer-needed mounted images.
5. [Unmount and Delete Actifio Mounted Images Used for the ASM Rebalance Operation](#) on page 84.

Stop the Database to be Recovered on All Nodes

1. From the protected node, log in as Oracle OS user and set the database environment.

```
srvctl stop database -d <DB Name>
```
2. (Optional) If you want to retain the production disk group name post-recovery, then connect to the ASM instance on each non-primary node and dismount the disk group:

```
SQL> alter diskgroup <DG name> dismount;
```


Then connect to the ASM instance on the primary node and drop the ASM diskgroup:

```
SQL> drop diskgroup <DG name> including contents;
```


To list and check the existing disk path for diskgroup:

```
set lines 200
set pages 500
col name format a15
col path format a25
select PATH from v$asm_disk where group_number in (select group_number from
v$asm_diskgroup where name in upper('<disk group name','disk group name') ) order by
group_number
```

Select and Mount an Image

3. Mount the image as described in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#). Mount the image to all RAC nodes of the source database. During the mount, select node 1 as the target host and under RAC Node List provide the IP of the RAC nodes in order node1:node2:node3. The first IP address in RAC Node list must be the selected host's IP address.

Switch the Database to the Newly Mounted Disk Group

This script can be run as root OS user, Oracle OS user, or service account user. The behavior of the script running as different OS user is:

Root OS user: Script will not prompt for any OS user password. The script will use su to Oracle and grid user account to run the required commands.

Oracle OS user: Script will prompt for grid OS user password if the grid account is not using the Oracle OS user (Oracle database instance and ASM instance is not running under Oracle OS user).

Service account user: Script will prompt for Oracle OS user and grid OS user password.

4. From the protected node:

To run the script as root OS user:

Login as root OS user.

cd to script folder /act/act_scripts/asmrestore

change permission to 755

chmod -R /act/act_scripts/asmrestore

or:

To run the script as an Oracle OS user or as a service account user:

- o Login as OS user.
- o create a folder tempASMRestore

mkdir tempASMRestore

cd tempASMRestore

- o copy the ASM restore script to this local folder

cp /act/act_scripts/asmrestore/*

- o change permission to 755

chmod -R *

5. In the same folder, edit the asm_node_config.conf file to create the list of RAC nodes to switch, restore, and start the database on all nodes. Each line should have only one node entry. For example for a 3 node RAC add one line for each node to asm_node_config.conf:

```
ASM_NODE1:rhel137-14.dev.actifio.com
```

```
ASM_NODE2:rhel137-15.dev.actifio.com
```

```
ASM_NODE3:rhel137-16.dev.actifio.com
```

6. Run the ASMOracleRestore.sh script.

./ASMOracleRestore.sh uses 5 required and two optional input parameters:

usage: ASMOracleRestore.sh <Oracle Home path> <Disk Group Name> <Oracle User> <Grid User>
<standalone yes/no> <log mount path> <timestamp yyyyymmddhh24mi>

For example:

```
ASMOracleRestore.sh /home/oracle/app/oracle/product/11.2.0/dbhome_1 DATA1 oracle grid  
no /act/mnt/vdbxlog 201404021435
```

Note: "log mount path", and "timestamp" are optional parameters that should be used only when the database and archive log are both protected using a policy with log protection enabled.

log mount path	timestamp	behavior
not provided	not provided	Database will be recovered to the database backup point.
provided	not provided	Database will be recovered to the database backup point. and Archive log will be rolled forward to all available logs under the log mount path.
provided	provided	Database will be recovered to the database backup point. and Archive log will be rolled forward to the provided timestamp under the log mount path. The timestamp format is <code>yyyymmddhh24mi</code> .

The production database will be up and running in the same configuration as the source database was in when the Actifio presented copy snapshot was taken.

Migrate the Database Back to Production Storage (Rebalance Operation)

From the protected node, use the same OS user account and script directory that were used in [Switch the Database to the Newly Mounted Disk Group](#) on page 82. This script can be run as root OS user or Oracle OS user or non-Oracle OS user. The behavior of the script depends upon the role that runs it:

Running the script as:

Root OS user: Script will not prompt for any OS user password. The script will use su to Oracle and grid user accounts to run the required commands.

Oracle OS user: Script will prompt for grid OS user password if the grid account is not using the Oracle OS user i.e. Oracle database instance and ASM instance is not running under oracle OS user.

Service account user: Script will prompt for Oracle OS user and grid OS user password

1. Create a file called asm_disks.conf, with the list of disk paths for the production disks to be added to the Actifio mounted disk group for rebalancing. The file should include a single line for each disk path:

```
/dev/sda1  
/dev/sdc1
```

2. Run the script ./asmDgRebalance.sh

This script needs 3 input parameters:

```
asmDgRebalance.sh <Disk group Name> <Oracle user> <Grid User>
```

For example: asmDgRebalance.sh DATA1 oracle grid

This script adds the list of disks from asm_disks.conf file to the Actifio mounted disk group, performs the rebalance operation, and then drops the Actifio disks from the disk group.

Note: The Oracle ASM rebalance operation uses a range of power (from 1 to 11) to run the rebalance in the background. The Actifio script can run with power of 11. To change the power, edit asmDgRebalance.sh

3. To verify the ASM rebalance operation status, run the query below from node 1 as grid OS user.
4. From the script folder, set the ASM environment and connect to sqlplus as sysasm:

```
Sqlplus / as sysasm  
SQL> @checkRebalanceOperation.sql
```

When prompted, enter the value for dg_name: <Actifio mounted Disk Group name>

5. When finished, unmount and delete the image.

Unmount and Delete Actifio Mounted Images Used for the ASM Rebalance Operation

On all RAC nodes:

1. Login as root user.
2. cd to /act/touch
3. Open the hidden file (dot).<dg_name>_switch_conf (to list the file, run the command ls -la)
4. Set the value of UNMAP to YES: UNMAP=YES
5. Unmount the images.

Note: Even after a rebalance-based restore is completed and Actifio disks are dropped from the production disk group, Oracle ASM still keeps hold on the underlying block devices for some time. This is due to a known ASM bug. This may result in unmount/delete job failures for some time as the LUNs could not be cleaned from the host gracefully. To work around this ASM issue, retry unmount/delete jobs after some time.

14 Performing an Oracle ASM Switch and Rebalance

You can protect an Oracle ASM instance either as an Oracle ASM diskgroup or as a filesystem. If an Oracle ASM instance uses an ASM diskgroup as backup destination, you gain the capability for restore and recovery by using ASM switch. This is particularly useful for very large databases where traditional RMAN restore would take too long to satisfy the RTO requirements, since RMAN restore has to physically move data from backup to original diskgroup.

If the database is backed up under file system, then you must use the traditional RMAN recovery method in [Mounting an Oracle Database Image Protected Under a File System for Data Access](#) on page 56.

Note: The ASM switch and rebalance operation cannot be performed on protected virtual databases (application aware mounts).

Note: For the switch and rebalance procedure to work, /etc/hosts must be appropriately populated.

The switch and rebalance procedure has two stages:

1. Mounting the image as single diskgroup and then **switching** the database running out of an Actifio mounted diskgroup. You can provide a preferred disk group name which will remain as the production disk group after the rebalance operation where the data gets moved to production storage from Actifio storage.
2. **Rebalancing** moves the data to production storage from Actifio storage. This is an online operation where data movement happens in the background by the Oracle ASM API.

This procedure works for:

- Databases configured as standalone ASM, single node RAC with ASM, multi-node RAC with ASM.
- ASM disk groups configured as: database using single disk group for datafile, database using multiple disk group for datafile, and database using multiple disk group and sharing the disk group with another database on the same server.

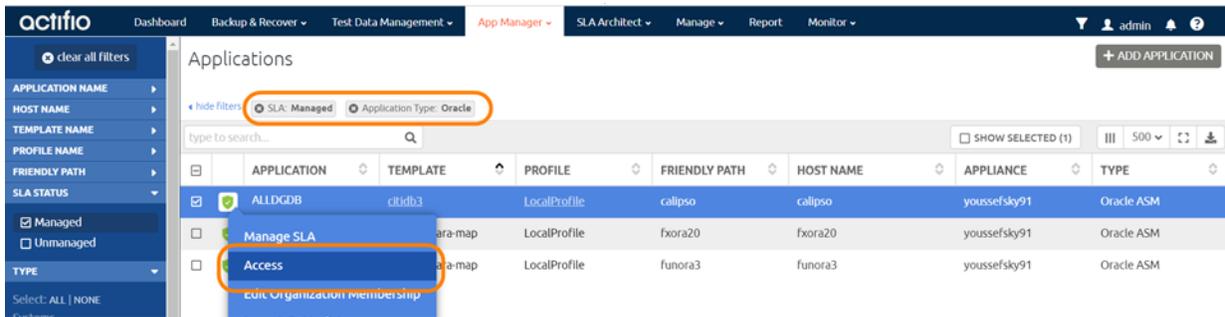
The RMAN image copy of all data files for the entire database is captured on an Actifio-presented ASM disk group retaining the ASM header information. A snapshot of the staging disk with ASM header information is taken.

Procedure

To instantly recover an Oracle ASM database from AGM:

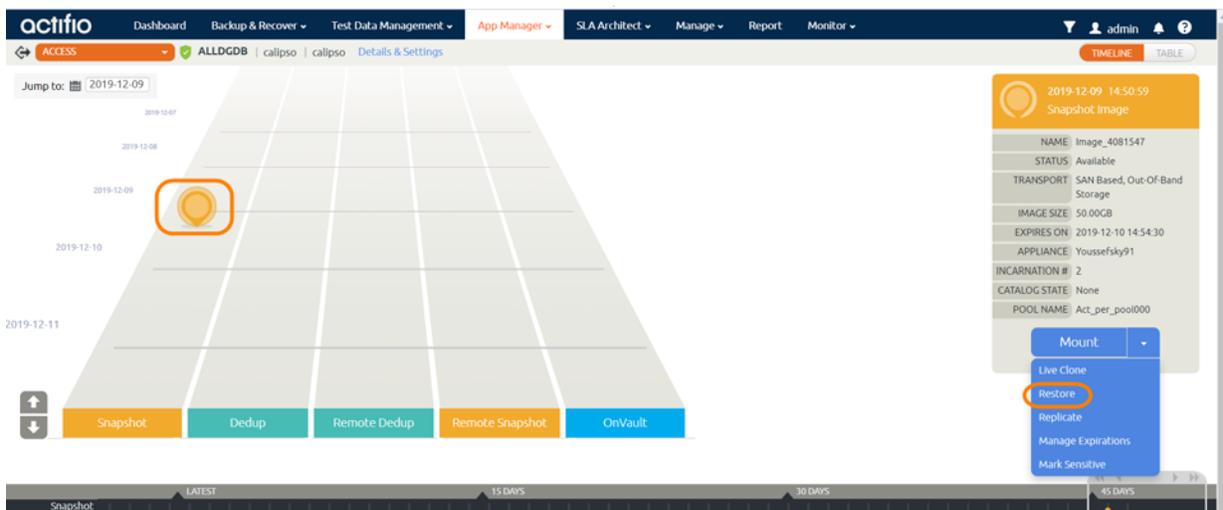
1. Create a candidate ASM diskgroup for the restored database. You can:
 - o Enter a new diskgroup name
 - o Use the failed diskgroup name: first delete the failed diskgroup and then create a new diskgroup with the same name, and prepare it as an ASM candidate.

- Open the AGM to the App Manager > Applications.
- Use the filter feature to search for the desired database. Select the database to be recovered and click **Access**.



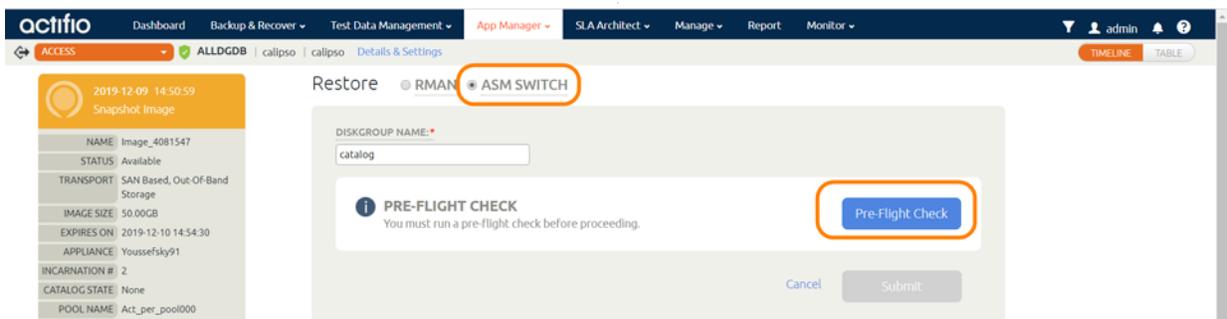
Right-Click the Selected Database and Click Access

- On the Access page, select the desired image and click **Restore** under the Mount menu.



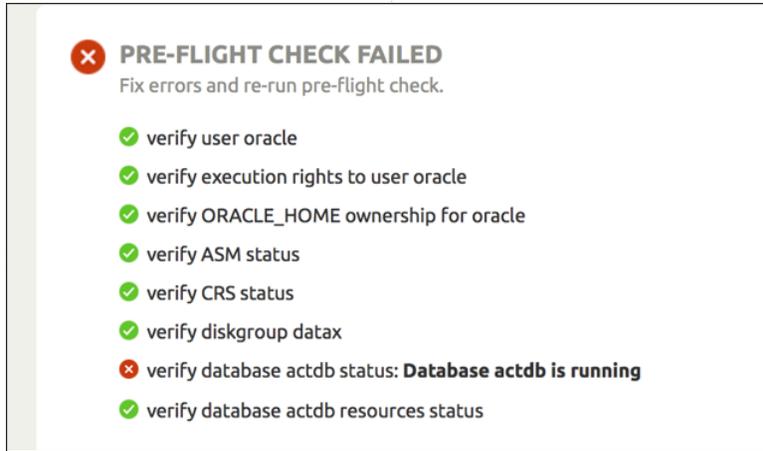
Select the Desired Image and Click Restore

- On the Restore page, select **ASM Switch**.
- Select a time if needed.
- Under **Diskgroup Name**, enter the name of the diskgroup candidate.



Select ASM Switch and Enter the Diskgroup Name

- Run the **Preflight Check**. The results will point you to any remedial steps.

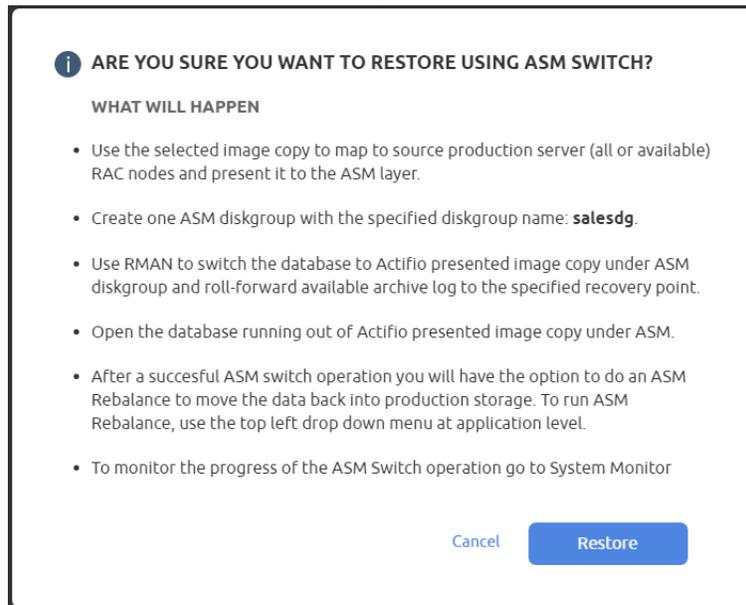


PRE-FLIGHT CHECK FAILED
Fix errors and re-run pre-flight check.

- ✓ verify user oracle
- ✓ verify execution rights to user oracle
- ✓ verify ORACLE_HOME ownership for oracle
- ✓ verify ASM status
- ✓ verify CRS status
- ✓ verify diskgroup datax
- ✗ verify database actdb status: **Database actdb is running**
- ✓ verify database actdb resources status

Address Any Errors, then Run the Pre-Flight Check Again

9. Address any errors, then run the pre-flight check again. When the pre-flight check passes without issues, click **Submit**. You see an informational screen.



ARE YOU SURE YOU WANT TO RESTORE USING ASM SWITCH?

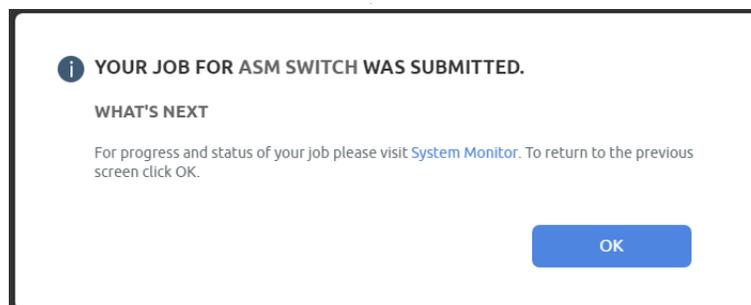
WHAT WILL HAPPEN

- Use the selected image copy to map to source production server (all or available) RAC nodes and present it to the ASM layer.
- Create one ASM diskgroup with the specified diskgroup name: **salesdg**.
- Use RMAN to switch the database to Actifio presented image copy under ASM diskgroup and roll-forward available archive log to the specified recovery point.
- Open the database running out of Actifio presented image copy under ASM.
- After a successful ASM switch operation you will have the option to do an ASM Rebalance to move the data back into production storage. To run ASM Rebalance, use the top left drop down menu at application level.
- To monitor the progress of the ASM Switch operation go to System Monitor

Cancel Restore

This Informational Screen Tells You What Will Happen

10. Click **Restore**. The job begins.



YOUR JOB FOR ASM SWITCH WAS SUBMITTED.

WHAT'S NEXT

For progress and status of your job please visit [System Monitor](#). To return to the previous screen click OK.

OK

Waiting for the Restore/Switch Job to Complete

- You can view progress on the Jobs Monitor in another browser instance. The Job Type is Restore (ASM Switch).

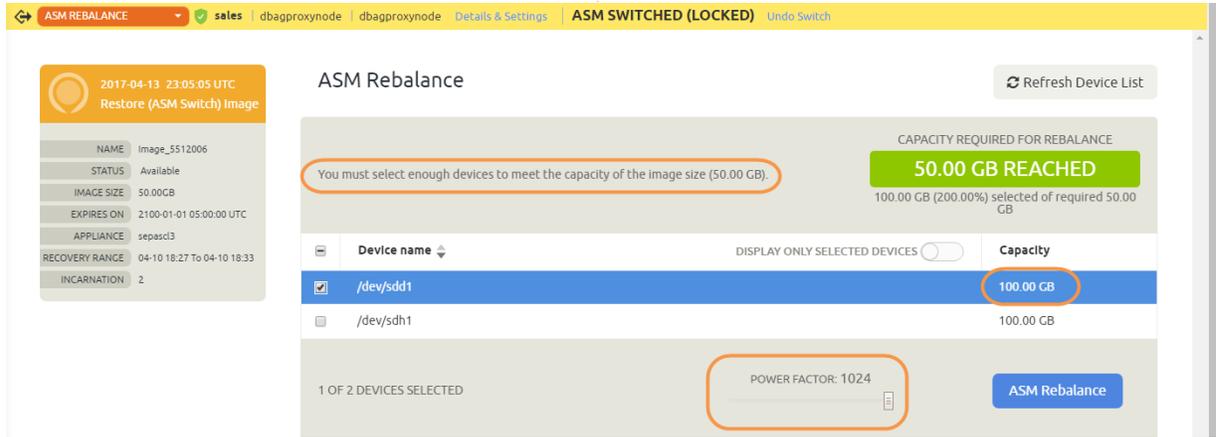
Job_5512021	Running: 4%	oravmn2	cchssdb	04-13 23:05:...	Production to...	snapshot (DB...	Oracle_ASM_...	0
Job_5511785b	Queued	bb6aix7	aixfs	04-13 23:20:...	dbSnap	snapshot	LogSmart	0
Job_5512006	Running: 97%	dbagproxyno...	sales	04-13 23:04:...	Production to...	Restore (ASM...	Oracle_ASM_...	0
Job_5511782a	Queued	oravmn2	accwfdb	04-13 23:05:...	Production to...	snapshot	test_dedup	0

Viewing Job Progress in the Jobs Monitor

- When the restore job has finished, go back to the original AGM browser instance and click **OK**.
- The next step is the rebalance operation. At the top of the window is an ASM Rebalance button. (Beside it is an Undo Switch button in case you have reason to stop this process.) To continue, click the **ASM Rebalance** button.

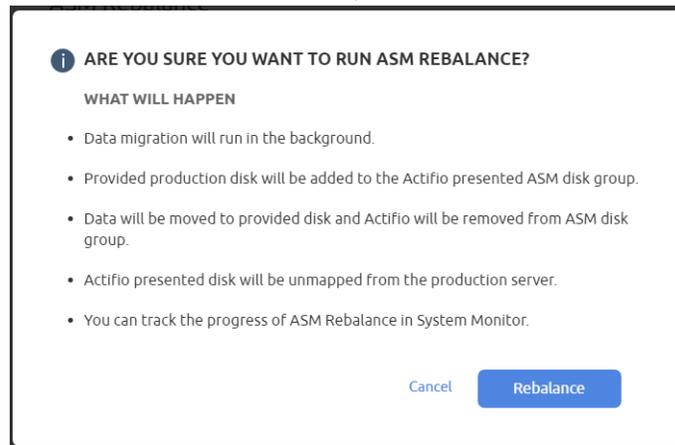
The Database is Switched and Ready to be Rebalanced

- The ASM Rebalance screen appears. You can see the size of the image and the capacity of the available storage devices. When you have selected a storage device that can accommodate the image, the Capacity Required indicator turns green and the ASM Rebalance button turns blue. Now you can proceed.
- Before starting the ASM Rebalance operation, select a power factor at the bottom of the display. Lower values use fewer system resources, but they take much longer. In a recovery operation, you may want to select the highest value for the fastest results. Select a **Power Factor** and then click **ASM Rebalance**.



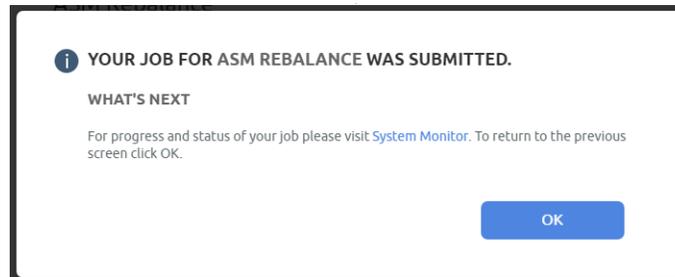
Selecting a Storage Device and a Power Factor

16. A screen appears explaining what happens next. Click **Rebalance** to submit the rebalance job.



Submit the Rebalance Job

17. As with the switch job, you can view progress on the Jobs Monitor in another browser instance. The Job Type is Restore (ASM Rebalance).



Waiting for the Rebalance Job to Complete

When the job is finished, the database is ready for use.

15 Protecting and Recovering Oracle Databases in a Windows Environment

Oracle database protection in a Microsoft Windows environment has these two limitations:

- All Oracle databases, both those under file system and those under ASM disk group, are protected under file system only.
- OS Authentication is not available for databases in a Windows environment. In Application Details & Settings, at **Username and Password** enter an Oracle user `act_rman_user` username and password for database authentication. Make sure the database user account has the proper role granted based on the **User Role in the Database** advanced setting. Application Details & Settings are detailed in [Application Details & Settings for Oracle Databases](#) on page 35.

This section includes:

[Preparing Oracle Protection in a Windows Environment](#) on page 91

[Identifying Database Instances On Windows](#) on page 92

[Backing Up an Oracle Database in a Windows Environment](#) on page 93

[Watch Script to Watch for Database Volumes Being Mounted](#) on page 94

[Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point if the archive log is Not Protected through Actifio](#) on page 95

[Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archive log](#) on page 96

Preparing Oracle Protection in a Windows Environment

Before you can discover, protect, and mount out-of-band Oracle databases, be sure to:

1. Check the following settings on the database server:
 - a. The Oracle database and the Oracle Listener are up and running (check Windows services).
 - b. Confirm that there is a tns entry with the name SID. The file `tnsnames.ora` is under `%ORACLE_HOME%\network\admin`
 - c. Verify tns entry is valid by running:
`%ORACLE_HOME%\bin\tnsping <SID>`
 - d. Verify the database is running with spfile. From sqlplus login as sysdba:
`sqlplus / as sysdba`
`sql> show parameter spfile;`
 - e. Verify the database is in archive mode. From sqlplus login as sysdba:
`sqlplus / as sysdba`

```
sql> archive log list;
```

2. Get an RMAN user account with “sysdba” and “create session” or “connect” privileges for configuring the RMAN backup. To verify the connection, as Oracle OS user set the ORACLE environment by running from the command line:

```
sqlplus <RMAN user account>/<password>@<SID> as sysdba;
```

3. Enable database change block tracking. With database BCT off, incremental backup time is impacted. Change block tracking feature is available in Oracle Enterprise Edition. Run a SQL query to check that change block tracking is enabled. Run the query:

```
sqlplus / as sysdba  
sql> select * from v$block_change_tracking;
```

Identifying Database Instances On Windows

1. On Windows, to find out what databases are on a host, use:

```
reg query HKLM\System\CurrentControlSet\Services | findstr OracleService
```

2. This returns a line out output that looks like this:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\OracleServiceORCL
```

3. This shows a database called ORCL. To identify the ORACLE_HOME directory and other details, run:

```
reg query HKLM\Software\Oracle /v ORA* /s
```

4. This returns the following which tells us the ORACLE_HOME directory and whether the database is set to start on boot.

```
HKEY_LOCAL_MACHINE\Software\Oracle\KEY_OraDb11g_home1  
ORACLE_HOME REG_SZ D:\oracle\product\11.2.0\dbhome_1  
ORACLE_HOME_NAME REG_SZ OraDb11g_home1  
ORACLE_GROUP_NAME REG_SZ Oracle - OraDb11g_home1  
ORACLE_BUNDLE_NAME REG_SZ Enterprise  
ORAMTS_CP_TRACE_LEVEL REG_SZ 0  
ORAMTS_CP_TRACE_DIR REG_SZ D:\oracle\product\11.2.0\dbhome_1\oramts\Trace  
ORAMTS_CONN_POOL_TIMEOUT REG_SZ 120  
ORAMTS_SESS_TXNTIMETOLIVE REG_SZ 120  
ORAMTS_NET_CACHE_MAXFREE REG_SZ 5  
ORAMTS_NET_CACHE_TIMEOUT REG_SZ 120000  
ORAMTS_OSCREDS_MATCH_LEVEL REG_SZ OS_AUTH_LOGIN  
ORACLE_SID REG_SZ orcl  
ORACLE_HOME_KEY REG_SZ SOFTWARE\ORACLE\KEY_OraDb11g_home1  
ORACLE_BASE REG_SZ D:\oracle  
ORA_ORCL_AUTOSTART REG_EXPAND_SZ TRUE  
ORA_ORCL_SHUTDOWN REG_EXPAND_SZ TRUE  
ORA_ORCL_SHUTDOWNTYPE REG_EXPAND_SZ immediate  
ORA_ORCL_SHUTDOWN_TIMEOUT REG_EXPAND_SZ 90
```

5. Next, we can see if the database is running, using this command:

```
tasklist /SVC | findstr oracle
```

6. If the instance is started, you should see a line of output like this:

```
oracle.exe 1492 OracleServiceORCL
```

Backing Up an Oracle Database in a Windows Environment

Actifio VDP does not yet support OS authentication for Oracle databases in a Windows environment. You can create a username and password inside the database and grant it rights to perform the backup:

1. Launch SQLPlus:
Set ORACLE_SID=orc1
sqlplus / as sysdba
2. Create the user and grant the necessary rights:
sql> create user ACT_RMAN_USER identified by mypassword;
sql> grant create session, resource, sysdba to act_rman_user;

Note: In an Oracle 12c environment, you can grant sysbackup role instead.

3. Check if block change tracking is enabled:
sql> select * from v\$dbblock_change_tracking;
If the status is disabled, enable it (optional but recommended for optimal backup performance):
sql> alter database enable block change tracking using file
'D:\oracle\product\11.2.0\dbhome_1\dfs\orc1.bct';

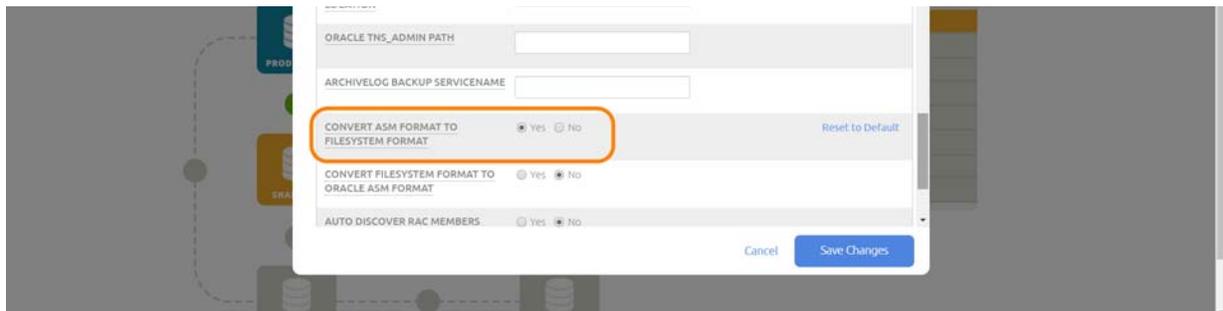
4. Check if the log mode is set to Archive Log mode:
sql> archive log list
If the database is in no-archive log mode, then return it to archive log mode:

Note: This will take down the database.

```
sql> shutdown immediate;  
sql> startup mount;  
sql> alter database archive log;  
sql> alter database open;
```

5. Set the Application Details & Settings. In particular:
 - o Open the AGM and enter the database username and password (defined above) in the Application Details & Settings of the database.
 - o If the database is under Oracle ASM Disk Group, then set Convert ASM Format to Filesystem Format to **Yes**.

For full information on Application Details & Settings, see [Application Details & Settings for Oracle Databases](#) on page 35.



6. Apply an SLA to protect the database.

Watch Script to Watch for Database Volumes Being Mounted

If you create an application-aware mount, then you can use a watch script to show the volumes being mounted from Actifio, and the Oracle processes running. Application-aware mounts are described in [Mounting an Oracle Database as a Virtual Application](#) on page 58.

When performing an application-aware mount, you can use this watch script. The script location must be: C:\Program Files\Actifio\scripts. Scripts run on Windows hosts must be .bat or .vbs files.

```
@echo off
:loop
echo. > watchtemp
echo ----- >> watchtemp
echo Oracle Processes >> watchtemp
echo ----- >> watchtemp
tasklist /svc | findstr oracle >> watchtemp
echo. >> watchtemp
echo ----- >> watchtemp
echo Actifio Mounts >> watchtemp
echo ----- >> watchtemp
wmic volume get label, name | findstr Actifio >> watchtemp
echo. >> watchtemp
cls
type watchtemp
timeout 2 > null
goto loop
```

Which produces output like this:

```
Oracle Processes
-----
oracle.exe           1492 OracleServiceORCL
oracle.exe           3768 OracleServiceTestDB
oracle.exe           872 OracleServiceTestDB2

-----
Actifio Mounts
-----
Actifio-Backup-ORCL  D:\mount_1
Actifio-Backup-ORCL  Y:
```

Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point if the archivelog is Not Protected through Actifio

To recover an entire out-of-band Oracle database in a Windows environment:

1. Mount the database backup snapshot from Actifio back to the Oracle server as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
2. Set the database environment and start the database in no-mount state using the parameter file from the Actifio mounted volume (for example for a mounted database volume to E:\):

```
sqlplus / as sysdba
sql> startup nomount pfile='E:\<sid>__backup.ora';
```
3. Create spfile from pfile:

```
sql> create spfile='%ORACLE_HOME%\database\spfile<sid>.ora' from
pfile='E:\<sid>__backup.ora';
```
4. Start the database with spfile in the nomount state:

```
sql> shutdown immediate;
sql> startup nomount;
```
5. Restore the control file using RMAN from the Actifio mounted volume:

```
rman target /
rman> restore controlfile from 'E:\cf-D_<sid>-id_<id>.ctl' ;
```
6. Mount the database:

```
rman> alter database mount;
```
7. Catalog the datafile and the archive file folder from the Actifio mounted volume to RMAN:

```
rman> run
{
catalog start with 'E:\datafile' noprompt;
catalog start with 'E:\archivelog' noprompt;
}
```
8. Restore and recover the database:

```
rman> run
{
restore database;
recover database;
}
```
9. Roll forward the logs as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
10. Open the database with reset log option:

```
rman> alter database open resetlogs;
```

Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point with Roll-Forward of Actifio-Protected archivelog

1. Mount the image from the source database server to recover. Mount the backup snapshot from Actifio back to the Oracle server as detailed in [Chapter 10, Accessing, Recovering, or Restoring an Oracle Database](#).
2. Set the database environment and start the database in no-mount state using the parameter file from the Actifio mounted volume. In this example, the database backup image will be mounted at: E:\ and the protected archive log will be mounted at: Z:\.

```
sqlplus / as sysdba
sql> startup nomount pfile='E:\<sid>__backup.ora';
```

3. Create spfile from pfile:

```
sql> create spfile='%ORACLE_HOME%\database\spfile<sid>.ora' from
pfile='E:\<sid>__backup.ora';
```

4. Start the database with spfile in the nomount state:

```
sql> shutdown immediate;
sql> startup nomount;
```

5. Restore the control file using RMAN from the Actifio mounted archive log image:

```
rman target /
rman> restore controlfile from 'Z:\cf-D_<sid>-id_<id>.ctl' ;
```

6. Mount the database:

```
rman> alter database mount;
```

7. Catalog the datafile and archivelog folder from Actifio mounted database image and archive log image to RMAN:

```
rman> run
{
catalog start with 'E:\datafile' noprompt;
catalog start with 'E:\archivelog' noprompt;
catalog start with 'Z:\archivelog' noprompt;
}
```

8. Restore and recover the database:

```
rman> run
{
restore database;
recover database;
}
```

For a specific point in time recovery run the recover command as below:

```
rman> run
{
restore database;
recover database until time "to_date('<desired time stamp>', 'yyyymmddhh24mi)";
}
```

9. Open the database with reset log option:

```
rman> alter database open resetlogs;
```

The database is available for read and write.

16 Using Actifio VDP with Oracle Exadata Database or Oracle ExaCC

This chapter includes:

[Using Actifio VDP with Oracle Exadata in an iSCSI Environment](#) on page 97

[Using Oracle Exadata with Actifio VDP in a dNFS Environment](#) on page 98

Actifio Appliances support capture and presentation of Exadata data over iSCSI or Oracle dNFS protocols.

- The Actifio Appliance is connected over iSCSI or Oracle dNFS in the network (not in the data path).
- RMAN backup uses RMAN to directly write to copy data store presented by Actifio as a file system or as an ASM Disk Group.
- Data Capture Formats: under ASM Disk Group (iSCSI only) or under File System (dNFS or iSCSI).
- Actifio incremental-forever backup uses RMAN Incrementally Updated Backups, rolling forward image copy backups.

Actifio Capture of Exadata Data

The Actifio Connector must be installed on the Exadata server to facilitate communication with the Actifio Appliance and to invoke the RMAN API for database backup.

Using Actifio VDP with Oracle Exadata in an iSCSI Environment

The Actifio Connector exposes and maps Actifio disk(s) to the Exadata server as an iSCSI target. Data Capture format can be under ASM Disk Group or under the File System.

Install the Actifio Connector on each Exadata host under user space to facilitate the communication with Actifio Appliance and to invoke the RMAN API for database backup.

During a backup, the connector will:

1. Map and expose the logical disk to the Exadata server as an iSCSI target.
2. Add the Actifio disk path to the ASM disk string.
3. Make sure the ASM disk string is added to the parameter file and does not exist in the CRS profile.
4. Create an ASM disk group as an external redundancy using Actifio disk.
 - o RMAN backup using RMAN to directly write to copy data store presented by Actifio Appliance as ASM Disk Group or as File system
 - o Incremental forever backup using RMAN Incrementally Updated Backups, rolling forward image copy backups

Using Oracle Exadata with Actifio VDP in a dNFS Environment

Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP). Direct NFS is built directly into the database kernel, just like ASM.

The dNFS protocol can be used for filesystem-based backup as an NFS share.

The Actifio Connector will expose and map Actifio disk(s) to Exadata server as NFS share.

Pre-requisites for dNFS on Exadata server:

- Enable dNFS on Exadata server:

```
cd $ORACLE_HOME/rdbms/lib  
make -f ins_rdbms.mk nfs on
```
- Restart the database

Use RMAN API to back up the database to filesystem on dNFS share presented by Actifio Appliance.

Note: dNFS supports backups under filesystem only; it does not support ASM Switch & Rebalance.

Actifio Connector

During a backup, the connector will:

1. Map and expose the logical disk to the Exadata server as an NFS target.
2. RMAN backup using RMAN to directly write to copy data store presented by Actifio Appliance using dNFS as File system
3. Incremental forever backup using RMAN Incrementally Updated Backups, rolling forward image copy backups

17 Protecting SAP ECC/BW with an Oracle Database

There are four steps to protecting SAP ECC/BW:

1. [Protecting the Oracle Database](#) on page 99
2. [Refreshing the Database](#) on page 99
3. [After the Refresh, on the Target Oracle Database](#) on page 99
4. [After the Refresh, on the Target SAP Application Server](#) on page 100

Protecting the Oracle Database

To protect the Oracle database, see [Chapter 9, Virtualizing an Oracle Database for Data Protection and Agility](#).

Refreshing the Database

Use an application aware mount to refresh the target test/dev database:

1. On the target SAP database server and application server, stop the SAP application and database.
2. Follow the pre-refresh and post-refresh activity for system copy in the SAP system copy guide: 1738258 - **System Copy of Systems Based on SAP NetWeaver**.
<https://websmp104.sap-ag.de/public/instguides>
3. Use an Actifio application aware mount to refresh the target Oracle database as detailed in [Mounting an Oracle Database as a Virtual Application](#) on page 58.

After the Refresh, on the Target Oracle Database

On Target SAP Database server:

1. Check the OPSS\$<OS_USER> in the database: (always enter <os_user> in uppercase)

```
SELECT * FROM DBA_USERS WHERE USERNAME = 'OPSS$<os_user>';
```
2. If the system does not return an entry, create the user:

```
CREATE USER "OPSS$<os_user>" DEFAULT TABLESPACE <user_tsp>  
TEMPORARY TABLESPACE PSAPTEMP IDENTIFIED EXTERNALLY;
```
3. Ensure that the name of the OPSS\$ user is specified entirely in uppercase
4. The table SAPUSER must occur in the system only once and it must be assigned to the user OPSS\$<sid>ADM. Use the following query to check this:

```
SELECT OWNER FROM DBA_TABLES WHERE TABLE_NAME = 'SAPUSER';
```

If the system returns an owner <owner> other than OPSS<sid>ADM, delete the relevant SAPUSER tables:

```
DROP TABLE "<owner>".SAPUSER;
```

If the system does not return OPSS<sid>ADM, then create the table SAPUSER as <sid>adm and enter the password:

```
CREATE TABLE "OPS${<sid>ADM}".SAPUSER  
(USERID VARCHAR2(256), PASSWD VARCHAR2(256));
```

```
INSERT INTO "OPS${<sid>ADM}".SAPUSER VALUES ('<sapowner>', '<password>');
```

After the Refresh, on the Target SAP Application Server

1. Modify the profiles for dbs_ora_schema to the right schema name
2. Run R3trans -d on the application server and make sure the return code is 000
3. Import the license key:

```
saplikey pf=/usr/sap/<SID>/SYS/profile/<instance_profile> -install /<directory>/  
license.txt
```

The license.txt file can be generated from the SAP site for the application server. This is specific for the hardware key that identifies the application server from SAP point of view.

18 Oracle RMAN Logs

This chapter details:

[Oracle Protection Logs on Linux](#) on page 101

[Oracle Protection Logs on Windows](#) on page 102

[Retrieving a Specific Oracle DB Archive Log Sequence Number from an Actifio Backup Image](#) on page 103

Oracle Protection Logs on Linux

These are the logs that you might need to consult:

Table 1: Oracle Protection Logs: Standard Mounts

Log	Location	What's In It	What to Look For
Connector log	/var/act/log/UDSAgent.log	For any mount job.	Any error with ORA-
Database RMAN backup log	/var/act/log/<database name>_rman.log	All the backup command and output for database and archive log backup.	ORA- and RMAN-errors in the log file.

Application Aware mounts produce additional logs. Check the below logs on the target database server:

Table 2: Oracle Protection Logs: Application Aware Mounts

Log	Location	What's In It	What to Look For
Connector log	/var/act/log/UDSAgent.log	For any mount job.	Any error with ORA-
Database under file system	/act/act_scripts/oracleclone/dbrecover_<dbname>.txt	Progress of the archive log roll-forward.	Errors EXCEPT those relating to a log looking for an archive during the roll-forward.
Database under ASM Disk Group	/act/act_scripts/asmclone/dbrecover_<dbname>.txt /act/act_scripts/asmclone/openDBlog_<dbname>.txt	Steps to configure target database post-roll-forward.	Any error with ORA-

Oracle Protection Logs on Windows

These are the logs that you might need to consult:

Table 3: Oracle Protection Logs: Standard Mounts

Log	Location	What's In It	What to Look For
Connector log	C:\Program Files\Actifio\log If the Actifio Connector is installed on a different drive, then use that drive letter.	For any mount job.	Any error with ORA-
Database RMAN backup log	C:\act_tmp\log	All the backup command and output for database and archive log backup.	ORA- and RMAN-errors in the log file.

Application Aware mounts produce additional logs. Check the below logs on the target database server:

Table 4: Oracle Protection Logs: Application Aware Mounts

Log	Location	What's In It	What to Look For
Connector log	C:\Program Files\Actifio\log If the Actifio Connector is installed on a different drive, then use that drive letter.	For any mount job.	Any error with ORA-
Database protected under file system	C:\Program Files\Actifio\act_scripts\oracleclone\	Progress of the archive log roll-forward.	Errors EXCEPT those relating to a log looking for an archive during the roll-forward.
Database protected under ASM Disk Group	C:\Program Files\Actifio\act_scripts\oracleclone\ C:\Program Files\Actifio\act_scripts\oracleclone\	Steps to configure target database post-roll-forward.	Any error with ORA-

Retrieving a Specific Oracle DB Archive Log Sequence Number from an Actifio Backup Image

Environment: Primary database and standby database.

Use case: Actifio backup job is running from the primary database. The standby database is stuck due to a missing archive log. You need the missing archive log from the Actifio backup at the primary database.

Example: Retrieving archive log sequence 74343 from Actifio backup image

On the primary database:

```
RMAN> list backupset 155219; List of Backup Sets ===== BS Key Size Device
Type Elapsed Time Completion Time -----
----- 155219 11.10M DISK 00:00:00 04-SEP-19 BP Key: 156789 Status: AVAILABLE
Compressed: NO Tag: TAG20190904T081447 Piece Name: /act/mnt/Staging_268638/archivelog/
5auatbjn_1_1 List of Archived Logs in backup set 155219 Thrd Seq Low SCN Low Time Next
SCN Next Time ----
----- 1 74343
324399126320 04-SEP-19 324399162514 04-SEP-19
```

To apply archive logs from the backup piece:

1. Mount backup image as a standard mount (not AppAware) on standby database host.
2. Catalog the backup piece from Actifio log staging disk using the command.
"catalog start with ?/<mountpoint>/archivelog' noprompt;"
or catalog the specific archive log backup piece using the command:
"catalog '/<mountpoint>/archivelog/5auatbjn_1_1'; "
3. List backup of archivelog logseq=74343;
4. Then apply the logs by running the command:
{ set archivelog destination to '/ora_backup/rman/arch/'; restore archivelog from
logseq=74343 until logseq=<endlogsequence number> thread=<threadno>; }`
5. Apply the logs in the standby database.
6. Unmount the Actifio mount locations from standby host.

This works for all supported Oracle Database versions irrespective of the operating system. Using the simple mount feature of Actifio of the protected primary database on the standby database host, you can apply the archives on the standby database host to sync the primary database.

Note: A missing archive log is more likely to happen in the standby database host.

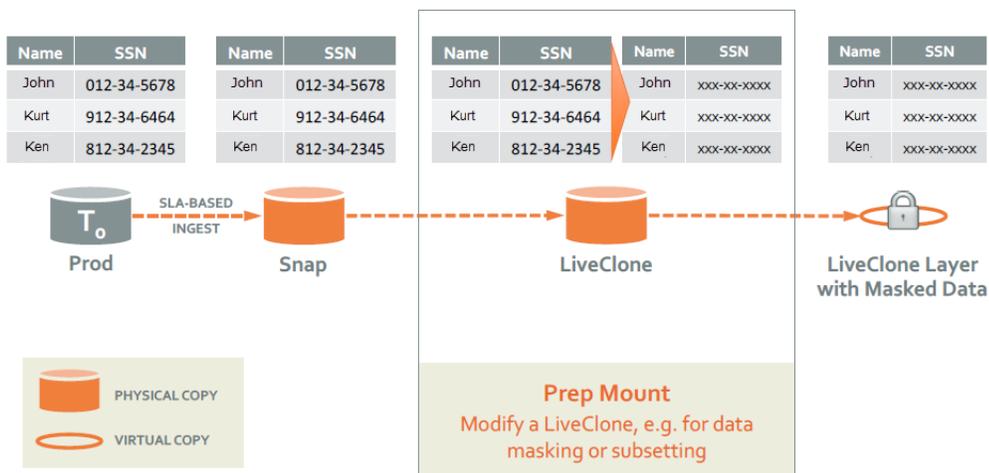
19 Introduction to Provisioning Environments With Workflows

Actifio Workflows automate access to captured data. Workflows can run according to a schedule or on demand. Workflows are built with captured production data. Workflows can present the captured data as a LiveClone, a virtual application, or as just the application data.

For those applications that contain sensitive data, a Workflow can include a step that creates a LiveClone and then automatically runs a script to mask the sensitive data.

The LiveClone with the masked data can then be mounted as a virtual application or the LiveClone's data to provision a work environment. For example, the following figure illustrates a Workflow that:

- Captures production data that contains Social Security numbers.
- Prep-mounts the captured data as a LiveClone so a script can scrub the Social Security numbers from the data.
- Mounts the scrubbed LiveClone as a virtual application via an Application Aware mount.
- Flags the scrubbed virtual application as non-sensitive.



Once a Workflow completes, users can access the server on which the virtual application has been mounted and use the data. Access to the data is controlled by the server on which the data is mounted. This chapter provides an overview of Actifio Workflows, including:

- [Workflow Benefits](#) on page 106
- [Workflow Data Access Methods](#) on page 106
- [Workflow Access Control](#) on page 107

Workflow Benefits

The Actifio Appliance allows you to capture production data via a schedule or on demand. Once the data is captured it can be used in a Workflow to provision work environments.

With Actifio Workflows:

- Work environments can be updated with production data as soon as new data is available without having to wait days or even weeks for a DBA or system administrator to do the provisioning.
- Work environments can be updated on a predefined schedule or on demand.
- User Acceptance Testing (UAT) is streamlined. Actifio can provision virtual applications (virtual copies of production applications) in UAT environments using the same protocols and environments that exist in production.
- DBAs and system administrators are freed from constantly provisioning and refreshing work environments.
- Users can quickly and efficiently reproduce and address issues using the latest production data in a secure, isolated environment.
- Software updates can be applied and tested, using the latest production data before being released in to a production environment.
- Multiple teams can work in parallel and instantly access their own copy or virtual copy of production data.

Actifio Workflows take provisioning one step further by providing the ability to scrub images of sensitive data. This ability allows users to self-provision their environments with virtual applications of production data while at the same time maintaining data security.

Workflow Data Access Methods

Workflows present captured data as Application Aware Mounts, LiveClones, or Standard Mounts.

Application Aware Mount

The Actifio Application Aware mount function provides instant access to an application as a virtual application without actually moving data. Applications can be mounted on any application server. If the application is a database, and its logs have been captured, you can use the logs to roll the database forward to a desired point in time. Application Aware mounts are performed from the Actifio Appliance and do not require manual intervention by application, database, server, or storage administrators.

Once an Application Aware Mount finishes, the virtual application can be captured like any other application. This capability is particularly useful when a problem is encountered with a virtual application. Users can roll back their copy of the virtual application to a point where the issue does not exist, while other users troubleshoot the problematic version.

LiveClone

A LiveClone is an independent copy of an application that can be refreshed on a schedule or on demand. A LiveClone can be scrubbed for sensitive data before being made available to users. This allows teams such as development and test to use production quality data without exposing sensitive data and interfering with the production environment. LiveClones are updated on demand or according to a schedule defined in the Workflow.

Direct Mount

Direct mounts present only an application's data. Direct mounts are updated on demand or according to a schedule defined in the Workflow.

Workflow Access Control

Actifio Appliances come with a predefined Actifio Admin user. The Actifio Admin is a super user who has full control of and access to all features, functions, and resources of an Actifio Appliance.

The Admin User can create users with various degrees of access.

When creating Workflows the Admin User must consider which users will have access to Workflows and what access those users will have to an Actifio Appliance's features, functions, and managed data - including access to sensitive data.

This chapter consists of the following topics:

- [Configuring Roles, Organizations, and Users](#) on page 107
- [Example Role for Limited Workflow Access](#) on page 108
- [Example Organization for Limited Workflow Access](#) on page 108
- [Example User for Limited Workflow Access](#) on page 109

Configuring Roles, Organizations, and Users

Team leaders are often tasked with running Workflows on demand and provisioning the work environment for their team. Team leaders typically are not the users who create Workflows, but are often responsible for running Workflows.

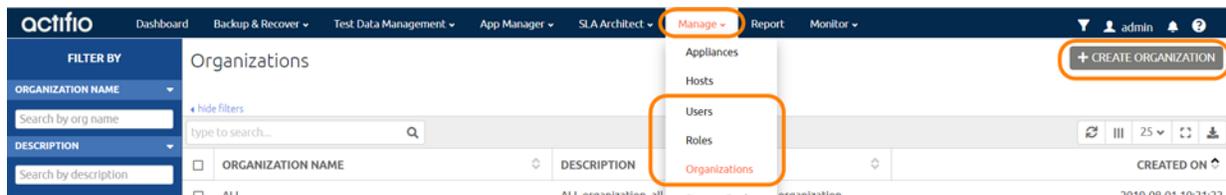
Giving team leaders the ability to run Workflows on demand frees DBAs and system administrators from having to update work environments.

A Workflow's on-demand capabilities allow team leaders to update their environments with refreshed production data as needed. For example, if a Workflow is scheduled to run every Sunday night, a team leader can run a Workflow on demand on a Wednesday and update the work environment with the latest production data.

To create an Actifio user who has just enough access to an Actifio Appliance to run Workflows and to provision their environment with the latest Actifio Workflow updates, you can create:

- A **role** that has the necessary privileges to log in to the Actifio Appliance, access the App Manager, manage mounts, run Workflows, and view the status of a Workflow's progress.
- An **organization** that limits the user to Workflow data captured from specified application servers and if necessary, specific applications.
- A **user** who may or may not have access to sensitive data, and is assigned to an appropriate role and organization.

New roles, organizations, and users are created in the AGM Manager:



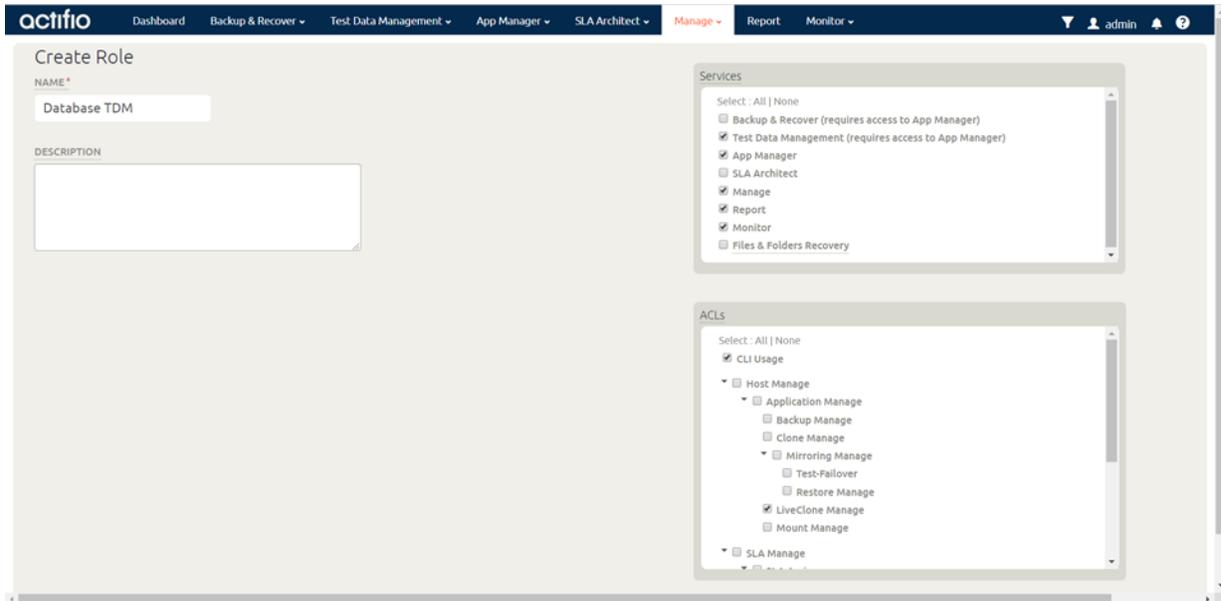
Details on creating new roles, organizations, and users can be found in the AGM online help.

Use the following examples as a quick reference for creating users who do not have access to sensitive data, and are assigned roles and organizations in such a way that they can only use an Actifio Appliance to provision their work environment.

Example Role for Limited Workflow Access

The following screen capture is an example of a user role that limits a user to accessing an Actifio Appliance's App Manager, Manage LiveClones, Manage Mounts, and to view and run Workflows. It also allows that user to monitor the progress of a Workflow via the Jobs Monitor. A single role can be assigned to multiple users.

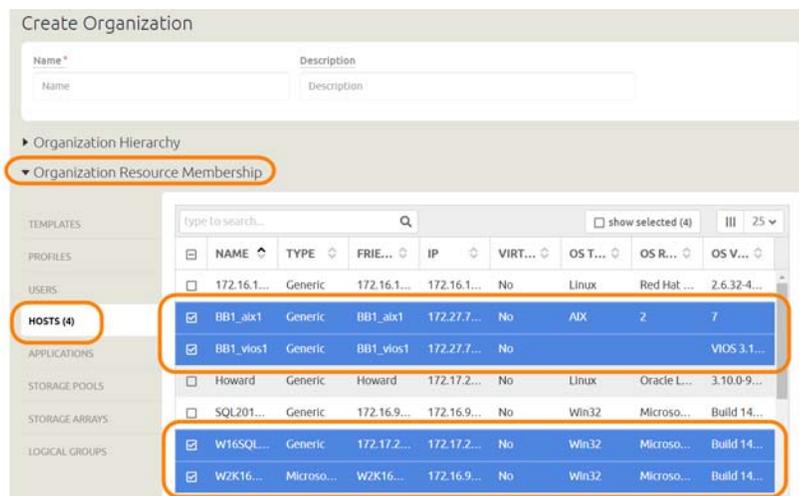
Note: Because a LiveClone in a LiveClone Workflow typically contains sensitive information, users assigned a role with Manage LiveClone rights must also have access to sensitive data. See [Example User for Limited Workflow Access](#) on page 109.



Example Role for Limited Workflow Access

Example Organization for Limited Workflow Access

The following screen capture is an example of an organization that limits a Workflow user to specific application servers. If necessary, the organization could limit a Workflow user to specific applications on the servers. A single organization can be assigned to multiple users.



Example Organization for Limited Workflow Access

Example User for Limited Workflow Access

The following screen capture is an example of a new user who does not have access to sensitive data.

Note: Such users do not have access to sensitive data in a Workflow even if their assigned organizations contain sensitive data.

Roles and Organizations are assigned to a user by clicking the tabs at the top of the page. A single user can be assigned multiple roles and multiple organizations.

Note: Because a LiveClone in a LiveClone Workflow typically contains sensitive information, users assigned a role with Manage LiveClone rights must also have access to sensitive data. See [Example Role for Limited Workflow Access on page 108](#).

The screenshot shows the 'Create User' form in the actifio interface. The form includes the following fields and options:

- USERNAME ***: DB Cooper
- AUTHENTICATE LOCALLY**:
- PASSWORD ***: [Redacted]
- CONFIRM ***: [Redacted]
- FIRST NAME**: Database
- LAST NAME**: Cooper
- EMAIL**: dbcooper@mycompany.com
- TIMEZONE ***: Africa/Brazzaville
- ENABLE CLI**:
- ACCESS TO SENSITIVE DATA**: (highlighted with an orange circle)
- COMMENTS**: [Empty text area]
- ROLES**: Select All / None. Search by Role. Selected roles: Basic, Test Data Management Admin.
- ORGANIZATIONS**: Select All / None. Search by Organization. Selected organizations: ALL, PUBLIC.

Buttons: Cancel, Create User

Example User for Limited Workflow Access

20 Useful Workflows

Workflows are defined in the App Manager. Workflows use captured production data as the source data from which they generate:

Direct Mounts: Direct Mounts are used when application data does not need to be scrubbed before it is mounted and made available to users. The application data can be mounted with a standard mount as just the data, as a virtual application, or as both the data and as a virtual application. Direct mounts make captured data available almost instantly without actually moving data.

LiveClones: LiveClones are typically used in Workflows when an application contains sensitive data which must be scrubbed before it is mounted and made available to users. The application data can be mounted with a standard mount as just the data, as a virtual application, or as both the data and as a virtual application. Updates to the LiveClone, scrubbing, and mounting can be done automatically via a schedule or on demand.

Note: Workflows consist of a number of options. Not specifying an option will allow a user to specify that option when running the Workflow on demand.

This chapter provides step-by-step examples for defining a:

- [Direct Mounting Application Data or a Virtual Application](#) on page 112
- [Creating Scrubbed Data or a Scrubbed Virtual Application with a LiveClone and Masking Tools](#) on page 115
- [Using an Actifio Workflow to Refresh Oracle Database Schemas](#) on page 119
- [Presenting an Oracle 12c Database PDB as a Virtual PDB to an Existing Database Container on a Target](#) on page 123

Direct Mounting Application Data or a Virtual Application

Direct Mount Workflows are used when the application image does not contain sensitive data. This allows the Workflow to mount an application's data as a virtual application or as just application data without first having to define a LiveClone and then scrub the LiveClone. Direct Mounts make application data and virtual applications available almost instantly.

In this example, you define a Workflow that:

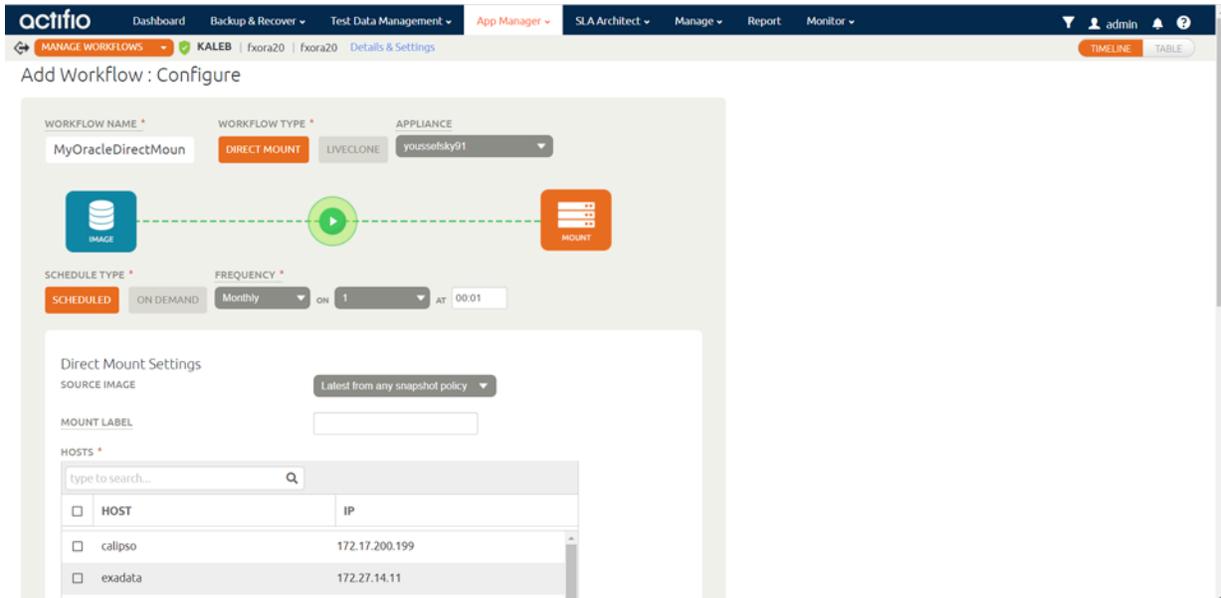
- Generates or updates mounted application data from a single selected production image.
- Defines a schedule for updating the application data with the latest production data. Scheduled Workflows can also be run on demand.
- Mounts the application data.
- Allows you to create a virtual application with the data.

To define this Workflow for a managed database:

1. Open the App Manager to Workflows.

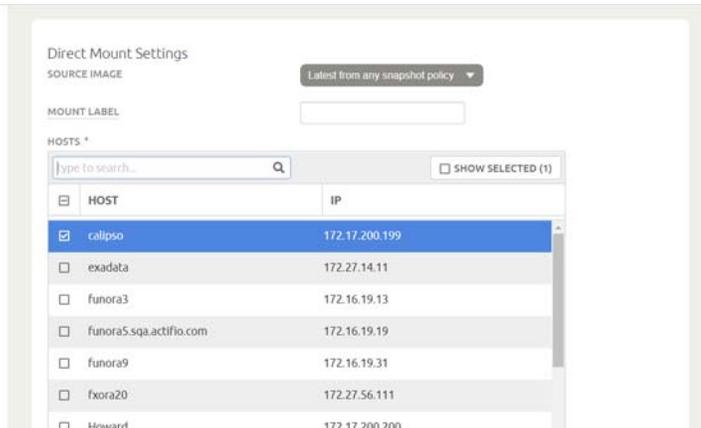
Note: To modify an existing workflow for an application, go to Applications and right-click the application. Then select Manage Workflows.

2. Click **+ Workflow** in the upper right corner of the page. The Add Workflow: Select an Application page is displayed. Right-click a managed database and click **Next**.
3. In the Add Workflow: Configure page, specify a name for the Workflow, select **Direct Mount**, and select the appliance to manage the workflow. Workflow names cannot include special characters.



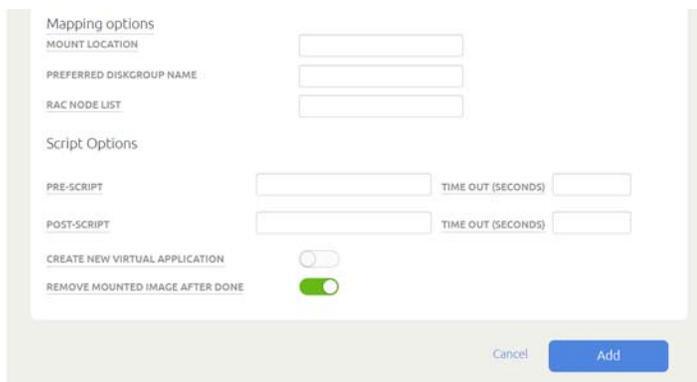
New Workflow Settings Page

4. For Schedule Type, select **Scheduled** and assign a frequency for the workflow to run. The time selector uses a 24-hour clock. Scheduled Workflows can also be run on demand.
5. In Direct Mount Settings, for **Source Image**, select the snapshot image to run the workflow on and give the mount a label to make it easily identifiable.
6. In the **Hosts** section, select the server(s) on which the data will be mounted. The user who runs the Workflow will be given the option of mounting the data to other servers as needed.



Direct Mount Settings Options

7. In the Mapping Options, add a **mount location**. If you do not specify a mount, then the Actifio Connector will choose a mount location.
8. Enter a **Preferred diskgroup name** as needed.
9. Enter a **RAC node list** as needed in a colon-separated list of IP addresses.



Mapping Options and Script Options

10. Specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting or unmounting the application. This script must reside in a folder named `/act/scripts` on the server that will host the mounted image. See [Workflow Pre and Post Scripts](#) on page 127 for scripting details.
11. Specify a **Post Script** as needed. The post script used to perform an operation on the data after it is mounted or unmounted, for example, initiating a copy to tape operation for long-term storage. This script must reside in a folder named `/act/scripts` on the server that will host the mounted image. See [Workflow Pre and Post Scripts](#) on page 127 for scripting details.
12. If you specified scripts, specify timeouts long enough for the scripts to complete.
13. **Remove mounted image after processing.** This is typically checked when you employ a script to process mounted data. Once the script finishes its task, this option will unmount and delete the virtual application.
14. Decide:
 - o If you need only the data, then click **Done** and the Workflow will run as scheduled. Users with proper access can also run this Workflow on demand to provision and re-provision their environments with the application data.
 - o If you want to create a virtual application, continue:

15. Check **Perform application aware mount** and the application specific options to perform an application aware mount are displayed.

Note: The **Remove mounted image after processing** option disappears when you select **Perform application aware mount**. The **Remove mounted image after processing** option is typically used with pre and/or post scripts. It unmounts the image from target server, and deletes it after the script(s) finish. This option is typically used when scripts perform operations on the mounted data.

CREATE NEW VIRTUAL APPLICATION

TARGET DATABASE SID * Must be unique and does not e...

USER NAME *

ORACLE HOME DIRECTORY *

MANAGE NEW APPLICATION

▼ Advanced Options

PASSWORD

TNS ADMIN DIRECTORY PATH

DATABASE MEMORY SIZE IN MB

SGA %

REDO SIZE

SHARED_POOL_SIZE IN MB

DB_CACHE_SIZE IN MB

DB_RECOVERY_FILE_DEST_SIZE IN MB

INMEMORY_SIZE IN MB FOR VESION 12C OR HIGHER

DIAGNOSTIC_DEST

MAX NUMBER OF PROCESSES

MAX NUMBER OF OPEN CURSORS

TNS LISTENER IP

Application Aware Mount Options

16. In the spaces provided, enter the Oracle related options. Click the question mark **?** next to an option for help text on that option.
17. Click the black arrow to open Advanced Options and scroll down to fill in the remainder of the options as needed.
18. Check **Restore with Recovery**. Doing so leaves the database in a state where if logs are available they can be applied to bring the database to a specific point in time.
19. At this point you can:

Check the **Protect new application** option. This allows you to apply an SLA Policy Template and a Resource Profile to protect the virtual application as a new application. For details on SLA Policy Templates and Resource Profiles see the AGM online help.

Click **Done** and the Workflow will run as scheduled. Users with proper access also can run this Workflow on demand to provision and re-provision their environments with the virtual application.

Creating Scrubbed Data or a Scrubbed Virtual Application with a LiveClone and Masking Tools

LiveClone Workflows are typically used when the production data to be used contains sensitive information. The LiveClone Workflow allows you to define a LiveClone from production data and then scrub the LiveClone's data before mounting it.

In the following example, a Workflow is defined that:

- Generates a new, or updates an existing LiveClone from a selected production image that is marked as sensitive.
- Defines a schedule for updating the LiveClone with the latest production data. Scheduled Workflows can also be run on demand.
- Employs scripts to scrub the LiveClone of sensitive data.
- Marks the scrubbed virtual application as non-sensitive.
- Mounts the scrubbed image as a virtual application (Application Aware Mount) to a specified host.

To define this Workflow:

1. Open the AGM to **App Manager > Workflows**.
2. In the upper right corner, click **+ Add Workflow**.
3. In the Add Workflow: Select an Application page, right-click the application to be scrubbed of its sensitive data and select **Next**. The Add Workflow: Configure page is displayed.

Note: If the application already has a Workflow, click the **Add** button.

4. Specify a name for the Workflow. Workflow names cannot include special characters.
5. For Workflow Type select **LiveClone**. The LiveClone Settings appear.

The screenshot displays the 'Add Workflow: Configure' interface in the Actifio App Manager. At the top, the navigation bar includes 'Dashboard', 'Backup & Recover', 'Test Data Management', 'App Manager', 'SLA Architect', and 'Manage'. Below the navigation bar, the breadcrumb trail shows 'MANAGE WORKFLOWS > FUNSDB | funora5.sqa.actifio.com | funora5.sqa.actifio.com Details & Settings'. The main form is titled 'Add Workflow: Configure' and contains the following sections:

- Workflow Name:** A text input field containing 'New Workflow'.
- Workflow Type:** A dropdown menu with 'DIRECT MOUNT' and 'LIVECLONE' (selected).
- Appliance:** A dropdown menu with 'yousebsky01'.
- Schedule Type:** A dropdown menu with 'SCHEDULED' (selected) and 'ON DEMAND'.
- Frequency:** A dropdown menu with 'Monthly' (selected), 'ON', and 'AT'.
- Live Clone Settings:** A section with the following fields:
 - SOURCE IMAGE:** A dropdown menu with 'Latest from any snapshot policy' (selected).
 - LIVECLONE LABEL:** An empty text input field.
 - DISK POOL:** A dropdown menu with 'act_per_pool000' (selected).
 - MOUNT FOR PRE-PROCESSING:** A toggle switch that is currently turned off.

At the bottom of the form, there are three buttons: 'Cancel', 'Next', and 'Add'.

Configuring a LiveClone Workflow

6. For Schedule Type, select either **Scheduled** or **On Demand**. Scheduled Workflows can also be run on demand.

Note: You can run a Scheduled Workflow on-demand at any time. You cannot convert an On-Demand Workflow to a Scheduled Workflow after it has been created.

7. If you selected **Scheduled**, define a schedule to run the Workflow. The time selector uses a 24-hour clock.
8. For **Source Image**, select the latest snapshot image from the policy that produces the image to scrub.

Note: If you selected On Demand, the user running the Workflow can select the source snapshot image at run time.

9. Give the LiveClone a label to make it easily identifiable.
10. Select the Snapshot Pool in which the LiveClone will reside. Ensure enough storage space is allocated to accommodate the LiveClone. See the AGM online help for details on configuring Snapshot Pools.

Enable **Mount for pre-processing** and the following screen provides the opportunity to mount the LiveClone to a host and then scrub the mounted image for sensitive data:

LiveClone Mount for Pre-Processing Options

11. The **Mark Data Non-Sensitive** option only appears if the selected application is marked as containing sensitive data. Check this option if the application will not contain sensitive data after scrubbing. This will allow users who do not have access to sensitive data to access the Workflow and provision their work environments. See [Example Organization for Limited Workflow Access](#) on page 108 for details.
12. Select a **Host** on which the LiveClone will be mounted.
13. In Mapping Options, add a **Mount Location** if you want the image mounted to a location other than the default. If you do not specify a mount location, then the Actifio Connector will assign a mount location.
14. Enter a **Preferred diskgroup name** as needed.
15. Enter a **RAC node list** as needed in a colon-separated list of IP addresses.

16. You can use IBM Optim Data Privacy Masking or you can use your own custom scripts:
 - o If you use Optim Data Privacy Masking, enter a Service Set and a Time Out value in minutes.
 - o If you use Custom Scripts, specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting or unmounting an application. Then specify a **Post Script** as needed. The post script in this example will be used to scrub the application of sensitive information after it has been mounted. Specify timeouts long enough for the scripts to complete.

Note: Custom scripts must reside in a folder named /act/scripts on the server that will host the mounted image. See [Workflow Pre and Post Scripts](#) on page 127 for scripting details.

17. If you do not need a virtual application, then click **Add** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand and provision and re-provision any server to which they have access with the application data.

If you need a virtual application of this LiveClone, continue here.

18. Enable Create New Virtual Application and click **Next**. Additional Oracle Settings options are displayed. Fill these in.

The screenshot shows a configuration form for creating a new virtual application. At the top, there is a toggle switch labeled "CREATE NEW VIRTUAL APPLICATION" which is currently turned on. Below this, there are several input fields for basic configuration: "TARGET DATABASE SID" (with a note "Must be unique and does not e..."), "USER NAME", "ORACLE HOME DIRECTORY", "PASSWORD", "TNS ADMIN DIRECTORY PATH", "DATABASE MEMORY SIZE IN MB", "SGA %", "REDO SIZE", "SHARED_POOL_SIZE IN MB", "DB_CACHE_SIZE IN MB", "DB_RECOVERY_FILE_DEST_SIZE IN MB", "INMEMORY_SIZE IN MB FOR VESION 12C OR HIGHER", "DIAGNOSTIC_DEST", "MAX NUMBER OF PROCESSES", "MAX NUMBER OF OPEN CURSORS", "TNS LISTENER IP", "TNS LISTENER PORT", "TNS DOMAIN NAME", and "PDB PREFIX". Below these fields is a section titled "Advanced Options" which is expanded to show several toggle switches: "DO NOT CHANGE DATABASE DBID", "NO ARCHIVE MODE", "CLEAR ARCHIVELOG", "DO NOT UPDATE TNSNAMES.ORA", and "DO NOT UPDATF ORATAB".

LiveClone Application Specific Options

19. Enable Manage New Application if you want VDP to protect it. If you select this, then you will have to select a policy template and a resource profile to apply to it.

20. Open the **Advanced Options** by clicking the arrow. Check **Restore with Recovery**. Doing so leaves the database in a state where if logs are available they can be applied to bring the database to a specific point in time.
21. Continue filling in the Oracle settings as needed for this database. Click on each label for help.
22. At the bottom, enable **Remove mounted image after done**. This is typically checked when you employ a script to process mounted data. Once the script finishes its task this option will unmount and delete the virtual application.
23. Click **Add**.

Using an Actifio Workflow to Refresh Oracle Database Schemas

If you are using Oracle 12c, then refer to [Presenting an Oracle 12c Database PDB as a Virtual PDB to an Existing Database Container on a Target](#) on page 123.

If you have a source database instance with multiple applications, and each application has its own schema, you can:

- Create and refresh a virtual copy at the schema level, and refresh each schema individually to the same target or to a different target.
- Create and refresh multiple virtual copies of a single schema to a single target under different schemas, each with its own refresh schedule.

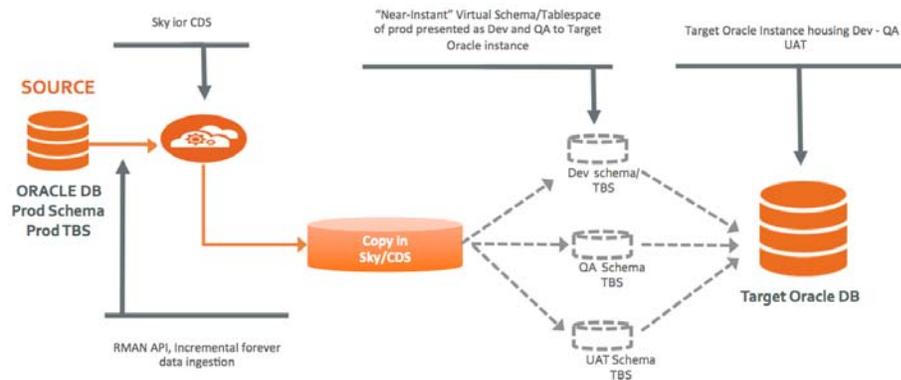
This section includes:

[Before You Begin](#) on page 120

[Creating the Workflow](#) on page 120

[Running the Workflow](#) on page 122

[Unmounting Mounted Images](#) on page 122



How it works: Actifio virtual copy with transportable tablespace

Create and refresh multiple virtual copies under a single target Oracle instance from the same or from multiple source instances.

1. Set up an application-aware mount workflow to create an application-aware mount with a temporary instance.
2. Set up a pre-script to run on the target instance (offline and drop the tablespace to be refreshed on the target instance).
3. Set up a post-script to be run on the temporary instance and the target instance.
4. On the temporary instance:
 - a. Rename the tablespace
 - b. Change the tablespace to read-only
 - c. EXPDP: Export the tablespace metadata.
 - d. Shut down the temporary instance.
5. On the target instance:
 - a. IMPDP: Import the tablespace metadata (TRANSPORT_DATAFILES with remap schema to target instance)
 - b. Change the tablespace to read-write

Before You Begin

Before you begin:

1. Discover and protect the database as described in the AGM online help.
2. Ensure database backups are running as per the SLA policy.
3. Set up the scripts on the target server:
 - a. Login to database server as root and cd to /act (cd /act)
 - b. Create /act/scripts directory if not there:

```
mkdir scripts
cd /act/scripts
```
 - c. Copy all files from /act/act_scripts/objectrefresh/
 - d. Move act_<schema>_refresh.conf to act_testuser_refresh.conf for a target schema name of testuser where <schema> is testuser.

```
[oracle@asmracnode3 scripts]$ cat act_testuser_refresh.conf
SOURCE_SCHEMA_NAME=TEST_USER
TARGET_SID=schpta
TARGET_SCHEMA_NAME=TEST_USER3
[oracle@asmracnode3 scripts]$
```

where:

SOURCE_SCHEMA_NAME: Source database schema to be presented to target.

TARGET_SID: Target database SID, where schema will be refreshed.

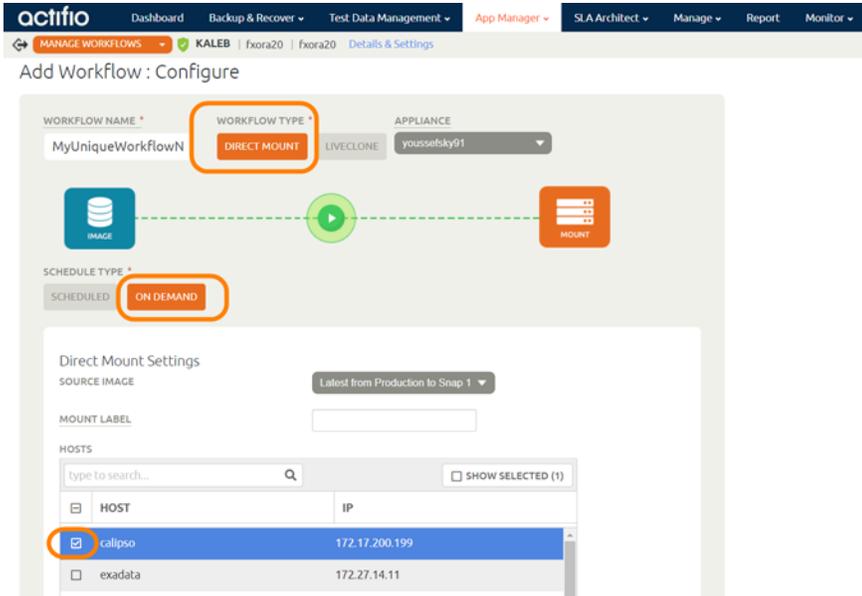
TARGET_SCHEMA_NAME: Target database schema to be refreshed with Source Schema.

Note: TARGET and SOURCE Schema names can be the same or different.

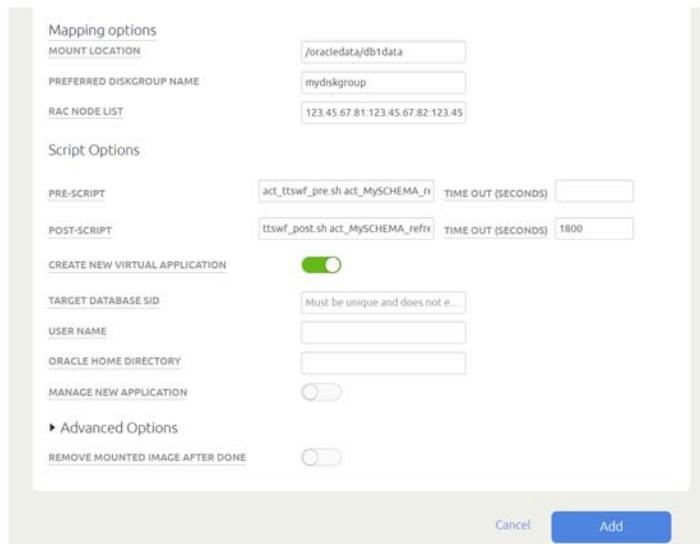
Creating the Workflow

Create a workflow to perform the schema refresh.

4. Open the AGM to App Manager > Workflows. In the upper right corner, select **+ Add Workflow**.
5. From the Add Workflow: Select an Application list, right-click the database to use and click **Next**.
6. The Add Workflow: Configure page appears:
 - o Provide a unique name for the workflow. Workflow names cannot include special characters.
 - o Select **Direct Mount**.
 - o Select the Actifio Appliance that will run the workflow.
 - o For Schedule Type, select **On Demand**.
 - o Select the Source image, either latest snapshot or latest snap from production.
 - o At Mount Label, provide a unique identifier to help identify mounts from this workflow when viewing active mounts.
 - o Select the target Host from the Hosts list.



7. In Mapping Options
 - o Enter the location for mounted drives (i.e. M:, D:\testdb1, /oracledata/db1data), if the temporary copy is going to be on NON – ASM.
 - o Add the Preferred Diskgroup Name.
 - o For RAC Node List, enter a colon-separated list of IP addresses.
8. Fill in the pre-script field:
`act_ttswf_pre.sh act_<SCHEMA>_refresh.conf`
 Where act_<SCHEMA>_refresh.conf is the file created from [Step 3](#).
9. Fill in the post-script field:
`ttswf_post.sh act_<SCHEMA>_refresh.conf`
 Where act_<SCHEMA>_refresh.conf is the file created from [Step 3](#).
10. Enter a timeout value of 1800 or more seconds:



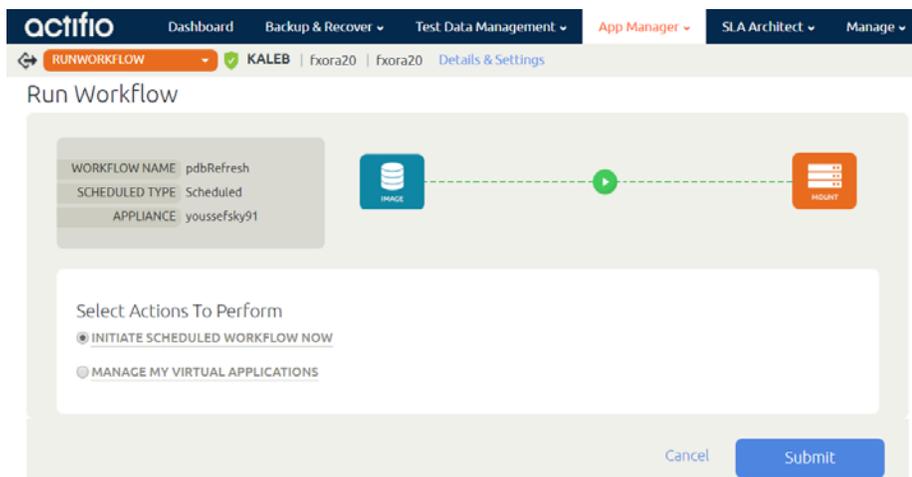
11. Select **Create New Virtual Application**.
12. Fill the target database SID (this is a temporary SID.)

13. Provide the OS Username who owns the Oracle software binary on the target machine.
14. Provide ORACLE_HOME as specified in the configuration file.
15. Open the Advanced Options by clicking the arrow. Provide the TNS ADMIN Directory Path on the target machine, and specify the minimum SGA% for the temporary instance. You can click the field labels for help.
16. Select **Standalone Non-RAC**.
17. Click **Add**.

Running the Workflow

Once the workflow is created:

1. Right-click the workflow and click **Run Now**.
2. Select **Initiate Scheduled Workflow Now** and click **Submit**.



Logs can be monitored on the target host at location: `/var/act/log`

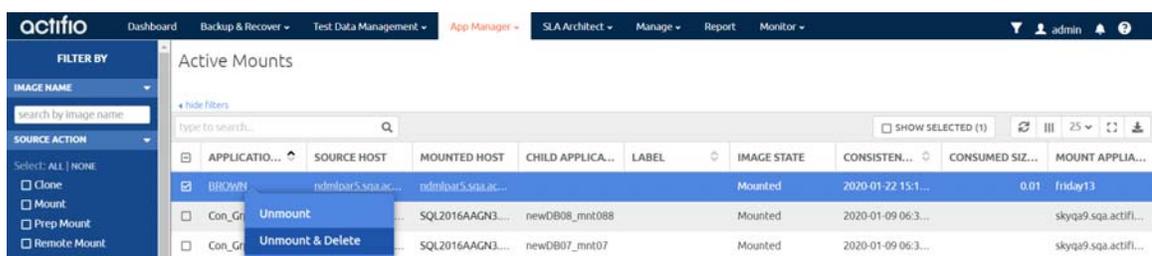
3. Once the workflow turns to success state, the schema on the target database is refreshed.
 - a. If the TARGET_SCHEMA_NAME does not exist on the target database, VDP will create a user for the first time and default user password is abc#1234.
 - b. Tablespaces for the refreshed schema on target database get presented to the target with the schema name as a prefix: `<Schema>_<Source_Tablespace_Name>`

Unmounting Mounted Images

To unmount mounted images:

1. Drop the tablespace and datafiles on the target database.


```
SQL>alter tablespace <TBS_NAME> offline immediate;
SQL>drop tablespace <TBS_NAME> including contents and datafiles;
```
2. Open App Manager > Active Mounts page. Right-click the image to be unmounted and select **Unmount and Delete** as shown below and then **Submit** the job.



Presenting an Oracle 12c Database PDB as a Virtual PDB to an Existing Database Container on a Target

If you are using an Oracle version earlier than 12c, then refer to [Using an Actifio Workflow to Refresh Oracle Database Schemas](#) on page 119.

Suppose you have a source database instance with multiple applications, and each application has its own PDB. You can create and refresh a virtual copy at the schema level, and refresh each PDB individually to the same target or to a different target. To do this:

[Before You Begin](#) on page 123

[Creating a Workflow to Perform the PDB Clone Job](#) on page 123

[Running the Workflow](#) on page 125

[Unmounting Mounted Images](#) on page 125

Before You Begin

Before you begin, set up the scripts on the target server:

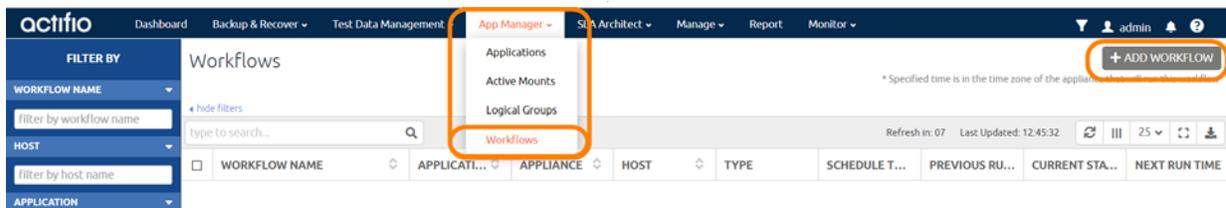
1. Get the script zip file from `/act/act_scripts/pdbrefresh`.
2. Login to database server as root.
3. Go to `/act` and create an `/act/scripts` directory (if it does not exist):

```
cd /act
mkdir -p /act/scripts
cd /act/scripts
```
4. Unzip the file.
5. Copy the script files into `/act/scripts`:

```
cp /act/act_scripts/pdbrefresh/act_pdbclone_pre.sh /act/scripts/
cp /act/act_scripts/pdbrefresh/act_pdbclone_post.sh /act/scripts/
```

Creating a Workflow to Perform the PDB Clone Job

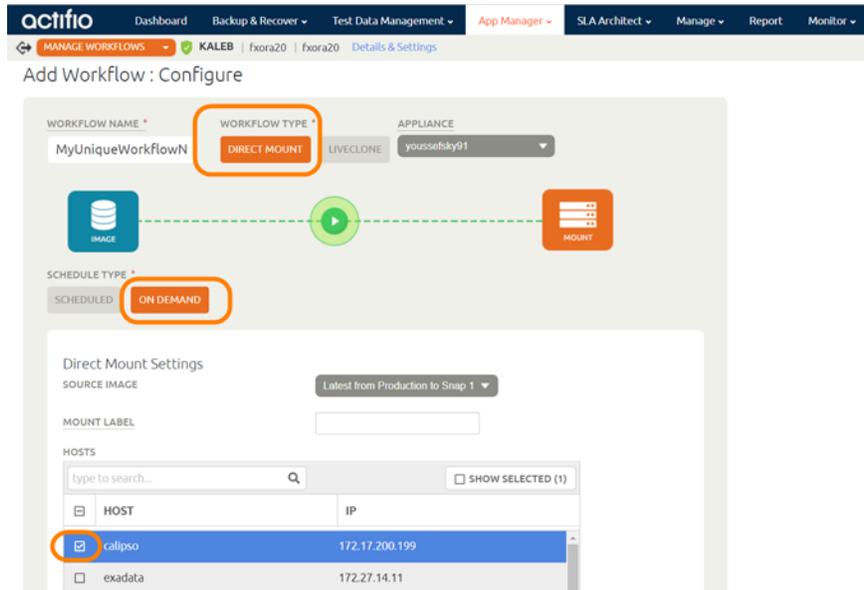
1. From the App Manager Workflows page, click **+ Add Workflow**.



2. Right-click the Oracle database and select **Next**.
3. Enter a name for the Workflow and select **Direct Mount** and **On Demand**.

Note: Workflow names cannot include special characters.

4. Select the target host by checking its checkbox.

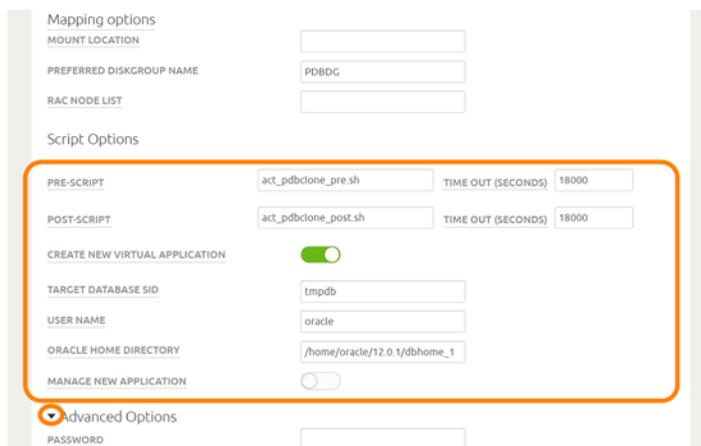


5. If the temporary copy is going to be on non-ASM, then under Mapping Options, select Specify Mount Location as the Mount Action and then provide a mount location.
6. Fill in the pre-script field and the post script field:

```
act_pdbclone_pre.sh "<TARGET_DB_SID>_<SOURCE_PDB_NAME>-<TARGET_PDB_NAME>"
```

```
act_pdbclone_post.sh "<TARGET_DB_SID>_<SOURCE_PDB_NAME>-<TARGET_PDB_NAME>"
```

 Where:
 TARGET_DB_SID = Target database SID where PDB should be attached
 SOURCE_PDB_NAME = Source PDB Name that needs to be attached to target container.
 TARGET_PDB_NAME = Target PDB Name to which source will be renamed.
7. Set **Timeout** for both scripts to 18000.
8. Select **Create New Virtual Application**.
9. Fill in the target database SID (this is temporary SID as defined in the act_pdb_config.conf file.)
10. Provide the OS username who owns the Oracle software binary on the target machine.
11. Provide ORACLE_HOME as specified in the configuration file.
12. Open the Advanced Options by clicking the arrow. Provide the TNS ADMIN Directory on the target machine and specify the minimum SGA for the temporary instance.

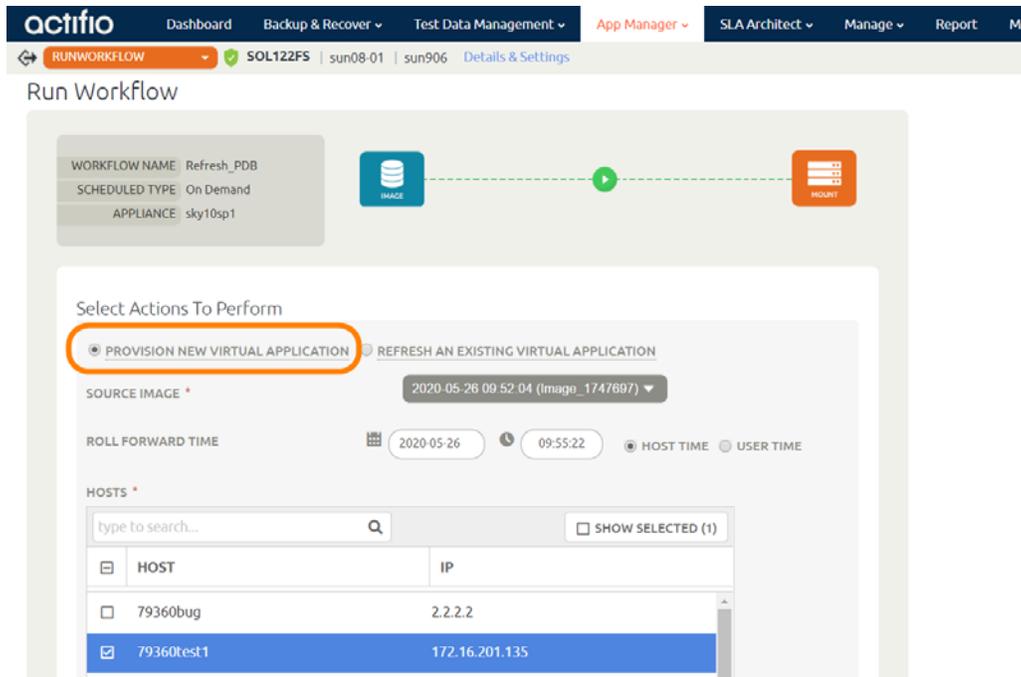


13. Select **Standalone Non-RAC**.
14. Click **Add** at the bottom of the page.

Running the Workflow

Once the workflow is created:

1. From App Manager, Workflows, right-click the workflow and click **Run Now**.
2. Select **Provision New Virtual Application**, fill in the details, and click **Submit**.



3. Logs can be monitored on the target host at: `/var/act/log`.
Once the workflow reaches success state, the PDB on the target database will be cloned.
If the target PDB must be refreshed again with latest or old source data, click on the Workflow **Run Now** button and select **Refresh Existing Virtual Application** and click **Done**.

Unmounting Mounted Images

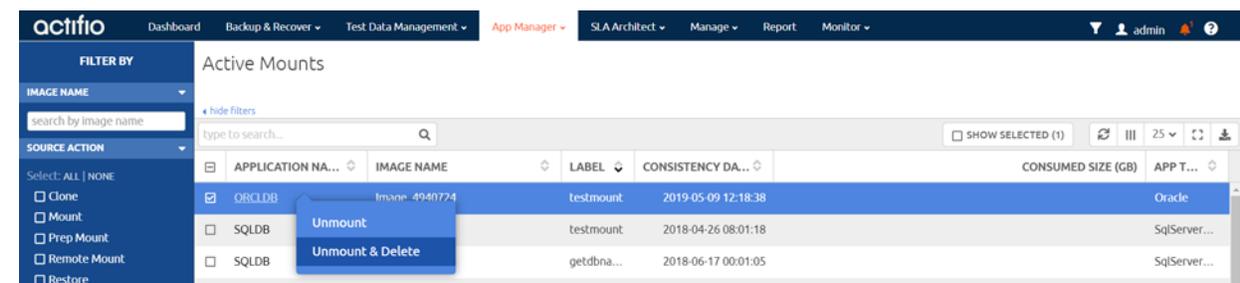
To unmount any mounted images:

1. From the Restore tab, select the image to unmount.
2. Drop the pluggable database and datafiles on the target database.
SQL>alter pluggable database <PDB_NAME> close immediate;
SQL>drop pluggable database <PDB_NAME> including datafiles;

Or, execute the Pre script:

```
cd /act/scripts  
./ act_pdbclone_pre.sh
```

3. From the App Manager **Active Mounts** list, right-click the image to be unmounted and select **Unmount and Delete** and then **Submit** the job.



21 Workflow Pre and Post Scripts

Workflows mount and unmount captured images according to a schedule or on demand. In a Workflow you can call:

- A pre script that runs before an image is mounted and/or unmounted
- A post script that runs after an image is mounted and/or unmounted

The ability to run a script before and after data is mounted or unmounted allows you to:

- Scrub for sensitive information
- Generate reports
- Warehouse data, especially when extracting, transforming, and loading (ETL) is required
- Copy to removable media

Scripts must reside in a folder named `/act/scripts` on the server that hosts the mounted Workflow image.

Note: For Workflows that involve both a database and its logs, you must configure scripts in such a way that pre and post script operations are not applied to the database's logs. See [Example Script](#) on page 128 for a script example that contains a command that will skip a database's log.

The following sections include:

- [Environment Variables](#) on page 127
- [Example Script](#) on page 128

Environment Variables

Environment variables allow you to invoke commands that apply to specific jobs, job types, or applications. Environment variables are prefixed with `ACT_`. For example, an environment variable for a database could look like: [`$ACT_APPNAME == "productiondb"`] or an environment variable for a mount operation could look like: [`$ACT_JOBTYPE == "mount"`]

The following is a list of common environment variables with sample values in parentheses.

- `JOBNAME`: The name of the job (e.g. Job_0123456)
- `APPID`: The ID of the application (e.g. 4186)
- `APPNAME`: Name of the application (e.g. My-DB)
- `HOSTNAME`: The name of the host which is the target of this job (e.g. Jupiter)
- `SOURCEHOST`: The name of the host that was the source for this application (e.g. Saturn)
- `JOBTYPE`: a text version of the job class (e.g. mount, unmount)
- `PHASE`: A text string that describes the job phase (e.g. pre, post)
- `TIMEOUT`: Define the duration of the script, how long the script is allowed to run
- `OPTIONS`: Policy options that apply to this job

Example Script

The following script example uses three environment variables:

- **ACT_JOBTYPE** - Identifies whether the job is a mount or unmount operation
- **ACT_PHASE** - Identifies whether the phase is either pre or post
- **ACT_MULTI_END** - Used only if both a database and its log are mounted. When this is "true" the database is in a state where it can be accessed

Note: The following example script is not meant to be used in a production environment, has not been tested, and is not warranted by Actifio.

```
#!/bin/sh
set +x
echo "*** Running user script: Job - $ACT_JOBNAME Type - $ACT_JOBTYPE Phase - $ACT_PHASE***"

# Use the first if clause to perform application specific operations during mount and in this example scrub-mount operation.

# Use the second if clause to perform any application specific operation during unmount and in this example, #scrub-unmount operation.

if [[ $ACT_JOBTYPE == "mount" ]] || [[ $ACT_JOBTYPE == "scrub-mount" ]]; then
if [[ $ACT_JOBTYPE == "unmount" ]] || [[ $ACT_JOBTYPE == "scrub-unmount" ]]; then
    echo "NO-OP for job type $ACT_JOBTYPE"
    exit 0
fi
fi

# Use the first if clause to perform application specific operations during the pre phase

# Use the second if clause to perform application specific operations during the post phase

if [[ $ACT_PHASE == "post" ]]; then
if [[ $ACT_PHASE == "pre" ]]; then
    echo "NO-OP for phase $ACT_PHASE"
    exit 0
fi
fi

# For multi-phase jobs (database and logs) check if the database has been mounted and the logs applied then #skip logs.

# If the operation needs to be performed in phases other than the last phase, modify the clause

if [[ -z "$ACT_MULTI_END" ]] && [[ $ACT_MULTI_END != "true" ]]; then
    echo "NO-OP for multi-phase operation"
    exit 0
fi

cd /act/scripts

echo "**** Running application specific logic: Job - $ACT_JOBNAME Type - $ACT_JOBTYPE Phase - $ACT_PHASE *"

# Any application specific commands will go here

echo "*** Finished running application specific logic : Job - $ACT_JOBNAME Type - $ACT_JOBTYPE Phase - $ACT_PHASE *"

exit $?
```

22 Oracle Database Management Using actDBM

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

You can also use the AGM version, detailed in [Chapter 23, Oracle Database Management Using ActAGMDBM](#).

This section includes:

[Installing and Configuring actDBM.pl on page 130](#)

[ActDBM Commands on page 132](#)

- o [listImageDetails](#)
- o [listApplication](#)
- o [listDiscoveredHost](#)
- o [backup](#)
- o [restore](#)
- o [clone](#) (both [Create a New Clone under Oracle ASM](#) and [Refresh a Clone](#))
- o [mount](#)
- o [cleanup](#) ([Unmount and Delete an Image](#))
- o [runwf](#) (both [Direct Mount Workflow](#) and [LiveClone Workflow](#))
- o [createliveclone](#)
- o [refreshliveclone](#)
- o [restoreASMswitch](#) ([Instant Oracle Database Recovery](#))
- o [restoreASMrebalance](#)

[actDBM.pl Script Template on page 150](#)

[Perl Examples of actDBM Usage and Results on page 154](#)

[RESTful API Examples of actDBM Usage and Results on page 155](#)

Installing and Configuring actDBM.pl

There are four steps to installing and configuring actDBM.pl:

[Installing actDBM.pl with the Actifio Connector](#) on page 130

[Installing and Verifying the Five Required Perl Modules](#) on page 130

[Enabling and Verifying Port 443](#) on page 130

[Blocking Main-in-the-Middle Attacks](#) on page 130

[Storing the Login Credentials for an Actifio Appliance \(CDSconfig\)](#) on page 131

Installing actDBM.pl with the Actifio Connector

When you install the Actifio Connector, the actDBM Perl script library is automatically installed on the Oracle host. To install the Actifio Connector, see **Network Administrator's Guide to Actifio GO**. It is available on the host under `/act/act_scripts/actdbm`.

The main script is `actDBM.pl`. You can see examples of actDBM usage and output in [RESTful API Examples of actDBM Usage and Results](#). If you want to create your own shell scripts to create a workflow to run backup, refresh clone, and other operations, follow the script template at [actDBM.pl Script Template](#).

Installing and Verifying the Five Required Perl Modules

These five Perl modules must be installed on the server where the actDBM tool is deployed and configured. Follow the verify step first to see if each module is already installed. If it is not, then install it:

Table 1: Installing the Five Required Perl Modules on a Linux System

PERL Module	Verify	Install on Red Hat or CentOS Linux	Install on SUSE or OpenSUSE Linux
JSON	<code>perl -e "use JSON"</code>	<code>yum install perl-JSON</code>	<code>zypper install perl-JSON</code>
LWP::UserAgent	<code>perl -e "use LWP::UserAgent"</code>	<code>yum install perl-LWP-UserAgent</code>	<code>zypper install perl-LWP-UserAgent</code>
IO::Socket::SSL	<code>perl -e "use IO::Socket::SSL"</code>	<code>yum install perl-IO-Socket-SSL</code>	<code>zypper install perl-IO-Socket-SSL</code>
Net::SSLeay	<code>perl -e "use Net::SSLeay"</code>	<code>yum install perl-Net-SSLeay</code>	<code>zypper install perl-Net-SSLeay</code>
LWP::Protocol::https	<code>perl -e "use LWP::Protocol::https"</code>	<code>yum install perl-LWP-Protocol-https</code>	<code>zypper install perl-LWP-Protocol-https</code>

Enabling and Verifying Port 443

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

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

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

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

Note: The escape character is `'\'`.

Blocking Main-in-the-Middle Attacks

If you are using CentOS with the installed `IO::Socket::SSL` version `> 1.78`, you may see a warning message:

```
Using the default of SSL_verify_mode of SSL_VERIFY_NONE for client is deprecated!
```

This is for security. Set `SSL_verify_mode` to `SSL_VERIFY_PEER` in your `SSL.pm` file:

1. Open `SSL.pm (/usr/share/perl5/vendor_perl/IO/Socket/SSL.pm)` file:
`vi /usr/share/perl5/vendor_perl/IO/Socket/SSL.pm`
2. Search for the following string/pattern in the file. There will be only one occurrence of this pattern in the file:
`SSL_verify_mode => SSL_VERIFY_NONE,`
3. Replace this line with the following line.
`SSL_verify_mode => SSL_VERIFY_PEER,`
4. Save the file after replacement. Run `actDBM.pl` after saving the file.

Storing the Login Credentials for an Actifio Appliance (CDSconfig)

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

```
perl actDBM.pl -type cdsconfig
--username <username>
--password <password>
--CDS <appliance ip>
```

Parameter	Use
<code>--username</code>	Actifio username to access the appliance. This is a required parameter.
<code>--password</code>	Password to access the appliance. This is a required parameter.

Repeat as needed for each affected Actifio Appliance. After storing the appliance details, you can run `actDBM.pl` from the Oracle host to the configured Actifio Appliance.

Running actDBM.pl

To run the `actDBM` tool, CD to `/act/act_scripts/actdbm` folder and invoke `./actDBM.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/act_scripts/actdbm/ /act/act_scripts/actdbm/actDBM.pl
```

Usage of actdbm.pl

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

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

The type options for `actDBM.pl` are:

```
[root@zoravmn1 actdbm]# perl actDBM.pl
Usage: actDBM
--type
    <backup>
    <restore>
    <clone>
    <mount>
    <cleanup>
    <listImageDetails>
    <runwf>
    <createliveclone>
    <refreshliveclone>
    <restoreASMswitch>
    <restoreASMrebalance>
    <cdsconfig>
```

ActDBM Commands

ActDBM.pl permits these operations:

To	See
Return a list of snapshot images with recovery range for the protected database application	listImageDetails on page 133
Return a list of applications on a host	listApplication on page 134
Return a list of hosts known to an Actifio Appliance	listDiscoveredHost on page 134
Back up a protected (source) database using traditional RMAN option	backup on page 135
Restore and recover the protected (source) database using traditional RMAN option	restore on page 136
Create or Refresh a copy of database with an Application Aware mount provisioning option depending on optype	clone on page 137
Mount the backup image to specified target server	mount on page 142
Stop and remove a copy of a database running out of a mounted image and remove the file system mount or drop an ASM disk group as part of cleanup	cleanup (Unmount and Delete an Image) on page 143
Run the pre-created DirectMountWorkflow or LiveCloneWorkflow based on --subtype option of runwf	runwf on page 144
Create a new LiveClone using the source image	createliveclone on page 146
Refresh an existing LiveClone using the source image	refreshliveclone on page 147
Instantly switch and start the database to Actifio mounted ASM Disk Group. The database will be up and running in its original configuration on your Actifio storage	restoreASMswitch (Instant Oracle Database Recovery) on page 148
Migrate the database back to production storage or to new storage for a switched database running out of Actifio image	restoreASMrebalance on page 149

listImageDetails

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

Example

```
perl actDBM.pl --type <listImageDetails>
--dbname <database name>
[--hostname <hostname>]
[--CDS <appliance name|IP>]
[--wait <yes|no>]
--remoteApp <no|yes>
```

Table 2: listImageDetails Parameters

Parameters	Use
<code>--dbname</code>	Name of database. This is a required parameter. The default value is <code>localHostname</code> .
<code>--hostname</code>	Name of source host. If not specified, host where script runs is used.
<code>--CDS</code>	The name or IP address of the Actifio Appliance
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".
<code>--remoteApp</code>	This is a flag to indicate whether the application data is to be retrieved from a local appliance or from a remote appliance. The values are <code>no</code> and <code>yes</code> . The default value is "no".

For a RESTful example of this command and sample output, see [Listing Images Without Archive Logs](#) and [Listing Images With Archive Logs Collected Between Oracle RMAN L1 Incremental Snapshots](#) on page 155.

listApplication

To return a list of applications on a host, use `--type listApplication`

Example

```
perl actDBM.pl --type <listApplication>
--hostname <hostname>
[--apptype <application type Oracle>]
[--protection <protectionState all/protected/disabled/notProtected>]
--CDS <CDS name|ip>
[--wait <yes|no>]
```

Table 3: listApplication Parameters

Parameters	Use
<code>--hostname</code>	Name of the source host. If not specified, host where script runs is used.
<code>--apptype</code>	Oracle
<code>--protection</code>	One of four protection states: all, protected, disabled, or notProtected
<code>--CDS</code>	The name or IP address of the Actifio Appliance
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".

For a Perl example of this command and sample output, see [Listing Applications Discovered on an Actifio Appliance](#) on page 154.

listDiscoveredHost

To return a list of discovered hosts known to a specific Actifio Appliance, use `--type listDiscoveredHost`

Example

```
perl actDBM.pl --type <listDiscoveredHost>
--CDS <CDS name|ip>
[--wait <yes|no>]
```

Table 4: listDiscoveredHost Parameters

Parameters	Use
<code>--CDS</code>	The name or IP address of the Actifio Appliance
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".

For a Perl example of this command and sample output, see [Listing Hosts Connected to an Actifio Appliance](#) on page 154.

backup

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

- db backup
- log backup
- dblog backup

Example

```
perl actDBM.pl --type backup
--dbname <database name>
[--hostname <hostname>]
[--backuptype <db|log|dblog>]
[--CDS <appliance name|ip>]
[--backuplabel <backup bookmark label name>]
[--wait <yes|no>]
```

Table 5: backup Parameters

Parameter	Use
--dbname	Name of database. This is a required parameter. The default value is localhostname.
--hostname (optional)	Name of database host. If not specified, host where script is running is used.
--backupType	Type of backup operation. This is an optional parameter. If not specified, the default type is database backup (db).
--backuplabel	Optional parameter. Use this to mark a backup image during backup.
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job is completed. The default value is No.

For a RESTful example of this command and sample output, see [Running an On-Demand Backup](#) on page 156.

restore

To restore a database, use `--type restore`

Example

```
perl actDBM.pl --type restore
--dbname <source database name>
--OSusername <OS oracle owner>
[--OSpassword <OS oracle owner password>]
[--image <Image name>]
[--backuplabel <backup label ID>]
[--sourceHost <source hostname>]
[--pointInTime <Recovery time 'yyyy-mm-dd hh24:mi:ss'>]
[--CDS <appliance ip|name>]
[--wait <yes|no>]
[--openDB <true|false>]
```

Table 6: restore Parameter

Parameter	Use
<code>--dbname</code>	Source database name
<code>--OSusername</code>	OS Oracle owner
<code>--OSpassword</code>	OS Oracle owner password. Optional, for backward compatibility with 6.2.x
<code>--image</code>	Image name. Optional. If not provided will use backuplabel to get the image for recovery. If both are not provided then the latest image will be used.
<code>--backuplabel</code>	Backup label ID. Optional. If image name is not provided will use backuplabel to get the image for recovery. If both are not provided then the latest image will be used.
<code>--sourceHost</code>	Source hostname
<code>--pointInTime</code>	Recovery time in the format yyyy-mm-dd hh24:mi:ss. Default: It will recover to all available archive logs.
<code>--CDS</code>	Actifio Appliance IP or name
<code>--wait</code>	Yes or No. The default value is "no".
<code>--openDB</code>	True or False. Default: True (recover and open the database at the end of this operation.)

clone

This creates or refreshes a copy of the database with an Application Aware mount provisioning option depending on which option you select:

`newclone` creates a new copy of the database to the specified target server. Set `optype` to `newclone` to create a new copy of the database to the specified target server. This will error out if the same copy of database already exists on specified target. For details, see [Create a New Clone under Oracle ASM](#) or [Create a New Clone under File System](#).

`refresh` refreshes the exiting clone copy of database created using New Clone option. Set `optype` to `refresh` to refresh the existing clone copy of a database created using the `newclone` option. This will error out if there is no copy of database to refresh on target. For details, see [Refresh a Clone](#) on page 141.

```
actDBM.pl --type clone
[--optype <newclone|refresh>]
```

Create a New Clone under Oracle ASM

To create a new clone under Oracle ASM, use `--type clone, optype <newclone>`

Example

```
perl actDBM.pl --type clone --optype newclone
[--optype <newclone>]
--dbname <source database name>
--TargetDBname <app-aware mount database name>
--OSusername <OS oracle owner>
[--image <Image name>]
[--sourceHost <source hostname>]
[--targetHost <target hostname>]
[--TargetOracleHome <Target Oracle Home>]
[--listenerpath <Target Oracle Listener path>]
[--pointInTime <Recovery time 'yyyy-mm-dd hh24:mi:ss'>]
[--ASMDiskgroup <ASM diskgroupname>]
[--RACNodeList <list of RAC nodes separated by colon>]
[--protectNewApplication <yes|no>]
[--totalmemory <Database memory size in MB>]
[--SGApct <SGA %>]
[--DBcacheSize <DB Cache Size in MB>]
[--SharedPoolSize <Shared Pool Size in MB>]
[--DBrecoveryFileDestSize <DB Recovery File Dest Size in MB>]
[--Processes <Max Number of Processes>]
[--OpenCursors <Max Number of Cursors>]
[--DiagnosticDestination <Diagnostic Destination>]
[--DBCharacterSet <Database Character Set>]
[--tnsip <TNS Listener IP/scan IP>]
[--tnsport <TNS listener port>]
[--tnsdomain <TNS Domain>]
[--redosize <redo size in MB>]
[--recovery <yes|no>]
[--envvar <post script env:post script db type>]
[--backuplabel <backup label ID>]
[--standalone <yes|no>]
[--template <child db protect template name>]
[--profile <child db protect profile name>]
[--CDS <CDS ip|name>]
[--remoteApp <no|yes>]
[--wait <yes|no>]
```

Create a New Clone under File System

To create a new clone under a file system, use `--type clone, optype <newclone>`

Example

```
perl actDBM.pl --type clone --optype newclone
[--optype <newclone>]
--dbname <source database name>
--TargetDBname <app-aware mount database name>
--OSusername <OS oracle owner>
[--image <Image name>]
[--sourceHost <source hostname>]
[--targetHost <target hostname>]
[--TargetOracleHome <Target Oracle Home>]
[--listenerpath <Target Oracle Listener path>]
[--pointInTime <Recovery time 'yyyy-mm-dd hh24:mi:ss'>]
[--mountpoint <filesystem mountpoint '/act/mnt'>]
[--protectNewApplication <yes|no>]
[--totalmemory <Database memory size in MB>]
[--SGApct <SGA %>]
[--DBcacheSize <DB Cache Size in MB>]
[--SharedPoolSize <Shared Pool Size in MB>]
[--DBrecoveryFileDestSize <DB Recovery File Dest Size in MB>]
[--Processes <Max Number of Processes>]
[--OpenCursors <Max Number of Cursors>]
[--DiagnosticDestination <Diagnostic Destination>]
[--DBCharacterSet <Database Character Set>]
[--tnsip <TNS Listener IP/scan IP>]
[--tnsport <TNS listener port>]
[--tnsdomain <TNS Domain>]
[--redosize <redo size in MB>]
[--recovery <yes|no>]
[--envvar <post script env:post script db type>]
[--backuplabel <backup label ID>]
[--standalone <yes|no>]
[--template <child db protect template name>]
[--profile <child db protect profile name>]
[--CDS <CDS ip|name>]
[--remoteApp <no|yes>]
[--wait <yes|no>]
```

Table 7: clone Parameters

Parameter	Use
<code>--optype</code>	Type of clone operation. This is optional parameter. If not specified, then a newclone operation will be performed.
<code>--dbname</code>	Name of database. This is a required parameter. The default value is localHostname.
<code>--TargetDBname</code>	Application aware mount clone database name.
<code>--OSusername</code>	OS Oracle owner name

Table 7: clone Parameters

Parameter	Use
--image	Image name to use for the database to be cloned. This is an optional parameter. If not provided, backuplabel will be used to get the image for recovery. If both are not provided then the latest image will be used.
--sourceHost	Name of source host. This is an optional parameter. If not specified, host where script is running is used.
--targetHost	Name of target host. This is an optional parameter. If not specified, host where script is running is used.
--TargetOracleHome	Oracle Home on target database Host
--listenerpath	Oracle listener path on target database host.
--pointInTime	Date and Time for clone the database. This is an optional parameter. If not specified, the database will be cloned with all available backup archive logs.
--mountpoint	Mountpoint name for filesystem application aware mountpoint name. This is not applicable for ASM application aware mount. Optional.
--protectNewApplication	Flag to indicate the new clone database is to be protected or not. The default value is No.
--ASMDiskgroup	Name of the mount ASM disk group name for clone
--RACNodeList	List of RAC node IP addresses separated by colon. This is a required parameter.
--totalmemory	The default value is the same as the memory on the source.
--SGApct	Percentage of Oracle System Global Area memory component. Optional.
--DBcacheSize	Oracle db_cache_size memory component in MB. Optional.
--SharedPoolSize	Oracle shared_pool_size memory parameter in MB. Optional.
--DBrecoveryFileDestSize	Oracle db_recovery_file_dest_size parameter value in MB. The default value is 51200MB (50 GB). Optional.
--Processes	Oracle processes parameter value. The default value is 500. Optional.
--OpenCursors	Oracle open_cursors init parameter value. If not specified, the value from the source database value will be used. Optional.
--DiagnosticDestination	Oracle diagnostic_dest parameter value. If not specified, Oracle will allocate automatically. Optional.
--DBCharacterSet	Oracle database character set parameter value.

Table 7: clone Parameters

Parameter	Use
--tnsip	TNS ip/scan/name for the target database.Default is target host IP address.
--tnsport	TNS port number for the target database. Default is 1521.
--tnsdomain	TNS domain name if any for the target database. Optional. By default, no domain name is used in creating service name entry under tnsnames.ora
--redosize	Oracle Redo log file size value. The default value is 1000 MB if not specified. Optional.
--recovery	Flag to open the database or leave it in a mounted state. Default is yes.
--envvar	Optional. Environment variables as arguments for running post script.
--backuplabel	Optional. If image name is not provided, then backuplabel is used to get the image for recovery. If both are not provided then the latest image is used.
--standalone	Flag to indicate if clone database is standalone or RAC (default) database.
--template	Child database protected template name. Required when protectNewApplication is set to yes.
--profile	Child database protected profile name. Required when protectNewApplication is set to yes.
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job got completed. The default value is No.
--remoteApp	This is a flag to indicate whether the application data is to be retrieved from a local appliance or from a remote appliance. The values are no and yes. The default value is "no".

For an example of this command and sample output, see [Cloning an Image](#) on page 156.

Refresh a Clone

To refresh a clone, use `--type clone, optype refresh`

Example

```
perl actDBM.pl --type clone --optype refresh
--dbname <source database name>
--TargetDBname <app-aware mount database name>
[--image <Image name>]
[--sourceHost <source hostname>]
[--targetHost <target hostname>]
[--pointInTime <Recovery time 'yyyy-mm-dd hh24:mi:ss'>]
[--backuplabel <backup label ID>]
[--CDS <CDS ip|name>]
[--wait <yes|no>]
```

Table 8: refresh Parameters

Parameter	Use
<code>--dbname</code>	Name of source database. This is a required parameter.
<code>--TargetDBname</code>	Application aware mount clone database name.
<code>--image</code>	Image name to use for refreshing the cloned database. This is an optional parameter. If not specified, the latest image will be used for refreshing the cloned database.
<code>--sourceHost</code>	Name of source host. This is an optional parameter. If not specified, host where script is running is used.
<code>--targetHost</code>	Name of target host. This is an optional parameter. If not specified, host where script is running is used.
<code>--pointInTime</code>	Date and Time to refresh the cloned database. This is an optional parameter. If not specified, the cloned database will be refreshed with all the available archive logs.
<code>--backuplabel</code>	Optional. If image name is not provided, then backuplabel is used to get the image for recovery. If both are not provided then the latest image is used.
<code>--CDS</code>	Actifio Appliance IP address
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".

mount

To mount a backup image, use `--type mount`

Example

```
perl actDBM.pl --type mount
--dbname <Database name>
[--image <Image name>]
[--sourceHost <Source host name>]
[--targetHost <Target host name>]
[--ASMDiskgroup <ASM diskgroup name>]
[--mountpoint <filesystem mountpoint '/act/mnt'>]
[--RACNodeList <RAC node list>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 9: mount Parameters

Parameter	Use
<code>--dbname</code>	Name of source database. This is a required parameter.
<code>--sourceHost</code> (optional)	Name of source host. If not specified, host where script is running is used.
<code>--targetHost</code> (optional)	Name of target host. If not specified, host where script is running is used.
<code>--image</code> (optional)	Image name to use for mount the database. If not specified, the latest image will be used for database mount.
<code>--ASMDiskgroup</code>	This parameter is for databases protected under ASM Disk Group only to provide a preferred Disk Group Name to mount a backup image under ASM.
<code>--RACNodeList</code>	This parameter is for databases protected under ASM Disk Group only to provide a list of node IP address separated by colon for RAC member nodes to mount a backup image under ASM to be available to the RAC members nodes listed here. This is a required parameter, even for mounting to a single node. The first IP address in RAC Node list must be the selected host's IP address.
<code>--mountpoint</code>	This is required for snapshot image backed up on filesystem. This parameter is for databases protected under File System only to provide a mount location to mount a backup image under provided file system mount location.
<code>--CDS</code>	Actifio Appliance IP address
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".

cleanup (Unmount and Delete an Image)

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 file system mount, or drop any ASM disk group as part of cleanup.

Example

```
perl actDBM.pl --type cleanup
--dbname <database name>
--clonedbname <target database name>
[--sourceHost <source host name>]
[--targetHost <target host name>]
[--CDS <CDS name|ip>]
[--wait <yes|no>]
```

Table 10: cleanup Parameters

Parameter	Use
<code>--dbname</code>	Name of source database for cleanup. This is a required parameter.
<code>--clonedbname</code>	Cloned database name or application aware mount database name which is created using Actifio.
<code>--sourceHost</code>	Optional. Name of source host. If not specified, host where script is running is used.
<code>--targetHost</code>	Optional. Name of target host. If not specified, host where script is running is used.
<code>--CDS</code>	Actifio Appliance IP address
<code>--wait</code>	Wait flag to wait until the job got completed. The default value is "no".

runwf

There are two types of workflows:

[Direct Mount Workflow](#) on page 144

[LiveClone Workflow](#) on page 145

Run Workflow runs the pre-created DirectMountWorkflow or LiveCloneWorkflow based on --subtype option, either directmount or liveclone. For an example of this command and sample output, see [Running an Oracle Workflow](#) on page 156.

Direct Mount Workflow

DirectMountWorkflow creates a new database copy or refreshes an existing database copy based on the re-provision option. To run a direct mount workflow, use --type runwf, subtype directmount

Example

```
perl actDBM.pl --type runwf
--subtype directmount
--dbname <source database name>
--hostname <sourcehostname>
--wfname <workflow name>
--reprovision <yes|no>
[--image <image name>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 11: directmount Parameters

Parameter	Use
--subtype	Type of workflow to be run, either directmount or liveclone.
--dbname	Name of source database. This is a required parameter.
--hostname (optional)	Name of source host. If not specified, host where script runs is used.
--wfname	Name of direct mount workflow
--reprovision	Reprovision flag to indicate new application aware mount or reprovision application aware mount.
--image	Image name to use for clone the database. This is an optional parameter. If not specified, the latest image will be used for database clone.
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job got completed. The default value is "no".

LiveClone Workflow

LiveCloneWorkflow first creates a new LiveClone or refreshes an existing LiveClone based on refreshliveclone or provisionliveclone option, and then creates or refreshes the target database copy based on the reprovision flag using the LiveClone image. To run a LiveClone workflow, use --type runwf, subtype liveclone

Example

```
perl actDBM.pl --type runwf
--subtype liveclone
--dbname <source database name>
--hostname <sourcehostname>
--wfname <workflow name>
--reprovision <yes|no>
[--image <Image name>]
[--refreshliveclone <yes|no>]
[--provisionliveclone <yes|no>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 12: liveclone Parameters

Parameter	Use
--subtype	Type of workflow to be run, either “directmount” or “liveclone”.
--dbname	Name of source database. This is a required parameter.
--hostname	Name of source host. If not specified, host where script runs is used. Optional
--wfname	Name of LiveClone workflow
--refreshliveclone	Refresh the LiveClone image for this workflow. The default is yes.
--provisionliveclone	Refresh the target copy of the database. The default is yes. If refreshliveclone is YES and provisionliveclone is YES, then running the workflow will first refresh the LiveClone image and then refresh the target copy of database using the refreshed image of the LiveClone. If refreshliveclone is YES and provisionliveclone is NO, then running the workflow will refresh only the LiveClone image. If refreshliveclone is NO and provisionliveclone is YES, then running the workflow will refresh only the target copy of database using the existing LiveClone image.
--reprovision	Flag to indicate a new clone or to provision already existing clone database. This MUST be set to NO the first time to create a copy of a database on the target. After the first time, set this to YES to refresh the copy of the database on the target.
--image	Image name to be used for LiveClone workflow. If not specified, then latest image will be used for LiveClone.
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job has completed. The default value is “no”.

createliveclone

To create a new LiveClone from a source image, use `--type createliveclone`

Example

```
perl actDBM.pl --type createliveclone
--dbname <database name>
--label <liveclone label>
--diskpool <disk pool name>
[--hostname <hostname>]
[--sourceimage <source image name>]
[--backuplabel <source backup image label>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 13: createliveclone Parameters

Parameter	Use
<code>--dbname</code>	Name of source database. This is a required parameter.
<code>--hostname</code> (optional)	Name of source host. If not specified, host where script runs is used.
<code>--label</code>	Name of LiveClone image.
<code>--diskpool</code>	Name of disk pool for LiveClone image.
<code>--sourceimage</code>	Name of source image to create a LiveClone image.
<code>--backuplabel</code>	Name of source image label instead of image name. If <i>neither</i> sourceimage nor backuplabel are provided, then the latest image will be used to create the LiveClone. If <i>both</i> sourceimage and backuplabel are provided then sourceimage is used and backuplabel is ignored.
<code>--CDS</code>	Actifio Appliance IP address
<code>--wait</code>	Wait flag to wait until the job has completed. The default value is "no".

refreshliveclone

To refresh an existing LiveClone from the source image, use `--type refreshliveclone`

Example

```
perl actDBM.pl --type refreshliveclone
--dbname <database name>
--label <liveclone label>
[--hostname <hostname>]
[--sourceimage <source image name>]
[--backuplabel <source backup image label>]
[--targetimage <target image name>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 14: refreshliveclone Parameters

Parameter	Use
<code>--dbname</code>	Name of source database. This is a required parameter.
<code>--hostname</code> (optional)	Name of source host. If not specified, host where script runs is used.
<code>--label</code>	Name of LiveClone image.
<code>--sourceimage</code>	Name of source image to refresh the LiveClone.
<code>--backuplabel</code>	Name of source image label instead of image name. If <i>neither</i> <code>sourceimage</code> nor <code>backuplabel</code> is provided then latest image will be used to create the liveclone. If <i>both</i> <code>sourceimage</code> and <code>backuplabel</code> are provided then <code>sourceimage</code> is used and <code>backuplabel</code> is ignored.
<code>--targetimage</code>	Target Image Name to be created for LiveClone image. Optional.
<code>--CDS</code>	Actifio Appliance IP address
<code>--wait</code>	Wait flag to wait until the job got completed. The default value is "no".

restoreASMswitch (Instant Oracle Database Recovery)

In case of storage failure, use this command to instantly switch to an Actifio-mounted ASM Disk Group and start the database. The database will be up and running in its original configuration on your Actifio storage.

To restore an ASM database using Oracle RMAN switch, use `--type restoreASMswitch`

Example

```
perl actDBM.pl --type restoreASMswitch
--dbname <database name>
--TargetOracleHome <Target Oracle Home>
--OSusername <OS oracle username>
--OSgridusername <OS grid username>
--standalone <yes/no>
--ASMDiskgroup <ASM diskgroup name>
--RACNodeList <RAC node list>
--sourceHost <source host name>
[--image <Image name>]
[--pointInTime <Recovery time 'yyyymmddhh24mi'>]
[--mountpoint <archivelog filesystem mountpoint '/act/mnt'>]
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 15: restoreASMswitch Parameters

Parameter	Use
--dbname	Name of source database. This is a required parameter.
--OSusername	OS oracle owner name
--OSgridusername	OS grid owner name
--pointInTime	Date and Time for clone the database. This is an optional parameter. If not specified, the database will be cloned with all the available backup archive logs.
--image	Image name to use for clone the database. This is an optional parameter. If not specified, the latest image will be used for database clone.
--TargetOracleHome	Oracle Home on target database host
--ASMDiskgroup	Name of the mount ASM disk group name for clone
--RACNodeList	List of RAC node IP addresses separated by colon. This is a required parameter.
--standalone	Standalone flag indicator to create target database as RAC or standalone.
--sourceHost	Name of source host. This is required parameter.
--mountpoint	Archivelog filesystem mountpoint
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job got completed. The default value is "no".

restoreASMrebalance

Use this option to migrate a database back to production storage or to new storage for a switched database running from an Actifio image.

To ASM rebalance use `--type restoreASMrebalance`

Example

```
perl actDBM.pl --type restoreASMrebalance
--dbname <Database name running on ASM diskgroup which is to be rebalanced>
--ASMDiskgroup <ASM diskgroup name>
--OSgridusername <OS grid username>
--OSusername <OS oracle username>
--ASMDiskString <ASM disk string separated by comma>
[--CDS <appliance name|ip>]
[--wait <yes|no>]
```

Table 16: restoreASMrebalance Parameters

Parameter	Use
--OSusername	OS Oracle owner name
--OSgridusername	OS grid owner name
--dbname	Database name running on the ASM disk group to be rebalanced
--ASMDiskgroup	Name of the mount ASM disk group name for clone
--ASMDiskString	ASM disk list of original disk group separated by comma
--CDS	Actifio Appliance IP address
--wait	Wait flag to wait until the job has completed. The default value is "no".

actDBM.pl Script Template

```
actDBM Script Template:
/*****
#!/bin/sh
set -x

# echo "create CDS config: One time only at for setting up the config"
# perl actDBM.pl --type cdsconfig --username admin --password password --CDS <CDS IP>

# for seting up the backup
echo "***** starting the backup *****"
perl actDBM.pl --type backup --dbname <source db name> --hostname <source host name
under CDS> --backuptype <db or log or dblog> --backuplabel <backup bookmark label> --
CDS <CDSIP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error backingup database *****"
exit 1;
fi
echo "***** end backup *****"

# for creating the first time clone when the database backup is under ASM disk group.
After clone is created use the refresh option syntex only.

# echo "***** start clone *****"

perl actDBM.pl --type clone --optype newclone --dbname <source database name> --
TargetDBname <target clone database name> --OSusername <oracle OS username> --
sourceHost <Source Hostname in CDS> --targetHost <Target Hostname in CDS> --image
<Source Image name> --backuplabel <source image backup label> --listenerpath <TNS
Listener path> --TargetOracleHome <Oracle Home on target server> --ASMDiskgroup
<Prefered ASM diskgroup name> --RACNodeList <RAC node list separated by colon(:)> --
mountpoint <filesystem mount location> --protectNewApplication <yes|no> - -totalmemory
<Memory in MB> --SGApct <SGA percentage> --DBcacheSize <DB Cache Size in MB> --
SharedPoolSize <Shared Pool Size in MB> --DBrecoveryFileDestSize <DB Recovery File
Dest Size in MB> --Processes <Max Number of Processes> --OpenCursors <Max Number of
Cursors> --DiagnosticDestination <Diagnostic Destination> --DBCharacterSet <Database
Character Set> --pointInTime <if want to clone to a specific time then specify this
('yyyy-mm-dd hh24:mi:ss')> --tnsip <TNS Listener IP> --tnsport <TNS listener port> --
tnsdomain <TNS Domain> --redosize <redo log size in MB> --recovery <yes|no> --
standalone <If the clone database is rac/standalone specify yes|no> --envvar
<environment variables> --template <child db protect template name> --profile <child
db protect profile name> --CDS <CDS IP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error clone/refresh database *****"
exit 1;
fi

echo "***** end clone *****"

echo "***** start refresh clone *****"

perl actDBM.pl --type clone --optype refresh --dbname <source database name> --
sourceHost <source Hostname in CDS> --targetHost <target Hostname in CDS> --
TargetDBname <cloned db name for refresh> --image <Image name> --backuplabel <backup
label ID> --pointInTime <if want to clone to a specific time then specify this ('yyyy-
mm-dd hh24:mi:ss')> --CDS <CDS IP> --wait <yes|no>
```

```

if [ $? -gt 0 ]; then
echo "***** error clone/refresh database *****"
exit 1;
fi
echo "***** refresh clone complete *****"

echo "***** start directmount workflow for clone *****"

perl actDBM.pl --type runwf --subtype directmount --dbname <source database name> --
hostname <source Hostname in CDS> --wfname <workflow name> --reprovision <reprovision
flag yes/no> --image <Image name> --CDS <CDS IP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error direct mount workflow *****"
exit 1;
fi
echo "***** end direct mount workflow *****"

echo "***** start liveclone workflow for new clone *****"

perl actDBM.pl --type runwf --subtype liveclone --dbname <source database name> --
hostname <source Hostname in CDS> --wfname <workflow name> --reprovision <reprovision
flag yes/no> --refreshliveclone yes --provisionliveclone yes --CDS <CDS IP> --wait
<yes|no>

if [ $? -gt 0 ]; then
echo "***** error live clone workflow *****"
exit 1;
fi
echo "***** end liveclone workflow *****"

echo "***** start createliveclone image *****"

perl actDBM.pl --type createliveclone --dbname <source database name> --label
<liveclone label> --diskpool <disk pool name> --sourceimage <source image name> --
hostname <source Hostname in CDS> --backuplabel <source backup image label> --CDS <CDS
name|ip> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running createliveclone image *****"
exit 1;
fi

echo "***** end createliveclone image *****"

echo "***** start refreshliveclone image *****"

perl actDBM.pl --type refreshliveclone --dbname <source database name> --label
<liveclone label> --hostname <source Hostname in CDS> --sourceimage <source image
name> --backuplabel <source backup image label> --targetimage <target image name> --
CDS <CDS name|ip> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running refreshliveclone image *****"
exit 1;
fi
echo "***** end refreshliveclone image *****"

echo "***** start list image details *****"

```

```

perl actDBM.pl --type listImageDetails --dbname <source database name> --hostname
<source Hostname in CDS> --CDS <CDS name|ip> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running list image details *****"
exit 1;
fi
echo "***** end list image details *****"

echo "***** start mount image *****"

perl actDBM.pl --type mount --dbname <source database name> --sourceHost <source
Hostname in CDS> --targetHost <target Hostname in CDS> --image <Image name> --
ASMDiskgroup <ASM diskgroup name> --RACNodeList <RAC node list> --mountpoint
<filesystem mountpoint> --CDS <CDS name|ip> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running mount image *****"
exit 1;
fi
echo "***** end mount image *****"

echo "***** start restore image *****"

perl actDBM.pl --type restore --dbname <source database name> --OSusername <oracle OS
username> --sourceHost <source Hostname in CDS> --OSpassword <OS oracle owner
password> --image <Image name> --backuplabel <source backup image label> --pointInTime
<if want to clone to a specific time then specify this ('yyy-mm-dd hh24:mi:ss')> --CDS
<CDS IP> --wait <yes|no> --openDB <flag to indicate the restored database need to be in
open or in mount state true|false>

if [ $? -gt 0 ]; then
echo "***** error running restore image *****"
exit 1;
fi
echo "***** end restore image *****"

echo "***** start unmount+delete image *****"

perl actDBM.pl --type cleanup --dbname <source database name> --clonedbname <target
database name> -- sourceHost <source Hostname in CDS> --targetHost <target Hostname in
CDS> --CDS <CDS IP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running unmount+delete image *****"
exit 1;
fi
echo "***** end unmount+delete image *****"

echo "***** start ASM restore database switch *****"

perl actDBM.pl --type restoreASMswitch --dbname <source database name> --sourceHost
<source Hostname in CDS> --TargetOracleHome <Oracle Home path for the database> --
OSusername <OS Oracle username> --OSgridusername <OS grid username> --standalone
<standalone flag yes|no> --image <Image name> --ASMDiskgroup <ASM diskgroup name> --
RACNodeList list of RAC nodes separated by colon> --pointInTime <if want to clone to a
specific time then specify this ('yyyymmddhh24mi')> --CDS <CDS IP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running ASM restore database switch *****"

```

```

exit 1;
fi
echo "***** end ASM restore database switch *****"

echo "***** start ASM restore rebalance *****"

perl actDBM.pl --type restoreASMrebalance --ASMDiskgroup <ASM diskgroup name> --
OSgridusername <OS grid username> --OSusername <OS Oracle username> --ASMDiskString
<ASM disk string seperated by comma 'ORCL:ASMDISK0'> --CDS <CDS IP> --wait <yes|no>

if [ $? -gt 0 ]; then
echo "***** error running ASM restore rebalance *****"
exit 1;
fi
echo "***** end ASM restore rebalance *****"

echo "***** completed all operation *****"
exit 0

/*****/

```

Perl Examples of actDBM Usage and Results

Listing Hosts Connected to an Actifio Appliance

```
perl actDBM.pl --type listDiscoveredHost --CDS 192.168.18.38
hostName: rac1node121
hostIP: 192.168.18.71
hostID: 6046

hostName: node1.rac
hostIP: 192.168.18.186
hostID: 12450

hostName: node2.rac
hostIP: 192.168.18.187
hostID: 12482

hostName: orafs.sreehari
hostIP: 192.168.18.185
hostID: 54845

hostName: WIN-OI5JAC5N5E8
hostIP: 192.168.18.231
hostID: 155712
```

Listing Applications Discovered on an Actifio Appliance

```
perl actDBM.pl --type listApplication --hostname rac1node121 --CDS 192.168.18.38 --
wait yes
rac1node121
  oracleApplicationName: targdb
  oracleApplicationID: 6583
  CDS: 192.168.18.38
  templateApplied: SreeHari_DB_Log
  protection: Protection Disabled

rac1node121
  oracleApplicationName: oracl
  oracleApplicationID: 92774
  CDS: 192.168.18.38
  templateApplied: SreeHari_DB_Log
  protection: Protection Disabled
```

RESTful API Examples of actDBM Usage and Results

These examples use:

- production host: **Oracle-Prod**
- Oracle databases running on Oracle-Prod: **bigdb** and **smalldb**

Listing Images Without Archive Logs

Listing images of bigdb on Oracle-Prod:

- To reduce the length of the output, only the first three images are listed.
- The Recover Start Time (ST) and End Time (ET) are the same as we are not collecting logs for this DB.

```
[ab@bb-oracle actdbm]# ./actDBM.pl --type listImageDetails --CDS 172.24.1.180 --dbname bigdb --hostname Oracle-Prod
Database Name: bigdb
Host Name: Oracle-Prod
Backup Image and Recovery Range:
Snapshot Pool Images:
Image_25785750:
RecoveryST: '2016-06-25 12:03:52' RecoveryET: '2016-06-25 12:03:52'
Image_25787410:
RecoveryST: '2016-06-26 00:03:47' RecoveryET: '2016-06-26 00:03:47'
Image_25790014:
RecoveryST: '2016-06-26 12:03:50' RecoveryET: '2016-06-26 12:03:50'
```

Listing Images With Archive Logs Collected Between Oracle RMAN LI Incremental Snapshots

Listing images of smalldb on Oracle-Prod:

- To reduce the length of the output, only the first three images are listed.
- The Recovery End Time and End Sequence number is the same for each image, because we make the full log range available for every image which includes OnVault images.

```
[ab@bb-oracle actdbm]# ./actDBM.pl --type listImageDetails --CDS 172.24.1.180 --hostname Oracle-Prod --dbname smalldb
Database Name: smalldb
Host Name: Oracle-Prod
Backup Image and Recovery Range:
Snapshot Pool Images:
Image_25839320:
RecoveryST: '2016-07-03 00:10:34' RecoveryET: '2016-07-04 20:09:38'
Thread1:          StartSequence: 13185          EndSequence: 13269
Image_25841040:
RecoveryST: '2016-07-03 12:11:09' RecoveryET: '2016-07-04 20:09:38'
Thread1:          StartSequence: 13207          EndSequence: 13269
Image_25842716:
RecoveryST: '2016-07-04 00:10:43' RecoveryET: '2016-07-04 20:09:38'
Thread1:          StartSequence: 13230          EndSequence: 13269

De-dup Images:
Image_25524341:
RecoveryST: '2016-06-12 12:08:11' RecoveryET: '2016-07-04 20:09:38'
Image_25734039:
RecoveryST: '2016-06-19 12:08:20' RecoveryET: '2016-07-04 20:09:38'
Image_25790261:
RecoveryST: '2016-06-26 12:12:32' RecoveryET: '2016-07-04 20:09:38'
```

Running an On-Demand Backup

Running an on-demand snapshot:

```
[ab@bb-oracle actdbm]# ./actDBM.pl --type backup --dbname bigdb --hostname Oracle-Prod
--CDS 172.24.1.180
policyid=8632
Job_25802477
[root@av-oracle actdbm]#
```

Cloning an Image

Cloning an image. We define a new SID for our mounted image (called **billed**) and a target host (**demo-oracle-4**). The end result is a new instance of a 2 TB Oracle database, available in less than three minutes, consuming no additional storage.

The command looks long, but most of it is just telling the Actifio Connector where to find ORACLE_HOME in case we need to use a different folder.

```
[ab@bb-oracle actdbm]# ./actDBM.pl --type clone --optype newclone --dbname bigdb --
sourceHost Oracle-Prod --TargetDBName billed --OSusername oracle --targetHost demo-
oracle-4 --TargetOracleHome /home/oracle/app/oracle/product/11.2.0/dbhome_1 --
listenerpath /home/oracle/app/oracle/product/11.2.0/dbhome_1/network/admin --
protectNewApplication no --recovery yes --CDS 172.24.1.180
backup array is HASH(0x27e7088)
backup name is Image_25802477
db name is bigdb
Image ID 25802479 name Image_25802477
image id is 25802479
image name is Image_25802477
Image details by name
Job_25802582 to mount Image_25802477 started
```

Now there is a new Oracle instance on target host demo-oracle-4:

```
[oracle@demo-oracle-4 ~]$ ps -ef | grep pmon
oracle 22319 1 0 21:41 ? 00:00:00 ora_pmon_billed
```

This command gives near-instant access to an off-production copy of a large production database with no need to buy or provision extra storage, engage the storage administrator, the backup administrator, or the DBA. Developers can get straight to the data.

Running an Oracle Workflow

We can re-run this command with `--optype refresh` to update our mounted copy.

You can embed nearly all of this in a workflow using the AGM and just call the workflow instead:

```
[ab@bb-oracle actdbm]# ./actDBM.pl --type runwf --subtype directmount --dbname bigdb -
-hostname Oracle-Prod --wfname EverWF --CDS 172.24.1.180
backup array is HASH(0x2db3000)
backup name is Image_25802477
db name is bigdb
Image ID 25802479 name Image_25802477
image id is 25802479
image name is Image_25802477
WF Id = 25763787
mount id: 25803103
database name is
clone work flow id: 25803199
```

Everything you have seen here is using RESTful API calls via a pre-built CLI. You could call these using your own scripts and integrate them into a portal.

23 Best Practices for Application Details & Settings

Actifio should manage both database/log backup and archive log purging.

Application Details & Settings Recommended Settings

- **STAGING DISK SIZE (GB):** Do not set this parameter. By default, the connector calculates the size as 1.5 times the maximum size of the database.

This is thin provisioned. The space will be consumed only when it gets written.
To change the default calculation set the policy option **Staging Disk Over-Allocation (in Percentage)** to desired value.
- **NUMBER OF CHANNELS:** Enter the number of RMAN channels based on the host computing power. Number of channels should be configured based on the number of cores available on the server, taking into account other database backups configured to run in parallel. The default number of channels is 1.
- **AU_SIZE:** Parameter to configure ASM Diskgroup AU size, in MB. Set this before the first snapshot (this only takes effect during diskgroup creation, during the first backup job). The recommended value is 4MB.
- **Do Not Uncatalog:** By default this is set to NO and Actifio datafile backup will be cataloged at the start of backup and then be uncataloged at the end of the backup. To keep RMAN datafile backup cataloged after each backup job, set it to YES, this will optimize the backup time for databases with a large number of datafiles.

Note: If Actifio GO must co-exist with other legacy backup products, then keep this set to NO; see [Chapter 2, Best Practices for Using Actifio GO with Other Backup Products](#).

- **Crosscheck Archivelog:** Select this to run crosscheck and delete expired archivelogs on archive backup.
- **Crosscheck Backup of Archivelog:** Select this to run crosscheck on the current backed up archivelog before the new logs are backed up, and delete expired archivelogs.
- **Number of Files per Backupset:** Specify the number of archivelogs to include in a backupset during archivelog backup. Recommended value is 4
- **Log Purging Retention Period:** In the space provided, enter the number of hours to retain archive logs in the primary log destination. For example, if this is set to 4, then archive logs older than four hours will be purged from the database primary archive destination. The default value is 24 hours.
- **Successful Log Backups Before Purge:** Recommended value is 1. By default, archive purging does not check for the number of successful log backups.

Staging disk size calculation

For the policy option called “Staging Disk Over-Allocation (in Percentage), the default value is 1.5.

If this value is set, it will be used instead of 1.5x, e.g. 60% will be 1.6x

```
If total_db_size*overhead < 50 -> 50GB
   total_db_size*overhead <100 -> 100GB
   total_db_size*overhead <200 -> 200GB
Else
   total_db_size*overhead
```