
An Oracle DBA's Guide to IBM InfoSphere Copy Data Management

Contents

Chapter 1 - Introduction to IBM InfoSphere Copy Data Management	1
Capturing Oracle Data	2
Replicating Data	3
Accessing Data	4
Known Oracle Bugs That Affect IBM InfoSphere Functionality	6
Chapter 2 - IBM InfoSphere Prerequisites for Protecting an Oracle Database	9
Prerequisites when Working from the IVGM	9
Chapter 3 - Oracle Authentication	11
Using Oracle with OS Authentication	11
Enabling Database Authentication for an Oracle Server from IVGM	12
Chapter 4 - Data Capture under File System and under ASM Disk Group	15
Protecting an Oracle Database Under a File System as a File System	16
Protecting an Oracle Database Under an ASM Disk Group as an ASM Disk Group	16
Protecting an Oracle Database Under a File System as an ASM Disk Group	18
Protecting an Oracle Database Under an ASM Disk Group as a File System	19
Chapter 5 - Preparing Oracle Databases for Protection	21
Preparing Oracle Databases in a Unix Environment Using OS Authentication	22
Preparing to Capture a Database from Oracle ASM to Oracle ASM	23
Preparing to Capture a Database from Oracle ASM to Filesystem	23
Preparing Oracle Database Authentication in a Unix Environment	24
Enable Database Block Change Tracking (optional)	26
Protecting from an Oracle Data Guard Node	26
Configuring RAC Transparent Failover of IBM InfoSphere RMAN Backup to Other Nodes	27
Oracle Archive Logs Compression	27
Manually Calculating Log Staging Disk Size (optional)	28
Configuring Oracle Database Services for Load Balancing across Multiple Nodes	29
Chapter 6 - Details and Settings for Oracle Databases in IVGM	31
SLA Policy Overrides for Oracle Databases	31
Application Details and Settings for Oracle Databases	33

Chapter 7 - Virtualizing an Oracle Database for Data Protection and Agility	37
Virtualizing an Oracle Database for Data Protection and Agility using the IVGM.....	37
Chapter 8 - Accessing, Recovering, or Restoring an Oracle Database via the IVGM	41
Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access from the IVGM.....	42
Mounting an Oracle Database Image Protected Under a File System for Data Access from the IVGM	45
Mounting an Oracle Database as a Virtual Application in the IVGM.....	47
Bringing IBM InfoSphere-Protected ASM Diskgroups Back Online after Reboot of a Target DB Server	51
Restoring a Database via IVGM, Overwriting the Production Database	52
Chapter 9 - Recovering an Oracle Database Manually Using RMAN	55
Recovering a Non-RAC Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere.....	56
Recovering a RAC ASM Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere.....	57
Recovering a Non-RAC Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog	59
Recovering a RAC ASM Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog	61
Recovering an Oracle Database to a Scheduled Backup Point if the archivelog is not Protected through IBM InfoSphere	63
Recovering an Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog	65
Chapter 10 - Recovering Tablespace and Data Files	67
Recovering a Single Tablespace of a Production Database on an ASM Disk Group	67
Recovering a Corrupt Database Block	68
Recovering Lost Control Files.....	69
Recovering an Oracle Pluggable Database.....	70
Chapter 11 - Performing an Oracle ASM Switch and Rebalance via IVGM	71
Chapter 12 - Protecting and Recovering Oracle Databases in a Windows Environment	77
Preparing Oracle Protection in a Windows Environment	77
Identifying Database Instances On Windows	78
Backing Up an Oracle Database in a Windows Environment	79
Watch Script to Watch for Database Volumes Being Mounted	80
Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere.....	81
Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog	82
Chapter 13 - Configurations for Using IBM InfoSphere in an Exadata Environment	83
Chapter 14 - Protecting SAP ECC/BW with an Oracle Database	85
Protecting the Oracle Database.....	85
Refreshing the Database.....	85
After the Refresh, on the Target Oracle Database.....	86

After the Refresh, on the Target SAP Application Server	86
Chapter 15 - Oracle RMAN Logs	87
Oracle Protection Logs on Linux/Unix	87
Oracle Protection Logs on Windows	88
Chapter 16 - Oracle Database Management Using actDBM	89
Installing and Configuring actDBM.pl	90
actDBM Commands	92
listImageDetails	93
listApplication	94
listDiscoveredHost	95
backup	96
restore	97
clone	98
Create a New Clone under Oracle ASM	98
Create a New Clone under File System	99
Refresh a Clone	102
mount	103
cleanup (Unmount and Delete an Image)	104
runwf	105
Direct Mount Workflow	105
LiveClone Workflow	106
createliveclone	107
refreshliveclone	108
restoreASMswitch (Instant Oracle Database Recovery)	109
restoreASMrebalance	110
actDBM.pl Script Template	111
Perl Examples of actDBM Usage and Results	115
RESTful API Examples of actDBM Usage and Results	116
Chapter 17 - Oracle Database Management Using actAGMDBM	119
Installing and Configuring actAGMDBM.pl	120
actAGMDBM Commands	121
listAppliance	121
listImageDetails	122
backup	123
cleanup	124
clone	125
workflow	126
actAGMDBM.pl Script Template	127
RESTful API Examples of actAGMDBM Usage and Results	128

1 Introduction to IBM InfoSphere Copy Data Management

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

[Capturing Oracle Data](#) on page 2

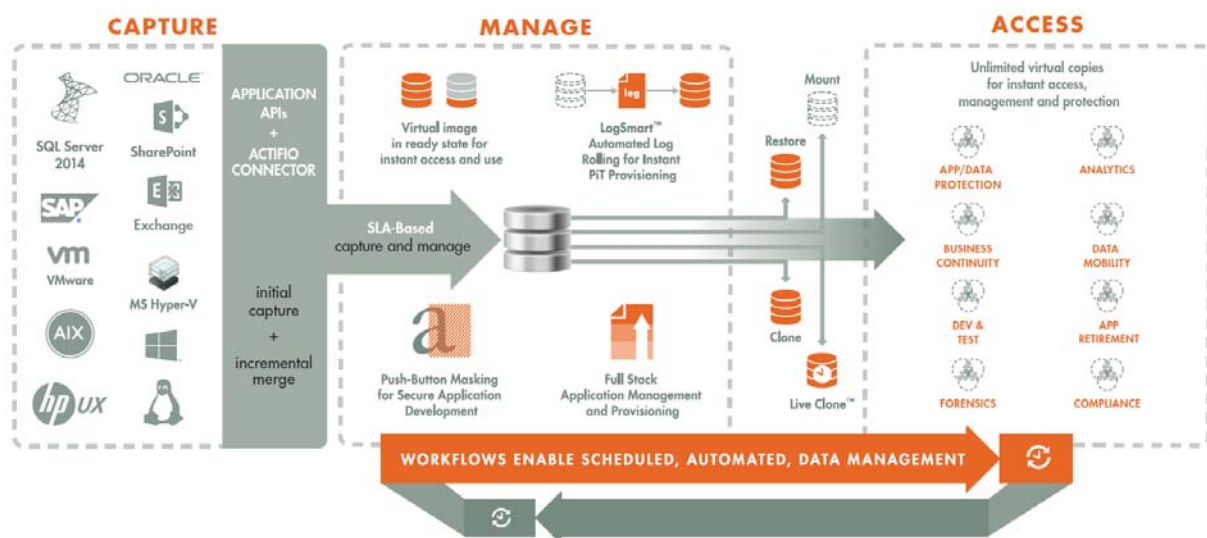
[Replicating Data](#) on page 3

[Accessing Data](#) on page 4

[Known Oracle Bugs That Affect IBM InfoSphere Functionality](#) on page 6

IBM InfoSphere Data Virtualization

An InfoSphere VDP 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 infrastructure.

Capturing Oracle Data

Capturing Oracle data consists of four simple steps:

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

The VDP Connector

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

Specifically, the VDP 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 IBM InfoSphere'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.

- IBM InfoSphere 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 IBM InfoSphere Consistency Group

A consistency group can contain only a single Oracle database application and any number of file system applications from the Oracle server. A consistency group is a very good choice if you are managing Oracle databases for test/dev and other business agility purposes.

Oracle Databases with TDE

IBM InfoSphere 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 TDE (Transparent Data Encryption). 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 requires the wallet to be copied to the appropriate location on the mount host.

Replicating Data

Data can be replicated to a second InfoSphere VDP 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. IBM InfoSphere replication addresses these issues with a global deduplication and compression approach that:

- Drives down overall network usage.
- Eliminates the need for a dedicated WAN accelerator/optimizer.
- Does not require storage array vendor licenses as data is sent from one InfoSphere VDP 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 InfoSphere VDP Appliances is performed using 1024-bit certificates.

Replication is controlled by IBM InfoSphere Policy Template policies. Production to Mirror policies have two options to replicate data to a second InfoSphere VDP Appliance:

- Dedup Backup to Dedup DR policies use a fixed, IBM InfoSphere proprietary replication engine to replicate data to a second InfoSphere VDP Appliance. In addition, Dedup Backup to Dedup DR policies allow you to replicate data to two locations.
- Production to OnVault policies use a fixed, IBM InfoSphere proprietary replication engine to replicate data to the cloud.

Accessing Data

The InfoSphere VDP Appliance can instantly present a copy of the database rolled forward to a specific point of time. The roll forward operation is performed from the IVGM.

Access options include:

[Mounts](#)

[LiveClones](#)

[Restores](#)

[Workflows](#)

Note: Cross platform presentation of Oracle images captured over NFS is not supported. For example, Oracle data captured from Solaris system cannot be presented on a Linux system.

Mounts

The IBM InfoSphere mount function provides instant access to data without moving data. Captured copies of databases can be rolled forward via the IBM InfoSphere user interface and mounted on any Oracle database server.

An InfoSphere VDP 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 InfoSphere VDP 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 in the IVGM](#) on page 47.

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

While SLAs govern the automated *capture* of a production Oracle database, Workflows automate *access* to the captured database.

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.

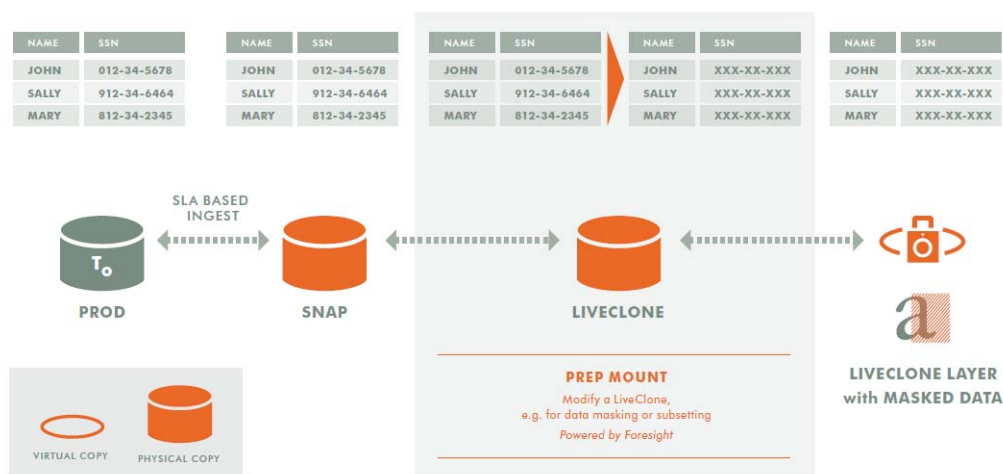
Combining IBM InfoSphere's automated Oracle data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now users can provision their own environments almost instantly.

For example, an IBM InfoSphere 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 the VDP Desktop.



Workflow With Masked Social Security Data

Known Oracle Bugs That Affect IBM InfoSphere Functionality

These are known Oracle bugs that can occur in an IBM InfoSphere environment:

Tracking 50661: Oracle Standalone ASM backups failing with error "Failed to create ASM staging disk group" :

Oracle Version Affected: 11.2.0.1

Oracle Standalone ASM backups failing with error "ORA-15238: 11.2.0.1.0 is not a valid value for attribute compatible.asm". Upgrade to 11.2.0.4 or higher.

Oracle Bug On IBM AIX on POWER Systems (64-bit) ORA-00449

Background process 'MMON' unexpectedly terminated with error 448 (Doc ID 2111305.1)

Oracle Version Affected 12.1.0.2 to 12.1.0.2

IBM InfoSphere Application Aware Mount may fail if the Patch 23089357 is not applied.

Apply Patch 23089357: APPSST12102:SHUTDOWN FAILS WITH ORA-449 "BACKGRND PROCESS "MMNL" UNEXPECTEDLY

Oracle Bug 13366202

DBNEWID [nid] does not allow TARGET=/ (NID-106) (Doc ID 13366202.8)

Oracle Version Affected: 11.2.0.3

IBM InfoSphere Application Aware mounts may fail if Oracle database 11.2.0.3 install does not include Patch 13366202.

Oracle Bug 8579188

RMAN RESTORE COMMAND FAILED WITH ORA-1861 (RMAN Recovery Session Fails with ORA-1861 (Doc ID 852723.1))

Oracle Version Affected: 10.2.0.1 to 11.2.0.1

IBM InfoSphere Oracle backup may fail for Oracle database 10.2.0.1 to 11.2.0.1.

Fixed in: 11.2.0.2

Oracle Bug 9216848

On AIX 6.1 64bit Power 7 Opatch Fails In \$ORACLE_HOME/jdk/jre/bin/libj9jit23.so (Unhandled Exception Type=Segmentation Error Vmstate=0x00000000 J9Generic_Signal_Number=...) (Doc ID 1307620.1)

Oracle Version Affected: 11.1.0.7

Apply Patch: 9216848

Oracle Bug 19621704

Oracle Database - Enterprise Edition - Version 11.2.0.1 to 12.1.0.2 [Release 11.2 to 12.1]

ORA-00600 [723] [memory Leak] Error With Leaked Memory For "mbr node memory" (Doc ID 1944511.1)

IBM InfoSphere Application Aware mounts may encounter a memory leak issue if patch 19621704 is not applied.

Apply Patch: 19621704

Oracle Database 12cR1 Bug 18845653

ORA-600 from PDB close if PDB renamed in another session [18845653.8]

Bug 19075256 - ORA-600 [kcfmis_internal: enq] from PDB RENAME [19075256.8]

This can happen on IBM InfoSphere created copy for test/dev if PDB rename operation is executed. Fixed in: April 2017 DB PSU

Oracle Database 12cR1 Bug 19404068

(ORA-1610 ON RECOVER DATABASE FOR CREATED CONTROLFILE)

Oracle running on Unix systems

IBM InfoSphere Application Aware mounts may fail if your Oracle 12c installation does not include Patch# 19404068 for 12.1.0.2.0

Oracle version affected: 12.1.0.2.0

Apply Patch: 19404068

Oracle running on WINDOWS systems

IBM InfoSphere Application Aware mounts may fail if your Oracle 12c installation does not include Patch# 22809813 for 12.1.0.2.0

Oracle version affected: 12.1.0.2.0

Apply Patch: 22809813

Oracle Patch release notes: <https://updates.oracle.com/Orion/Services/download?type=readme&aru=20122528#BABCGCAB>

Oracle Bug 13037524

Querying v\$asm_disk from database instance raises ORA-01455: converting column overflows integer datatype (Doc ID 1473647), caused by Oracle unpublished Bug 13037524

Confirmed Oracle Version Affected: 11.1.0.6 to 11.1.0.7

Fixed in: 11.2.0.1 and above.

IBM InfoSphere Oracle backup may fail if Oracle database 11.1.0.6 to 11.1.0.7 install does not include this patch

Apply Patch: 13037524 for 11.1.0.7

Oracle Known Issue

Oracle known bug for 12.2.0.1 DB on Windows published on Oracle Metalink: Rman Crosscheck on 12.2 Reports ORA-07445: Exception Encountered: Core Dump [skgfifi()+4307] [INT_DIVIDE_BY_ZERO] On Windows (Doc ID 2444452.1)

Confirmed Oracle Version Affected: 12.2 Windows only

Queries on DBA_FREE_SPACE are Slow (Doc ID 271169.1)

During IBM InfoSphere backup, the VDP 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 very 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. Below are the steps to purge the recycle bin.

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;
```


2 IBM InfoSphere Prerequisites for Protecting an Oracle Database

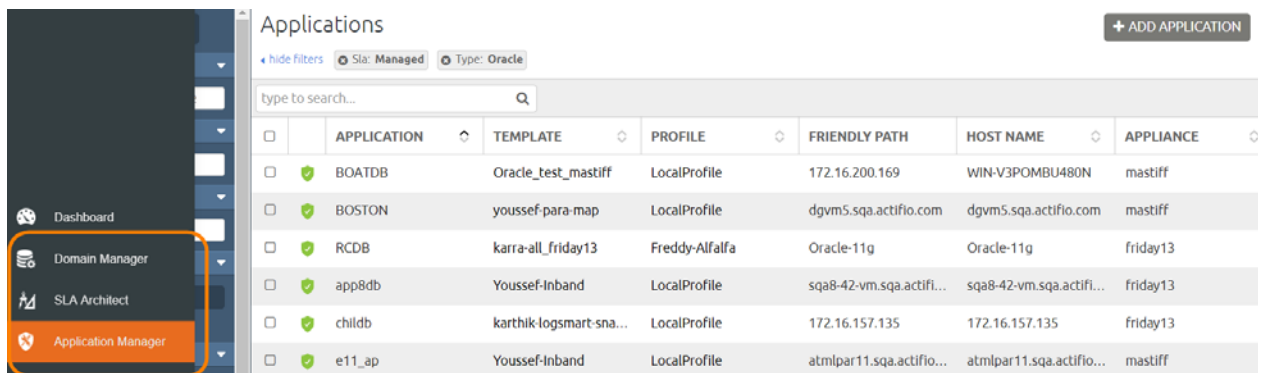
After the IBM InfoSphere preparation and before you can virtualize and access Oracle databases,

1. Review the concepts in [Chapter 3, Oracle Authentication](#) and [Chapter 4, Data Capture under File System and under ASM Disk Group](#),
2. Prepare the database according to the steps in [Chapter 5, Preparing Oracle Databases for Protection](#).

Then you can perform most database protection, access, and recovery operations from the IBM InfoSphere VDP - Global Manager (IVGM).

Prerequisites when Working from the IVGM

Step	Where	What	These procedures are in:
1	The Database Server	Install/upgrade the VDP Connector. Always use the most recent VDP Connector.	Connecting Hosts to IBM InfoSphere VDP Appliances
2	Domain Manager	The database server must be added as a host or as a VM.	IVGM Online Help
3	Application Manager	The database must be discovered as an application.	IVGM Online Help
4	SLA Architect	You need one or more suitable SLA templates and resource profiles for the database.	SLA Policy Overrides for Oracle Databases on page 31
5	Application Manager	There are many Oracle-specific Application Advanced Settings that must be set.	Application Details and Settings for Oracle Databases on page 33



3 Oracle Authentication

This section describes two forms of Oracle database user authentication that you can use from an InfoSphere VDP Appliance:

[Using Oracle with OS Authentication](#)

[Enabling Database Authentication for an Oracle Server from IVGM on page 12](#)

Note: IBM InfoSphere RMAN backup runs as the Oracle binary owner. If the user running the database instance is different than the Oracle OS owner and the group access privileges for the two users are not the same, then the backup will fail.

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 is required to 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

Note: OS Authentication is not supported in Windows environments.

OS Authentication is the default setting in Unix environments. No database user account and no service name are needed. An IBM InfoSphere 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 IVGM Application Details and settings, you can validate the authentication configuration, as shown at right.

Required Application Advanced Settings for an Oracle database when OS Authentication is configured are:

- **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 Oracle Data Guard. See [Protecting from an Oracle Data Guard Node](#) on page 26. In case of 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 is titled 'BOSTON' and contains a table of configuration details. The 'Authentication' section is highlighted with an orange border and contains fields for 'USERNAME' (set to 'username') and 'PASSWORD' (masked with asterisks). A 'Validate Configuration' button is next to the password field. Below the 'Authentication' section is a 'Settings' section with 'Cancel' and 'Save Changes' buttons.

Application Details & Settings	
Select options that will revert back to default. Settings Help	
BOSTON	
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
Authentication	
USERNAME	username
PASSWORD	*****
Validate Configuration	
Settings	
Cancel Save Changes	

Enabling Database Authentication for an Oracle Server from IVGM

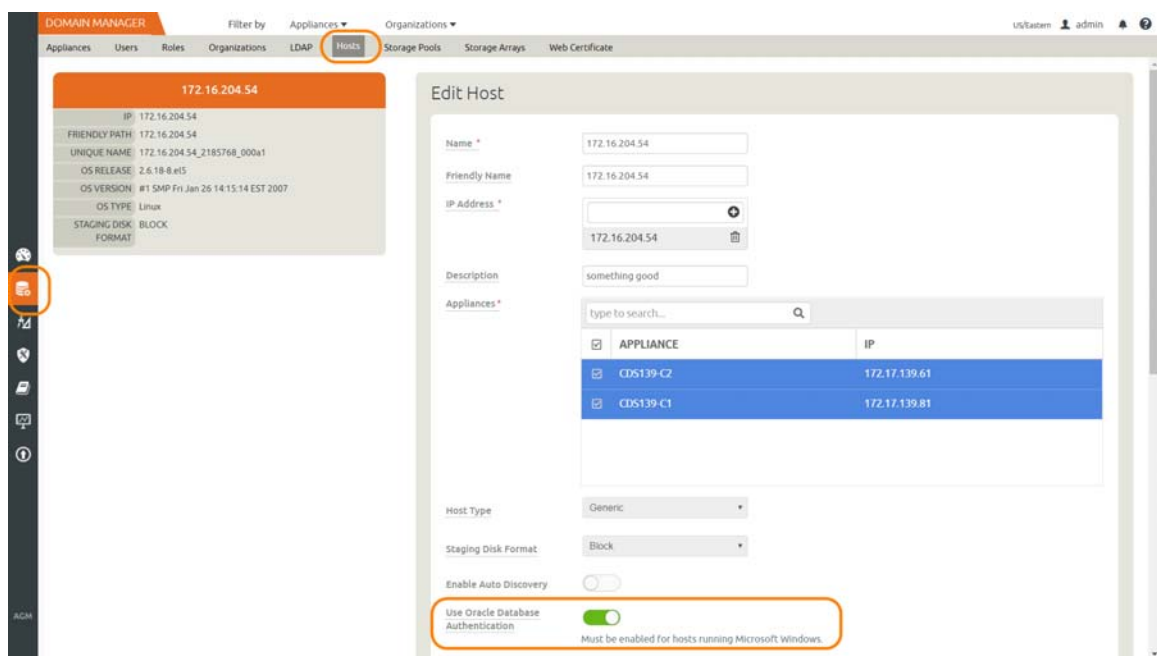
Oracle database authentication uses Oracle database credentials. With Oracle Database Authentication, you must provide database credentials to connect to the database with sysdba privilege (or sysbackup for Oracle 12c).

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

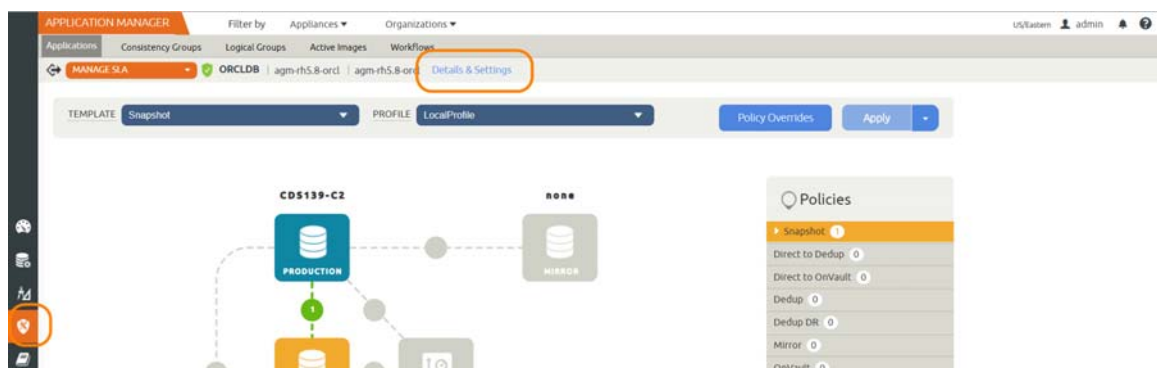
Enabling Oracle Database Authentication

To enable Oracle Database Authentication:

1. Open IVGM to the Domain Manager, Hosts tab.
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 Advanced Settings for the backup to succeed.



4. Go to the Application 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.

Application Details & Settings

STAGING DISK GRANULARITY (CB)

LAST STAGING DISK MINIMUM SIZE (CB)

STAGING DISK MOUNT POINT

RMAN LOG LOCATION

RESTORE VALIDATE ☐ Yes ☒ No

RMAN CATALOG DB NAME

RMAN CATALOG USER

RMAN CATALOG PASSWORD

ORACLE SERVICE NAME

ORACLE DATA GUARD PRIMARY NODE SERVICE NAME

CLUSTER NODES

[Settings Help](#)

[Cancel](#) [Save Changes](#)

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.

4 Data Capture under File System and under ASM Disk Group

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

File System: For all source database configurations other than Oracle ASM, the backup is under a file system. See [Protecting an Oracle Database Under a File System as a File System](#) on page 16.

ASM Disk Group: If a source database configuration is using ASM, the backup is under an ASM disk group. The application advanced 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 16.

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 18.

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 19.

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 IBM InfoSphere 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 IBM InfoSphere 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 IBM InfoSphere presented file system. A snapshot of the staging disk will be taken.

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 IBM InfoSphere 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 IBM InfoSphere-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
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 29.

The Application Advanced Settings that are 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 and Settings for Oracle Databases](#) on page 33.

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

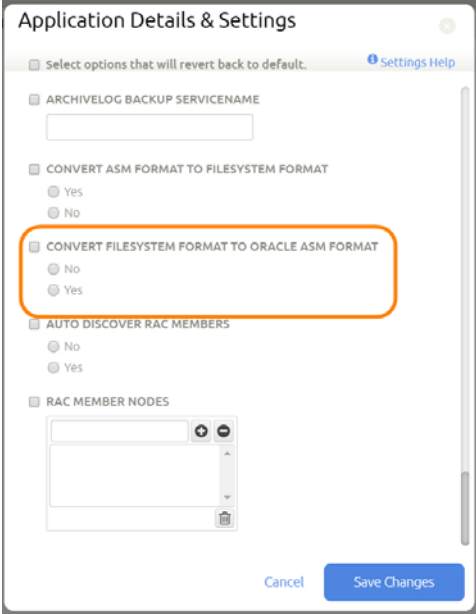
Without any Oracle Exadata Storage, Oracle ASM has the following storage limits if the COMPATIBLE.ASM and COMPATIBLE.RDBMS disk group attributes are set to 12.1 or greater:

- 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 IBM InfoSphere 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 Advance 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.



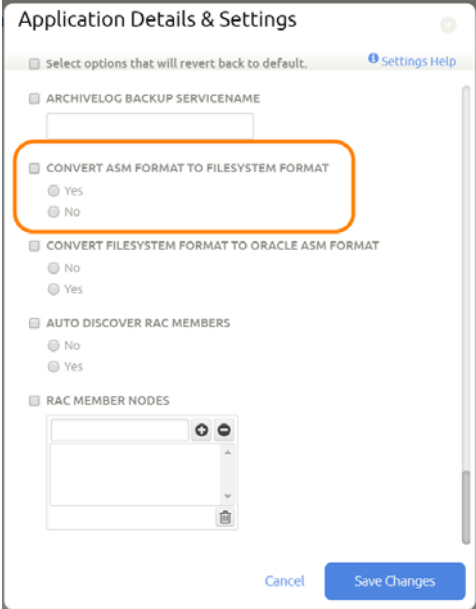
The screenshot shows the 'Application Details & Settings' window. At the top, there is a header bar with the title 'Application Details & Settings' and a 'Settings Help' link. Below the header, there is a section titled 'Select options that will revert back to default.' with a 'Settings Help' link. The settings are organized into sections: 'ARCHIVELOG BACKUP SERVICE NAME' with a text input field; 'CONVERT ASM FORMAT TO FILESYSTEM FORMAT' with 'Yes' and 'No' radio buttons; 'CONVERT FILESYSTEM FORMAT TO ORACLE ASM FORMAT' with 'No' and 'Yes' radio buttons (this section is highlighted with an orange rectangle); 'AUTO DISCOVER RAC MEMBERS' with 'No' and 'Yes' radio buttons; and 'RAC MEMBER NODES' with a list box and a 'Save' button. At the bottom right, there are 'Cancel' and 'Save Changes' buttons.

Application Advanced Settings to Capture an Oracle Database under File System to ASM Disk Group in the IVGM

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 Advanced Settings select **Yes** under **Convert ASM Format to Filesystem Format**. File system backup will be used for all source databases including ASM.

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 from the VDP Desktop or the IVGM. Application advanced settings are detailed in [Application Details and Settings for Oracle Databases](#) on page 33.



The screenshot shows the 'Application Details & Settings' window. At the top, there is a header bar with the title 'Application Details & Settings' and a 'Settings Help' link. Below the header, there is a section titled 'Select options that will revert back to default.' with a 'Settings Help' link. The main content area contains several settings:

- ARCHIVELOG BACKUP SERVICENAME**: A text input field.
- CONVERT ASM FORMAT TO FILESYSTEM FORMAT**: A radio button group with 'Yes' selected and 'No' unselected. This section is highlighted with an orange border.
- CONVERT FILESYSTEM FORMAT TO ORACLE ASM FORMAT**: A radio button group with 'No' selected and 'Yes' unselected.
- AUTO DISCOVER RAC MEMBERS**: A radio button group with 'No' selected and 'Yes' unselected.
- RAC MEMBER NODES**: A list box with a '+' button to add nodes and a '-' button to remove nodes.

At the bottom right, there are two buttons: 'Cancel' and 'Save Changes' (highlighted in blue).

Application Advanced Settings to Capture an Oracle Database under ASM Disk Group as a File System in the IVGM

5 Preparing Oracle Databases for Protection

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

Preparation Procedures for Oracle Databases in Unix Environments

Step	Preparation Procedure
1	Preparing Oracle Databases in a Unix Environment Using OS Authentication on page 22
2	Preparing to Capture a Database from Oracle ASM to Oracle ASM on page 23 (This is needed only for RAC or Standalone ASM configurations.) Preparing to Capture a Database from Oracle ASM to Filesystem on page 23
3	Preparing Oracle Database Authentication in a Unix Environment on page 24 Enable Database Block Change Tracking (optional) on page 26 Protecting from an Oracle Data Guard Node on page 26 Configuring RAC Transparent Failover of IBM InfoSphere RMAN Backup to Other Nodes on page 27 Oracle Archive Logs Compression on page 27 Manually Calculating Log Staging Disk Size (optional) on page 28 Configuring Oracle Database Services for Load Balancing across Multiple Nodes on page 29

Preparing Oracle Databases in a Unix 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 Unix Environment](#) on page 24.

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
```

Set the Oracle Database SID Entry

The database SID entry must be set in the /etc/oratab file (/var/opt/oracle/oratab on Solaris). For a database named “oasm” the entry looks like:

```
#cat /etc/oratab
oasm:/home/oracle/app/oracle/product/11.1.0/db_1:Y
```

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 or /var/opt/oracle/oratab on Solaris systems)
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. IBM InfoSphere 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 running 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 or /var/opt/oracle/
oratab on Solaris systems)
#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 IBM InfoSphere protection. The IBM InfoSphere 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 Advanced Setting Must Be Enabled

If you are protecting an Oracle ASM database to a filesystem, then check these Application Advanced Settings:

- **Force out-of-band backup** application advanced setting is enabled from the VDP Desktop or the IVGM.
- **Convert ASM to File System** is set to Yes. Application advanced settings are detailed in [Application Details and Settings for Oracle Databases](#) on page 33.

Preparing Oracle Database Authentication in a Unix Environment

These additional preparation steps are required only if you will use database authentication. Oracle database authentication is described in [Chapter 3, Oracle Authentication](#).

1. Follow the steps in [Preparing Oracle Databases in a Unix Environment Using OS Authentication](#) on page 22.
2. Create a database user account for IBM InfoSphere 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 USERNAMESYSDBSYSOPSYSAS
-----
1 SYSTRUETRUEFALSE
2 SYSTRUETRUEFALSE
1 ACT_RMAN_USERTRUEFALSE
2 ACT_RMAN_USERTRUEFALSE
```
5. Test the service name as described in:
[Creating and Verifying the Oracle Servicename in a non-RAC Environment](#) on page 24
[Creating and Verifying the Oracle Servicename in a RAC Environment](#) on page 25

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 Advanced Settings described in [Application Details and Settings for Oracle Databases](#) on page 33.
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 IBM InfoSphere 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 Advanced Settings described in [Application Details and Settings for Oracle Databases](#) on page 33.

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 IBM InfoSphere 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 Advanced Settings described in [Application Details and Settings for Oracle Databases](#) on page 33.

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 Advanced Settings described in [Application Details and Settings for Oracle Databases](#) on page 33.

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 Application Advanced Settings (see [Application Details and Settings for Oracle Databases](#) on page 33):

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;
```

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.

Configuring RAC Transparent Failover of IBM InfoSphere RMAN Backup to Other Nodes

The VDP 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.

In Advanced Settings, Cluster Nodes, specify the failover node choice in a Oracle RAC environment in this way:

Failover choice:Node IP:ServiceName:Role

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 IBM InfoSphere RMAN backup. This can be a new dedicated service created for IBM InfoSphere 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

- 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 Advanced Settings, the Cluster Nodes entry will be:

Failoverchoice:NodeIP:ServiceName:Role

1:172.1.1.1:act_svc_node2:F

Oracle Archive Logs Compression

IBM InfoSphere archivelog backup supports Oracle log backupset compression. The type of compression you select depends on the RMAN configuration settings described below. Select the right option based on your use case and available resources.

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

Selecting an Oracle Compression Algorithm

Algorithm Name	Use This For	Oracle Versions
BASIC	good compression ratio	9.2.0.0.0 and later
BZIP2	good compression ratio	9.2.0.0.0 and later
LOW	maximum possible compression speed	11.2.0.0.0 and later
ZLIB	balance between speed and compression ratio	11.0.0.0.0 and later
MEDIUM	balance between speed and compression ratio	11.0.0.0.0 and later
HIGH	maximum possible compression ratio	11.2.0.0.0 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;

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

Manually Calculating Log Staging Disk Size (optional)

The VDP 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 Advanced Settings, detailed in [Application Details and Settings for Oracle Databases](#) on page 33.

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. In this example, assume a four-node RAC environment; nodes 3 and 4 are to be load-balanced for backup use.

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 VDP Connector on all nodes.
2. Setup the ASM disk group mapping to node 3 and node 4 using Application Advanced Settings.
3. Create a database service using `srvctl` to run from node 3 and node 4.
4. Use this service to specify under Application Advanced 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 Advanced 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 IBM InfoSphere 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 Advanced 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 Advanced Settings Archivelog Backup Servicename.

6 Details and Settings for Oracle Databases in IVGM

There are two kinds of advanced settings:

[SLA Policy Overrides for Oracle Databases](#) on page 31

[Application Details and Settings for Oracle Databases](#) on page 33

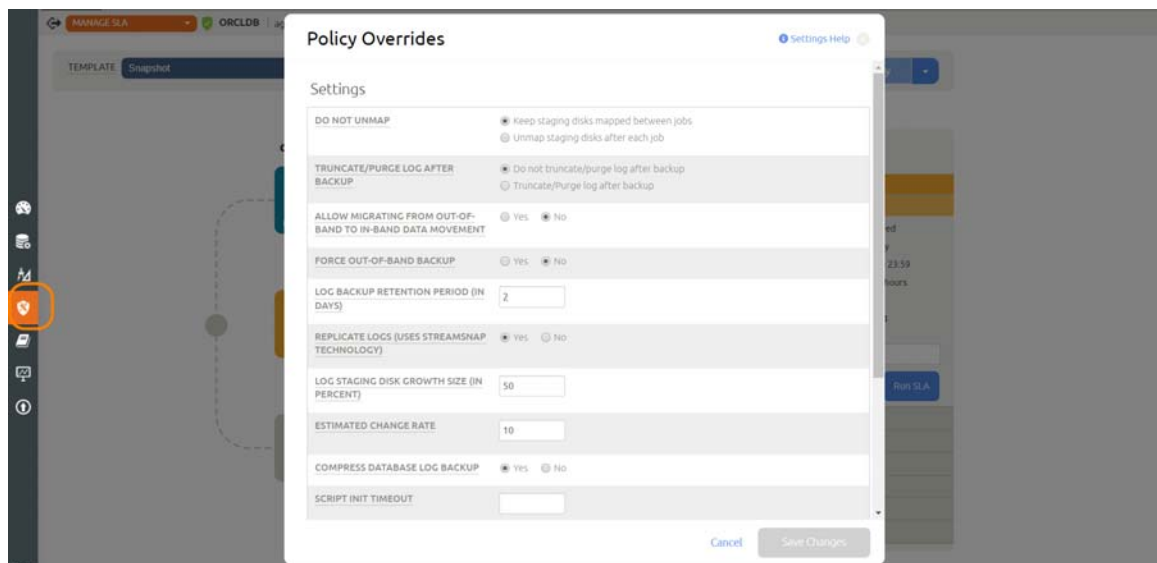
SLA Policy Overrides for Oracle Databases

SLA policy overrides enable you to customize an SLA policy for a specific application's needs. To set policy overrides:

1. Go to the Application Manager 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.



SLA policy overrides available for Oracle databases are:

Truncate/Purge Log After Backup: If you want to manage log purging, then select Truncate/Purge Log After Backup. 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.

Enable Database Log Backup: To enable database archive log backup select Yes. To set the frequency of log backup set the RPO (next).

RPO: Set an RPO value for archive logs in minutes based on the recovery point objective for the database being protected. For example, if the recovery point objective is one hour, set RPO to 60.

Log Backup Retention Period: The archive log backup under IBM InfoSphere staging disk will be retained to the value set here. Backup log retention can be different from snapshot retention. For example, if snapshot retention is 2 days and dedup retention is 7 days, then a value of 7 days will ensure that all the images under dedup will have the archive log backup to roll-forward the image to any point in time.

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

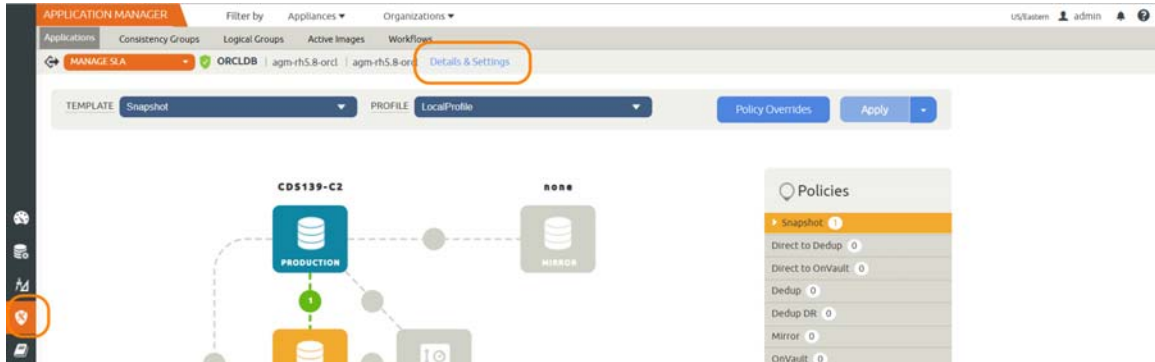
Estimated Change Rate: Estimate the percentage by which the database data changes daily.

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

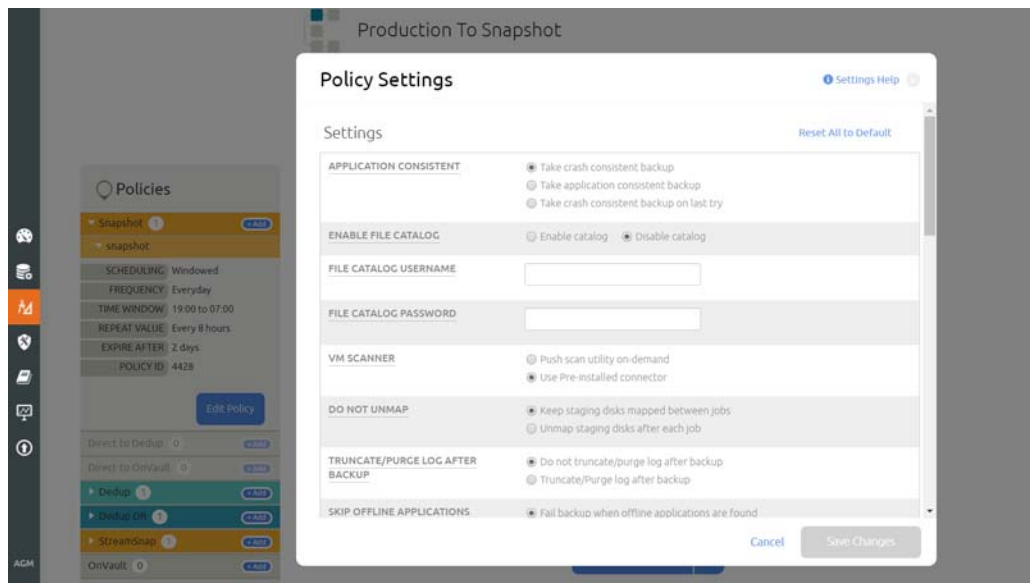
Application Details and Settings for Oracle Databases

To set database-specific details:

3. Go to the Application Manager and right-click a database. Select **Manage SLA**.
4. At the top of the page, click **Details and Settings**.



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



Application Details and Settings

The Application Details and Settings for an Oracle database are:

- **Username and Password:** 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:** The default value is `sysdba`, but select `sysbackup` for an Oracle 12c database.
- **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.

- **Staging Disk Size:** By default, the Oracle 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.
- **Staging Disk Granularity:** Default size is 1TB.
- **Last Staging Disk Minimum Size:** If an application might require multiple staging disks, enter the minimum size to be allocated for the staging disk used for the last part of that application.
- **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 IBM InfoSphere 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 IBM InfoSphere backup, as described in [Creating and Verifying the Oracle Servicename in a non-RAC Environment](#) on page 24 and in [Creating and Verifying the Oracle Servicename in a RAC Environment](#) on page 25. If not specified, then by default IBM InfoSphere 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 IBM InfoSphere 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 26.
- **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 IBM InfoSphere RMAN Backup to Other Nodes](#) on page 27.
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 InfoSphere VDP Appliance uses this as the backup node instead of using the original protected node.
- **Connector Options:** Use this only under the direction of IBM InfoSphere 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.

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.
- **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).

- **Log Staging Disk Size:** Used if log backup policy is set. By default IBM InfoSphere 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 28 for more information on determining this value.
- **Do Not Uncatalog:** To keep RMAN backup cataloged after each backup job. By default, IBM InfoSphere backup will be cataloged at the start of backup and then be un-cataloged at the end of the backup.
- **Force New Level 0 Backup:** If for any reason a full level 0 backup is required, overwriting the IBM InfoSphere 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.
- **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.
- **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.

Note: For security reasons, 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.

- **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.
- **Truncate/Purge Log After Backup:** This is to set the archive log purging. The current database archive logs will be purged with the next database backup job. If database logs are protected, then purging is based on the Log Purging Retention Period option. See [Archive Log Purge Behavior Policy Advanced Settings are Overridden by Application Advanced Settings](#) on page 36.
- Choose **Do Not Unmap** if you want temporary staging disks mapped to the host and used during data movement for backup to remain mapped to the host. LUNs are mapped during the first job and all the subsequent jobs reuse the same mapped LUN. By default, this option is selected.

- **Compress Database Log Backup:** Flag to enable the log backup compression. If selected, the backup of archive log will be run using RMAN compress option.
- **Estimated Change Rate:** Entering a percentage for the database's daily change rate will allow the InfoSphere VDP Appliance to better predict how to grow the staging VDisk for the database's log file.

Archive Log Purge Behavior Policy Advanced Settings are Overridden by Application Advanced Settings

SLA Policy Advanced Setting		Application Advanced Setting		Behavior
Truncate/ Purge Log After Backup	Enable Log Backup	Truncate/ Purge Log After Backup	Log Purging Retention	
Yes	No	No value	No value	Purge log will run at end of each database backup with a retention of 24 hours i.e. delete archivelog older then sysdate -1
Yes	Yes	No value	No value	Purge log will run at the end of each archive backup with a retention of 24 hours i.e. delete archivelog older then sysdate -1
Yes	No	Do not Truncate/ Purge	No Value	Archive log will not be purged
Yes	Yes	Do not Truncate/ Purge	No Value	Archive log will not be purged
No	No	Truncate/ Purge log	No value	Purge log will run at end of each database backup with a retention of 24 hours i.e. delete archivelog older then sysdate -1
No	No	Truncate/ Purge log	Value (n hours)	Purge log will run at end of each database backup with a retention of n hours i.e. delete archivelog older then sysdate -n/24
No	Yes	Truncate/ Purge log	No value	Purge log will run at the end of each archive backup with a retention of 24 hours i.e. delete archivelog older then sysdate -1
No	Yes	Truncate/ Purge log	Value (n hours)	Purge log will run at the end of each archive backup with a retention of n hours i.e. delete archivelog older then sysdate -n/24

7 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.

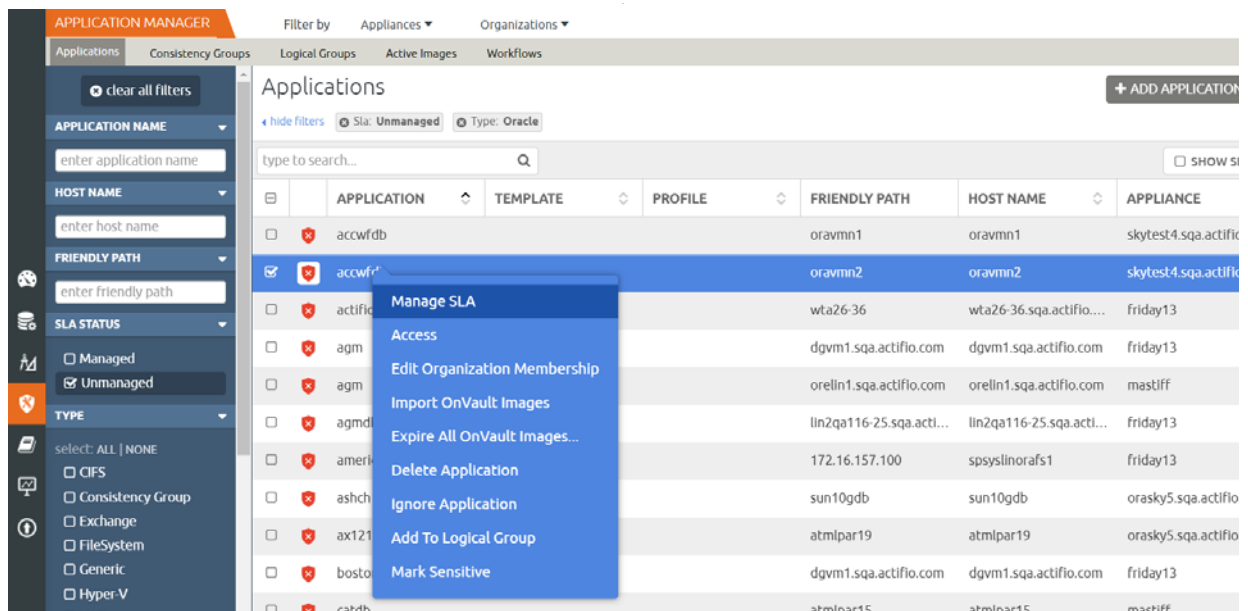
Before you can virtualize and protect Oracle databases, see:

- [Chapter 2, IBM InfoSphere Prerequisites for Protecting an Oracle Database](#)
- [Chapter 3, Oracle Authentication](#)
- [Chapter 4, Data Capture under File System and under ASM Disk Group](#)
- [Chapter 5, Preparing Oracle Databases for Protection.](#)

Virtualizing an Oracle Database for Data Protection and Agility using the IVGM

To capture an Oracle database and its logs:

1. Open the IVGM to the **Application Manager** 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**.



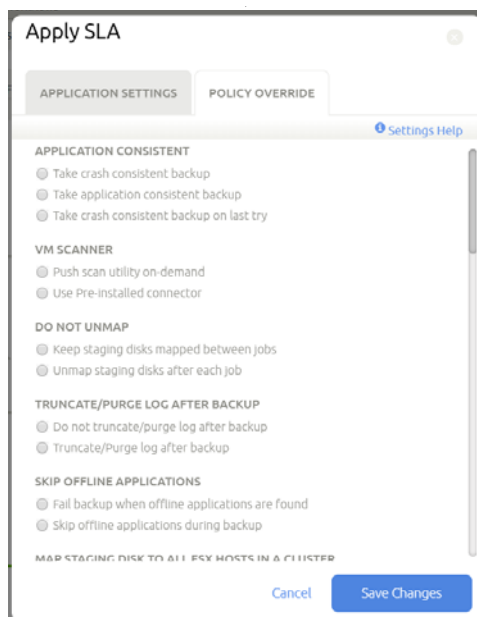
Selecting an Oracle Database in the IVGM

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



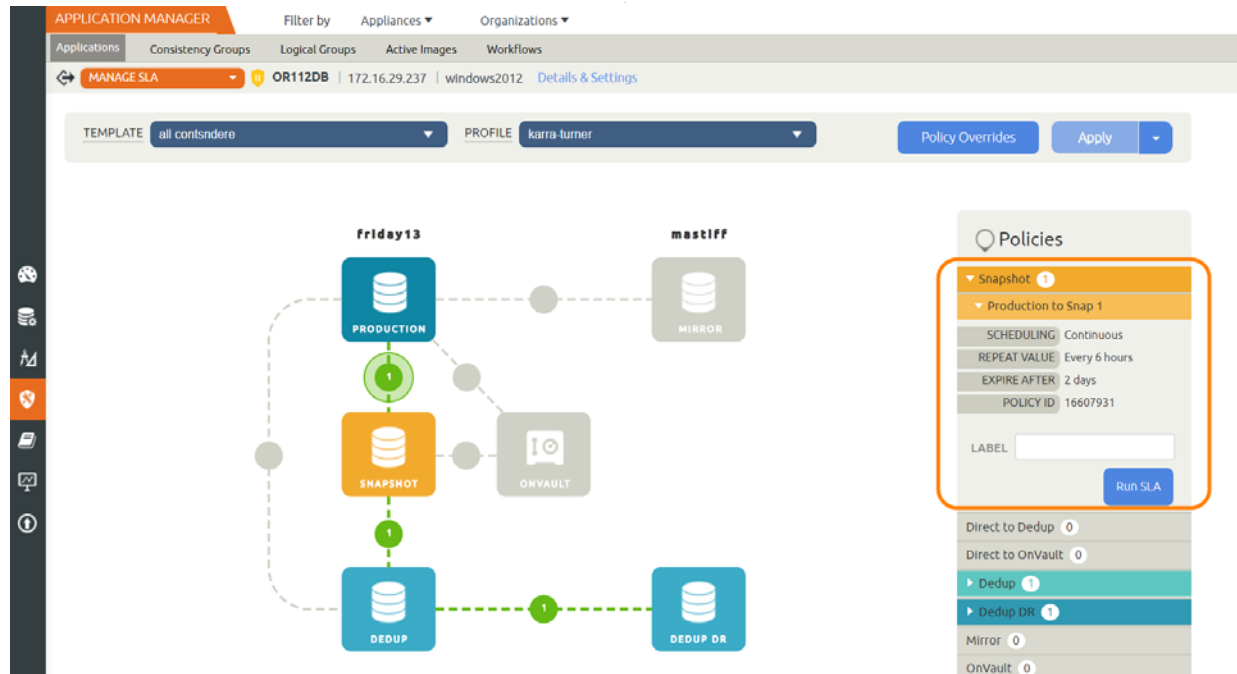
Applying an SLA in the IVGM

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



Applying SLA Policy Overrides in the IVGM

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**.



Running the SLA in the IVGM (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 System Monitor to view the progress and details of the job.

SYSTEM MONITOR Filter by Appliances Organizations Refresh in

Jobs Events

clear all filters

JOB NAME filter by job name

HOST filter by host

APPLICATION filter by application

STARTED

ENDED

QUEUED

STATUS

select: ALL | NONE

☐ Canceled

☐ Failed

☐ Not Run

☐ Queued

☐ Retry

☒ Running

☐ Succeeded

TYPE

select: ALL | NONE

☐ Catalog

☐ Clone

Jobs hide filters Status: Running

type to search...

JOB	STATUS	HOST	APPLI...	APPLI...	QUEUED	STARTED	ENDED	DURAT...	TYPE	TEMPL...	POLICY
Job_21...	Running: 60%	testvm3	TestVM3	littleras...	08-02 1...	08-02 1...		0d 00:0...	expirati...	catch u...	onvalut...
Job_21...	Running: 4%	testvm3	TestVM3	littleras...	08-02 1...	08-02 1...		0d 00:0...	OnVault	catch u...	onvalut...
Job_54...	Running: 17%	ATM-LP...	/	mastiff	08-02 1...	08-02 1...		0d 00:0...	snapshot	Youssef...	Product...
Job_21...	Running: 25%										
Job_21...	Running: 6%										
Job_54...	Running: 10%										
Job_54...	Running: 4%										
Job_00...	Running: 75%										
Job_00...	Running: 52%										
Job_00...	Running: 52%										
Job_00...	Running: 96%										
Job_00...	Running: 39%										

1 - 25 of 69 Jobs

Job_5468931a Details

ADDITIONAL VOLUME INFO RETRIES STATISTICS

ID: Job_5468931a_590021382708

PROGRESS: 17%

APPLIANCE: mastiff

POLICY NAME: Production to Snap 1

PRIORITY: medium

JOB TYPE: snapshot

STATUS: running

HOST NAME: ATM-LPAR7_scsi

TEMPLATE NAME: Youssef_vscsi_lpar-map

APPLICATION NAME: /

TARGET HOST: ATM-LPAR7_scsi

DURATION: 00:03:24

START DATE: 2018-08-02 19:31:04

CONSISTENCY DATE: 2018-08-02 19:31:04

EXPIRATION DATE: 2018-08-02 21:31:04

LAST CONSTRAINT DATE: 2018-08-02 19:27:04

QUEUE DATE: 2018-08-02 19:20:00

[Return to Jobs](#) [Cancel Job](#)

SLA Job Details for an Oracle Database in the IVGM

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 IBM InfoSphere backup job, the backup may have incomplete log information.

8 Accessing, Recovering, or Restoring an Oracle Database via the IVGM

IBM InfoSphere 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 IBM InfoSphere 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 from the IVGM on page 42](#)
- o [Mounting an Oracle Database Image Protected Under a File System for Data Access from the IVGM on page 45](#)

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 InfoSphere VDP 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 in the IVGM on page 47](#).

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.

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

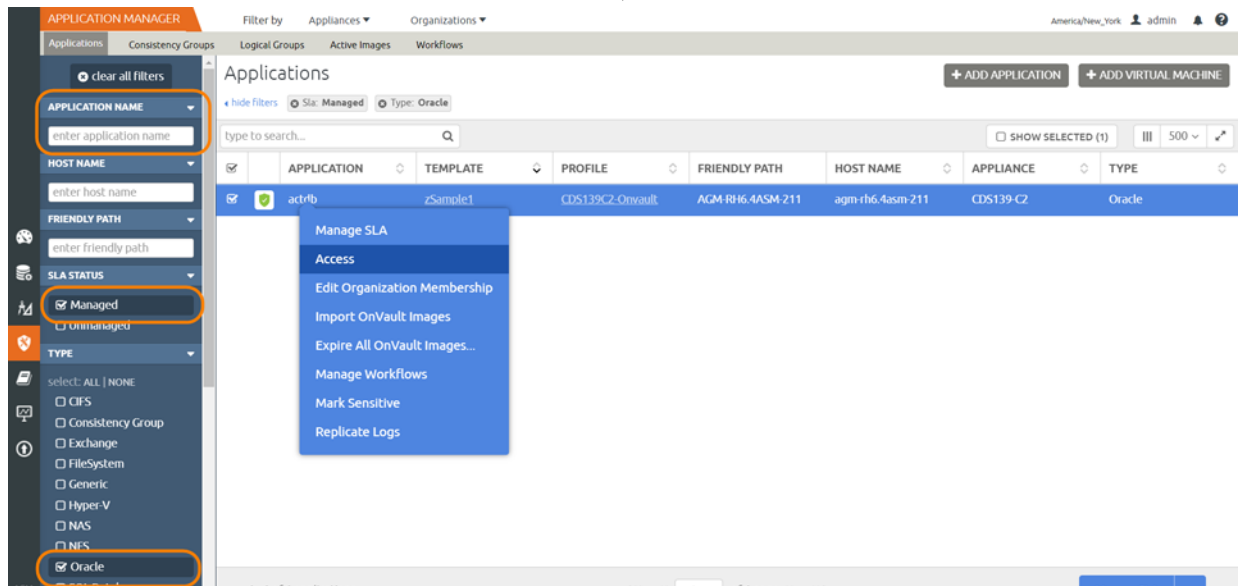
Oracle-specific workflows in the IVGM are detailed in the IVGM online help.

After any database server reboot where an IBM InfoSphere image is mounted, or if IBM InfoSphere backups are in progress for the database at the time of reboot/crash, please see [Bringing IBM InfoSphere-Protected ASM Diskgroups Back Online after Reboot of a Target DB Server on page 51](#).

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

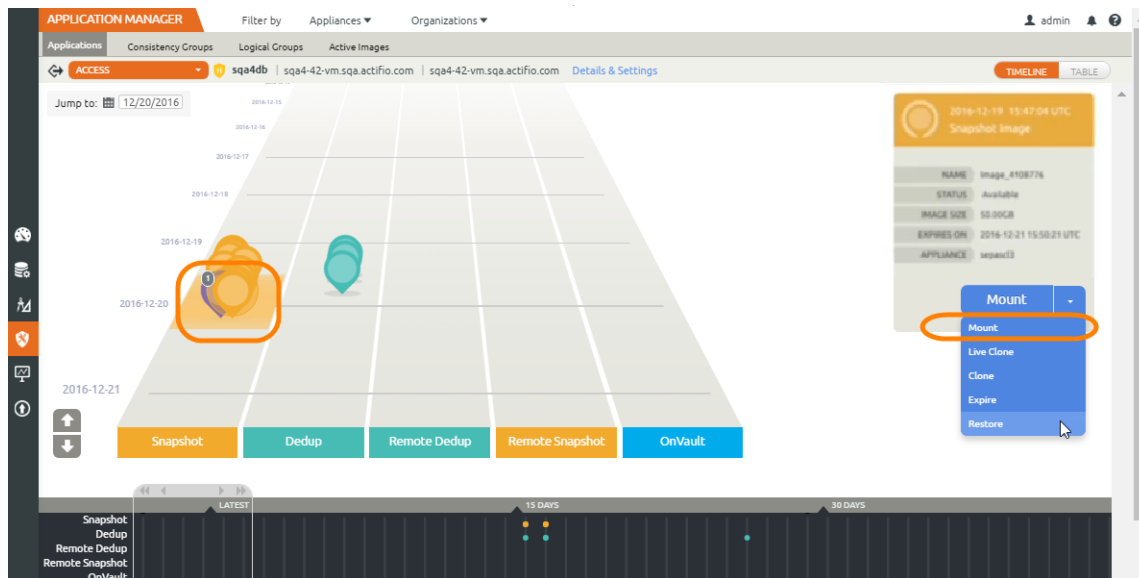
To mount an Oracle database image for data access:

1. Open the IVGM to the **Application 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**.



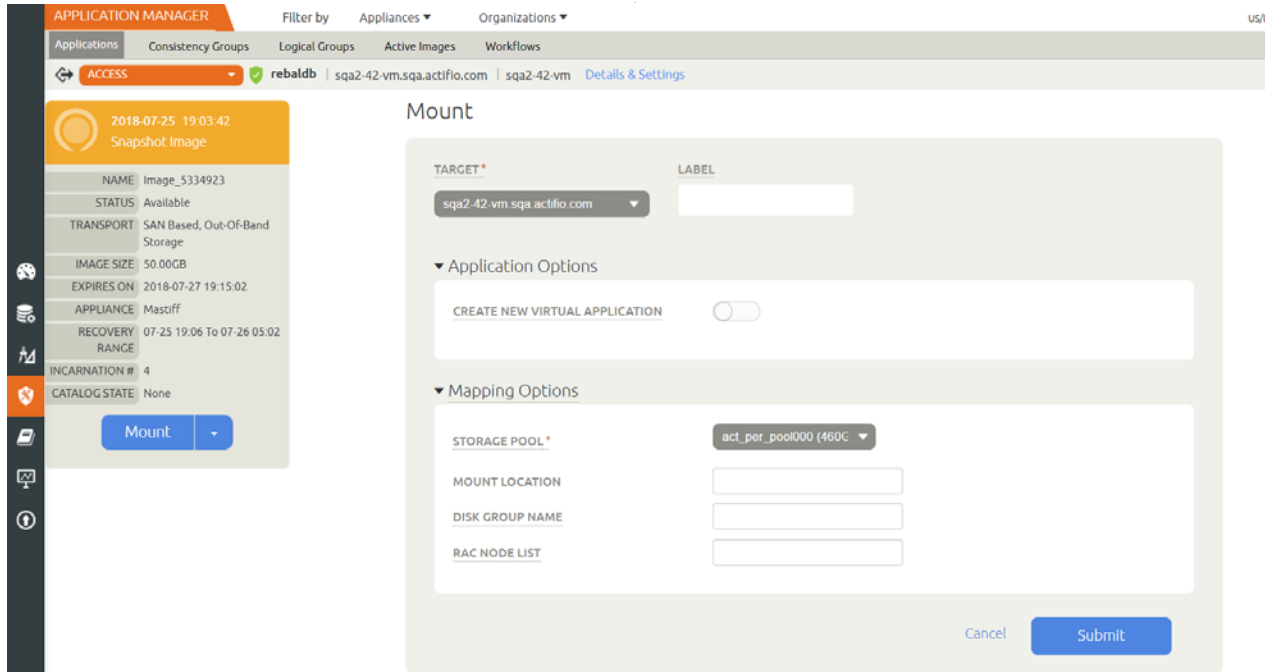
Selecting an Oracle Database in the IVGM

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



Selecting a Managed Oracle Database Image in the IVGM

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



Mount Details for an Oracle Database in the IVGM

- 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 in the IVGM](#) on page 47.
- 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 at which you want to mount the new database:
If the path exists as an empty folder, the VDP 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 point 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 IBM InfoSphere-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 System Monitor to view the progress and details of the job.

The screenshot displays the IBM InfoSphere System Monitor interface. On the left, a sidebar contains filter options for Job Name, Host, Application, Started, Ended, Queued, and Status. The main panel shows a table of jobs, with Job_1852961 highlighted. An orange arrow points from this job entry to a detailed view of the job. The detailed view, titled 'Job_1852961 Details', shows various attributes including ID, Appliance, Policy Name, Priority, Job Type, Status, Host Name, Template Name, Application Name, Target Host, Message, Duration, Start Date, Consistency Date, Expiration Date, End Date, and Queue Date.

JOB	STATUS	HOST	AP...	AP...	LO...	QU...	ST...	EN...	DURATION	TYPE	TE...
Job_1852961	Running: 58%	ag...	Con...	CD...	07...	07...			0d 00:13:23	mo...	Ora...

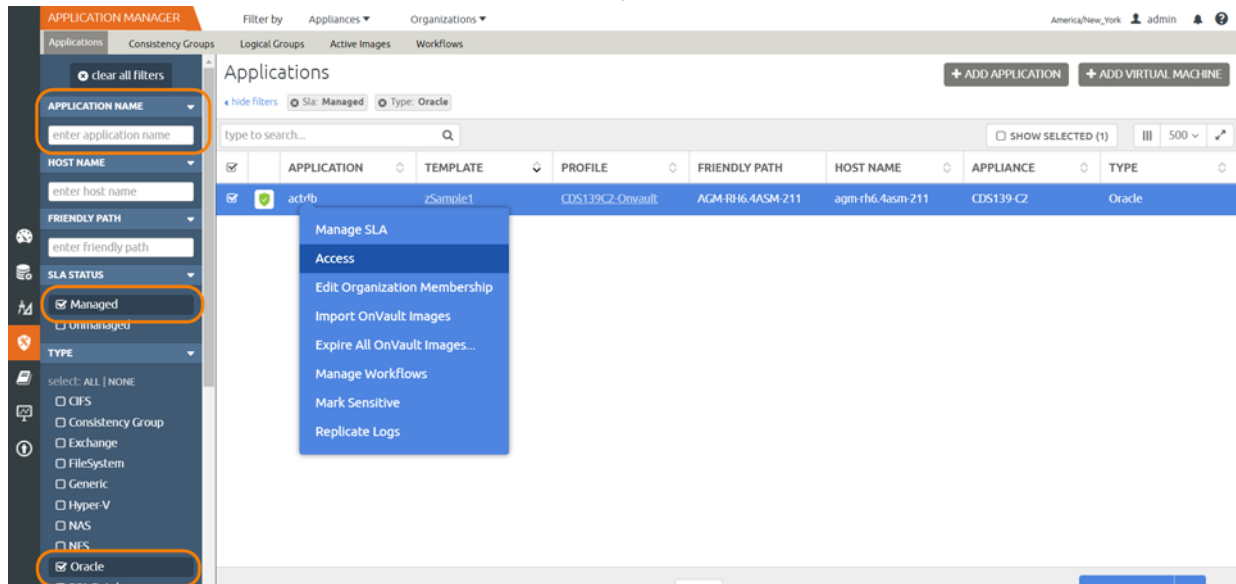
Job_1852961 Details	
	ADDITIONAL
ID	Job_1852961_590021132730
APPLIANCE	CDS139-C2
POLICY NAME	Production to Snap 1
PRIORITY	medium
JOB TYPE	mount
STATUS	succeeded
HOST NAME	agm-rhs.8-orcl
TEMPLATE NAME	OracleSnap
APPLICATION NAME	Con_Grp1528401955796
TARGET HOST	agm-rhs.8-orcl
MESSAGE	Success
DURATION	00:01:40
START DATE	2018-07-20 15:11:21
CONSISTENCY DATE	2018-07-11 02:10:39
EXPIRATION DATE	2100-01-01 00:00:00
END DATE	2018-07-20 15:13:02
QUEUE DATE	2018-07-20 15:11:21

Mount Job Details for an Oracle Database Image in the IVGM

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

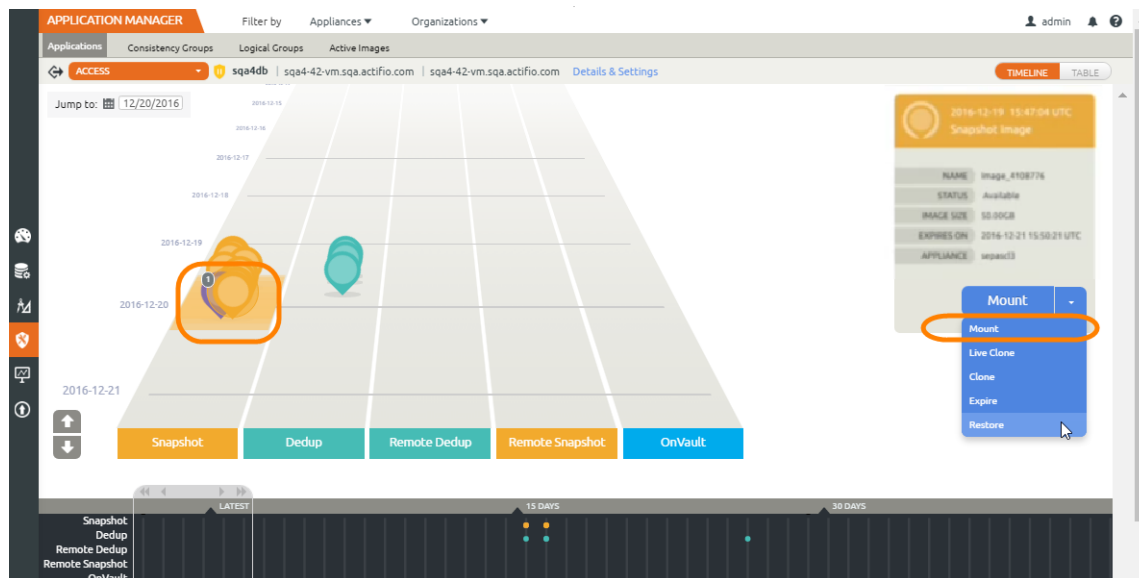
To mount an Oracle database image for data access:

1. Open the IVGM to the **Application 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 in the IVGM

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



Selecting a Managed Oracle Database Image in the IVGM

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 in the IVGM](#) on page 47.

The screenshot shows the 'APPLICATION MANAGER' interface. On the left, a sidebar lists various icons. The main area displays a 'Mount' dialog for a 'Snapshot Image' created on 2017-10-19 at 07:02:10. The dialog has a 'TARGET' dropdown set to 'karthiknode1' and an empty 'LABEL' field. Under 'Application Options', the 'CREATE NEW VIRTUAL APPLICATION' toggle is turned off. Under 'Mapping Options', 'MAP TO ALL ESX HOSTS' is turned off, 'MOUNT MODE' is set to 'vRDM', 'MARK DEPENDENT' is turned off, 'STORAGE POOL' is set to 'mdisk9 (5318GB Free)', and 'MOUNT LOCATION' is an empty text field. 'Cancel' and 'Submit' buttons are at the bottom right.

Mounting an Oracle Database to a VM in the IVGM

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 VDP 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 Point	No Mount Point
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 in the IVGM

Before You Begin

An IBM InfoSphere 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 the control file.

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 23.

Patching Oracle 12c

IBM InfoSphere 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
- (Patch 19404068) IBM AIX on POWER Systems (64-bit) for Oracle 12.1.0.2.0
- (Patch 19404068) Solaris on SPARC (64-bit) 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 source and target host database versions match, and make sure that there are enough resources (like memory & CPU) on the target database server to meet your performance requirements.

Mounting an Oracle Database as a Virtual Application

To mount an Oracle database as a virtual application:

1. Start mounting the image as detailed in [Mounting an Oracle Database Image Protected Under a File System for Data Access from the IVGM](#) on page 45. 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.
- 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.

Note: There is one exception to this: if the target server is a VMware VM, you must select "pRDM" when performing 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 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.
specified	specified	Set SGA and PGA for the database provisioned on the target to: $SGA_TARGET = \text{Database Memory Size in MB} \times (SGA\%/100)$ $PGA_TARGET = \text{Database Memory Size in MB} \times (100 - SGA\%/100)$ <hr/> <p>Note: Do not set SGA to 100. To avoid database slowness, be sure to reserve some memory space for PGA.</p> <hr/>

- 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 **Diagnostic_Dest:** Parameter to configure diagnostic destination on the host. If not specified, diagnostic_dest will be set to ORACLE_HOME.
- 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 **Character set:** Parameter to configure character set encoding, for the database being provisioned. If not specified, AL32UTF8 will be used.
- 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 **Do not change database DBID:** If selected, new database's DBID will not be changed.
- 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 **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 attempting to connect 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 RAC configuration and IBM InfoSphere stores the copy in ASM format. If selected, this performs an application aware mount to a stand alone ASM non-RAC instance. Do not select this option if a RAC node list has been provided.
- o **Environment variable:** If you have any user-defined environment variables to be passed to pre / post scripts, you can enter one here.

Application Aware Oracle Database Mount: Application Options and Advanced Options

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

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

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

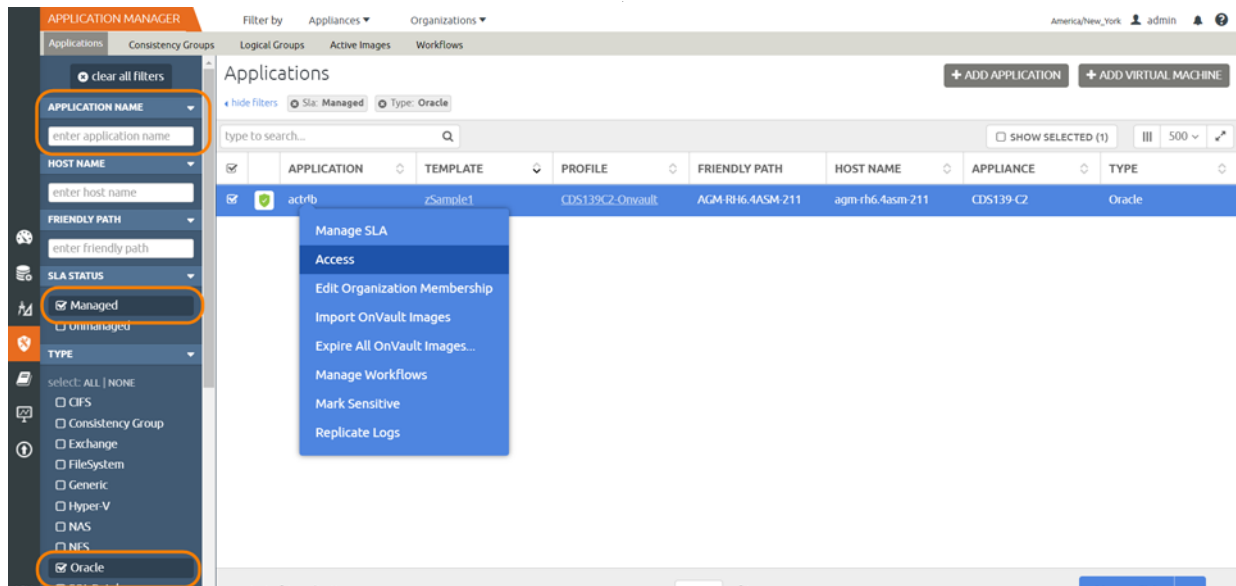
1. Check that the target database server is back up and that ASM and RAC system are also up.
2. Restart the VDP Connector (from root).
3. Set ASM environment.
4. Login to ASM sqlplus and check the disk group status:

```
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 via IVGM, Overwriting the Production Database

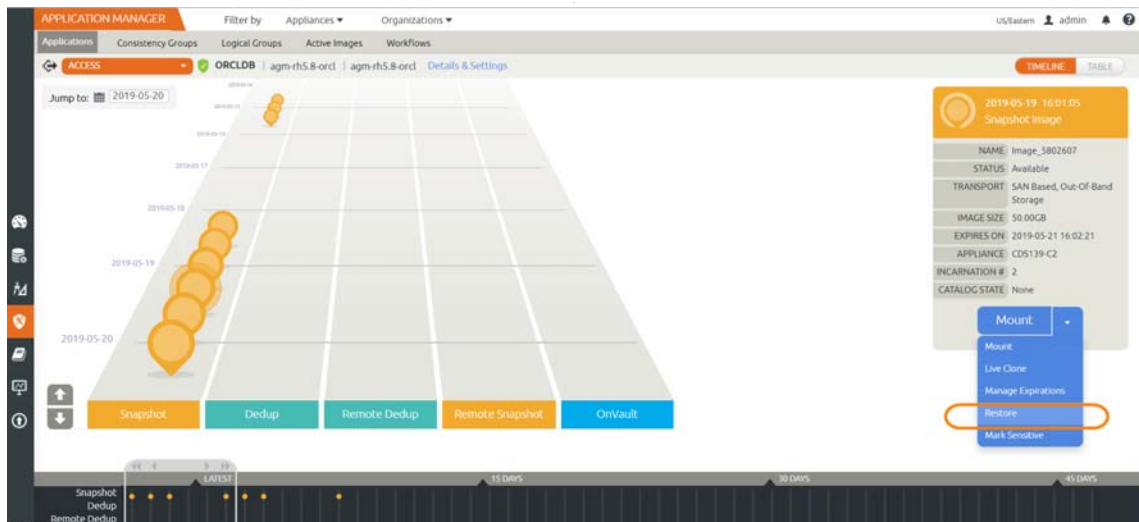
To restore an Oracle database out-of-band backup image, overwriting the original production database:

1. Open the IVGM to the **Application 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**.



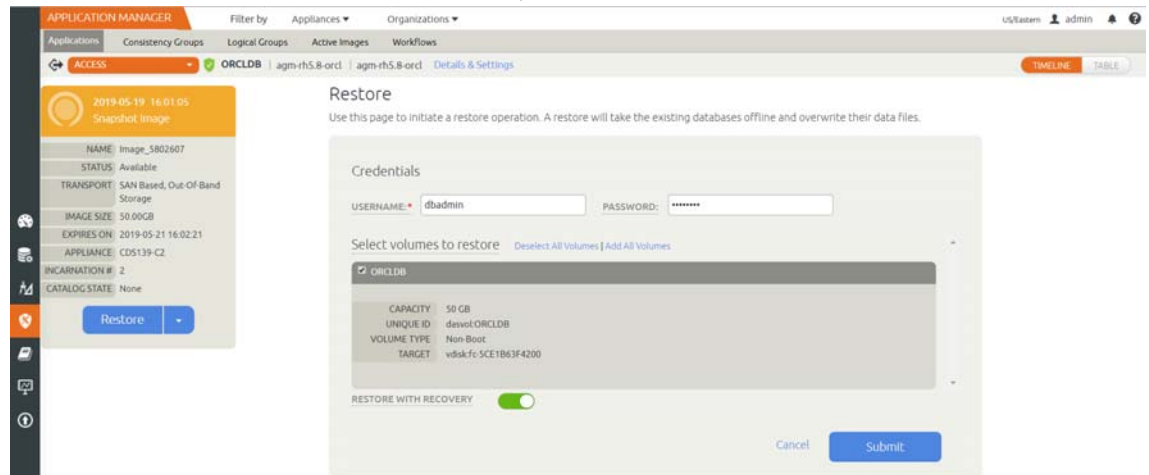
Selecting an Oracle Database to be Restored

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



The Restore Option in the IVGM

4. From the calendar tool, select the date/time range in which the database image was taken. If the policy that protected the database had log protection enabled, the backup represents both the database and its logs.
5. Fill in the Restore page options as needed, then click **Submit**.



Restoring an Oracle Database

6. A warning dialog appears. Read it and enter **DATA LOSS** to confirm.
7. The job is queued to start with the next available job slot. You can view progress from the System Monitor.

Note: When you restore a database, the SLA options (Run Schedule, Expire Data) of the database are disabled.

9 Recovering an Oracle Database Manually Using RMAN

The procedures to restore a database protected under a file system or an ASM Disk Group 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 IBM InfoSphere 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 IBM InfoSphere 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 IBM InfoSphere-protected or are not IBM InfoSphere-protected.

Six RMAN Procedures to Recover Databases

Source Database	Protected Under	Archivelog is	See
Non-RAC, Non-ASM	File System	Logs not IBM InfoSphere-Protected	Recovering a Non-RAC Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere on page 56
RAC or Standalone ASM	File System	Logs not IBM InfoSphere-Protected	Recovering a RAC ASM Oracle Database to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere on page 57
Non-RAC, Non-ASM	File System	Logs are IBM InfoSphere-Protected	Recovering a Non-RAC Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog on page 59
RAC or Standalone ASM	File System	Logs are IBM InfoSphere-Protected	Recovering a RAC ASM Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog on page 61
RAC or Standalone ASM	ASM Disk Group	Logs not IBM InfoSphere-Protected	Recovering an Oracle Database to a Scheduled Backup Point if the archivelog is not Protected through IBM InfoSphere on page 63
RAC or Standalone ASM	ASM Disk Group	Logs are IBM InfoSphere-Protected	Recovering an Oracle Database to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog on page 65

Note: IBM InfoSphere-Protected means that database log protection is enabled in the SLA policy advanced settings.

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

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

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

1. Mount the image from the source database server to recover. In the Mount window, provide a mount point 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 from the IVGM](#) on page 45.

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 IBM InfoSphere 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 IBM InfoSphere 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 any warning from RMAN looking for the next archivelog as 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 IBM InfoSphere

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

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

1. Mount the image from the source database server to recover. In the Mount window, provide a mount point 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 from the IVGM](#) on page 45.
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 IBM InfoSphere 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 IBM InfoSphere 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 archive log 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 IBM InfoSphere-Protected archivelog

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

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

1. Mount the image from the source database server to recover. In the Mount window, provide a mount point 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 from the IVGM](#) on page 45.

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 IBM InfoSphere 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 archive log folder from IBM InfoSphere mounted database image and archive log image to RMAN

```
rman> run { catalog start with '/acttestdb/datafile' noprompt;  
catalog start with '/acttestdb/archive log' 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 IBM InfoSphere-Protected archivelog

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

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

1. Mount the image from the source database server to recover. In the Mount window, provide a mount point 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 from the IVGM](#) on page 45.

The database backup image will be mounted at the mount point: /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 IBM InfoSphere 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 archive log folder from IBM InfoSphere 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 IBM InfoSphere

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

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

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 from the IVGM](#) on page 42.

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 IBM InfoSphere 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 IBM InfoSphere mounted ASM disk group to RMAN

```
rman> run { catalog start with '+acttestdg/<db name>/datafile' noprompt;  
catalog start with '+acttestdg/<db name>/archivelog' 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 IBM InfoSphere-Protected archivelog

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

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

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 from the IVGM](#) on page 42.

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 IBM InfoSphere 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 archive log folder from IBM InfoSphere 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.

10 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 InfoSphere VDP 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 67

[Recovering a Corrupt Database Block](#) on page 68

[Recovering Lost Control Files](#) on page 69

[Recovering an Oracle Pluggable Database](#) on page 70

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 [Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#) on page 41.
2. In the Mount window, provide a mount point 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 Point 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 from the IVGM](#) on page 45.
 - 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 from the IVGM](#) on page 42.
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>/archivelog' noprompt;  
rman>catalog start with '/<mountpoint_log>' noprompt; (If archive logs are protected by IBM InfoSphere)
```

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 IBM InfoSphere)
```

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](#)

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 8, Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#).
2. In the Mount window, provide a mount point 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 Point 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 from the IVGM](#) on page 45.
 - 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 from the IVGM](#) on page 42.

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 recovering lost control files:

1. Mount the database point-in-time snapshots detailed in [Chapter 8, Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#).
2. In the Mount window, provide a mount point 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 Point 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 from the IVGM](#) on page 45.
 - 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 from the IVGM](#) on page 42.
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 IBM InfoSphere 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 protected under ASM Disk Group, see [Mounting an Oracle Database Image Protected Under an ASM Disk Group for Data Access from the IVGM](#) on page 42.
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 IBM InfoSphere 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')";  
    }
```
5. Open the Pluggable database:

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

The Pluggable database is open for read and write.

11 Performing an Oracle ASM Switch and Rebalance via IVGM

You can protect an Oracle ASM instance either as an Oracle ASM diskgroup or as a filesystem. If the 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 traditional RMAN recovery method [Mounting an Oracle Database Image Protected Under a File System for Data Access from the IVGM](#) on page 45.

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 IBM InfoSphere 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 IBM InfoSphere storage.
2. **Rebalancing** moves the data to production storage from IBM InfoSphere storage. This is an online operation where data movement happens in the background by the Oracle ASM API.

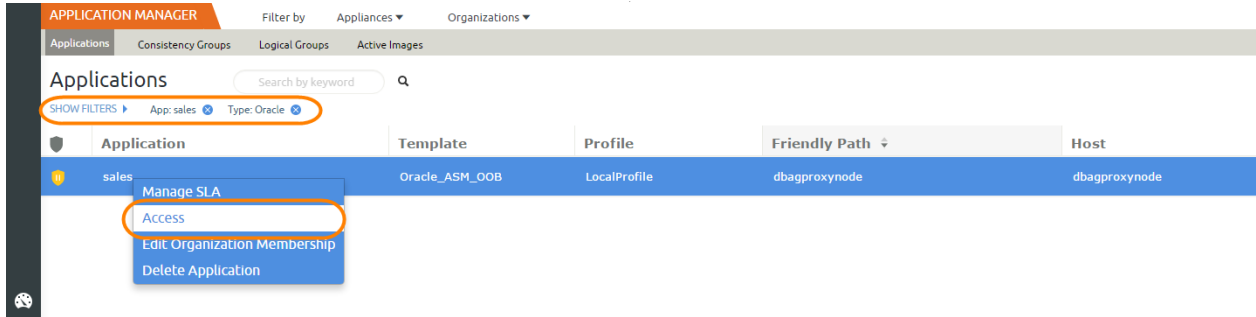
This procedure works for:

- Databases configured as: standalone ASM, single node RAC with ASM, and 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 IBM InfoSphere-presented ASM disk group retaining the ASM header information. A snapshot of the staging disk with ASM header information is taken.

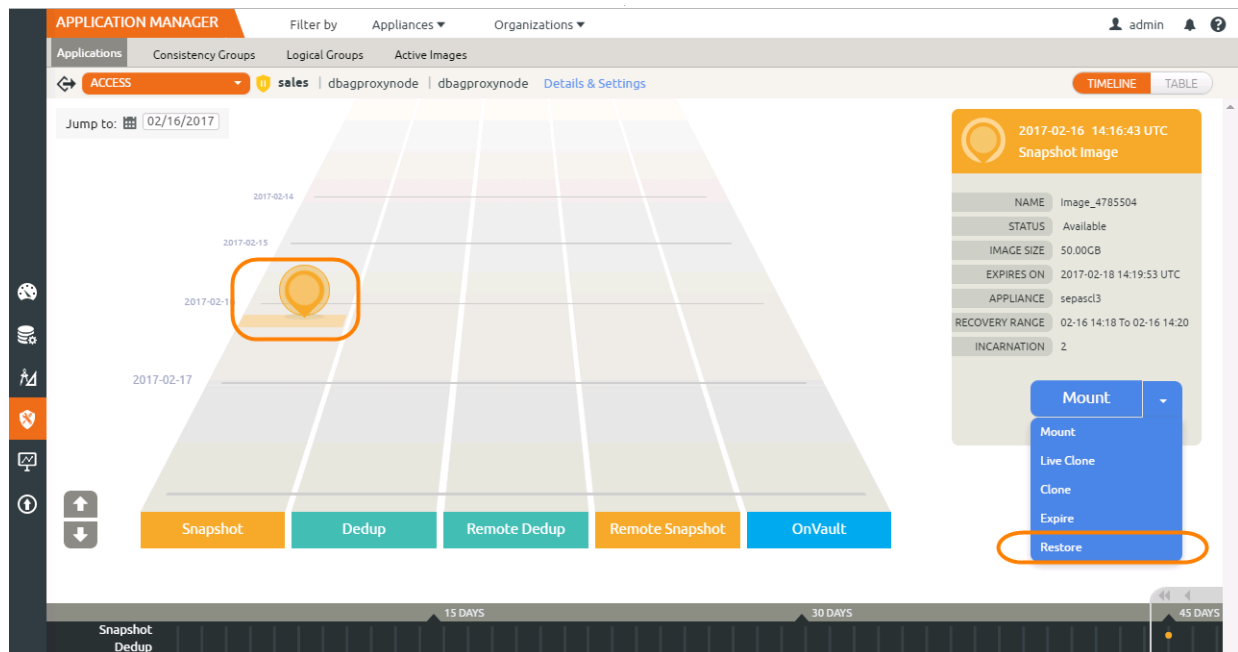
To instantly recover an Oracle ASM database from the IVGM:

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.
2. Open the IVGM to the Application Manager.
3. Use the filter feature to search for the desired database and click **Update Filters**.
4. Select the database to be recovered and click **Access**.



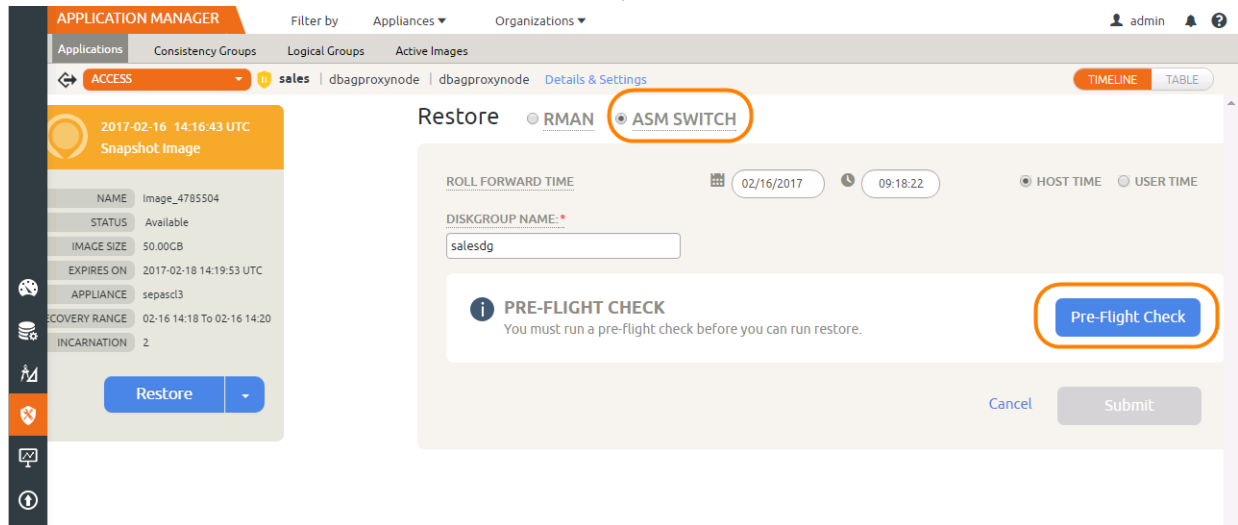
Right-Click the Selected Database and Click Access

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



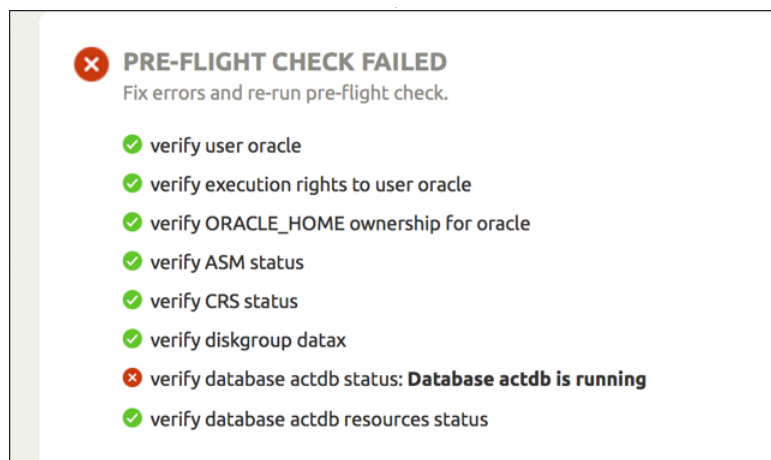
Select the Desired Image and Click Restore

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



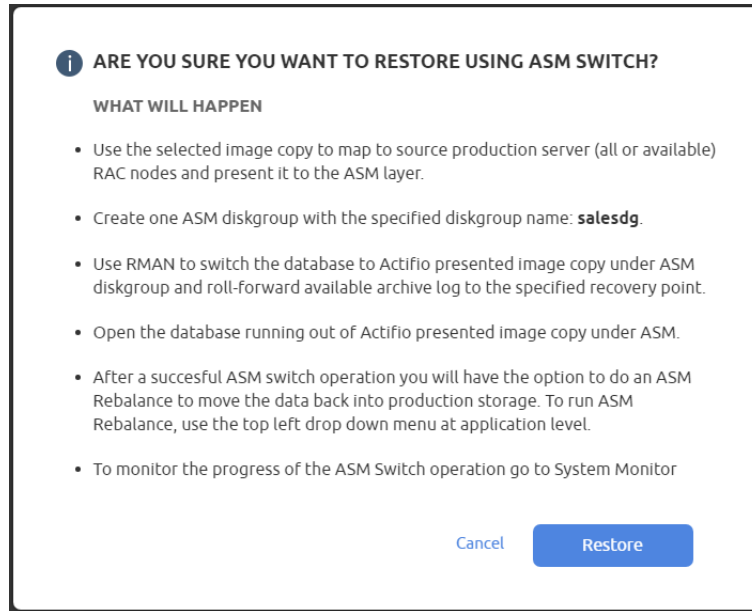
Select ASM Switch and Enter the Diskgroup Name

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



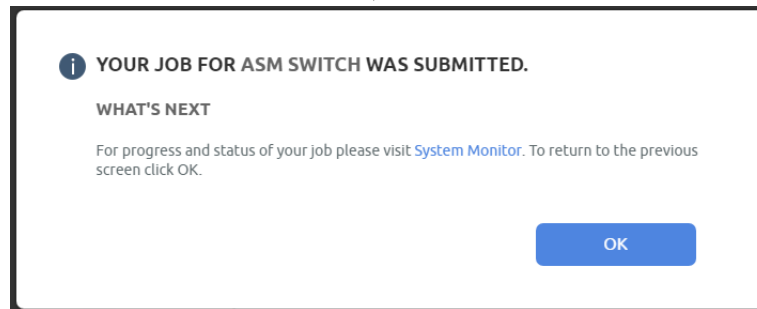
Address Any Errors, then Run the Pre-Flight Check Again

10. 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.



This Informational Screen Tells You What Will Happen

11. Click **Restore**. The job begins.



Waiting for the Restore/Switch Job to Complete

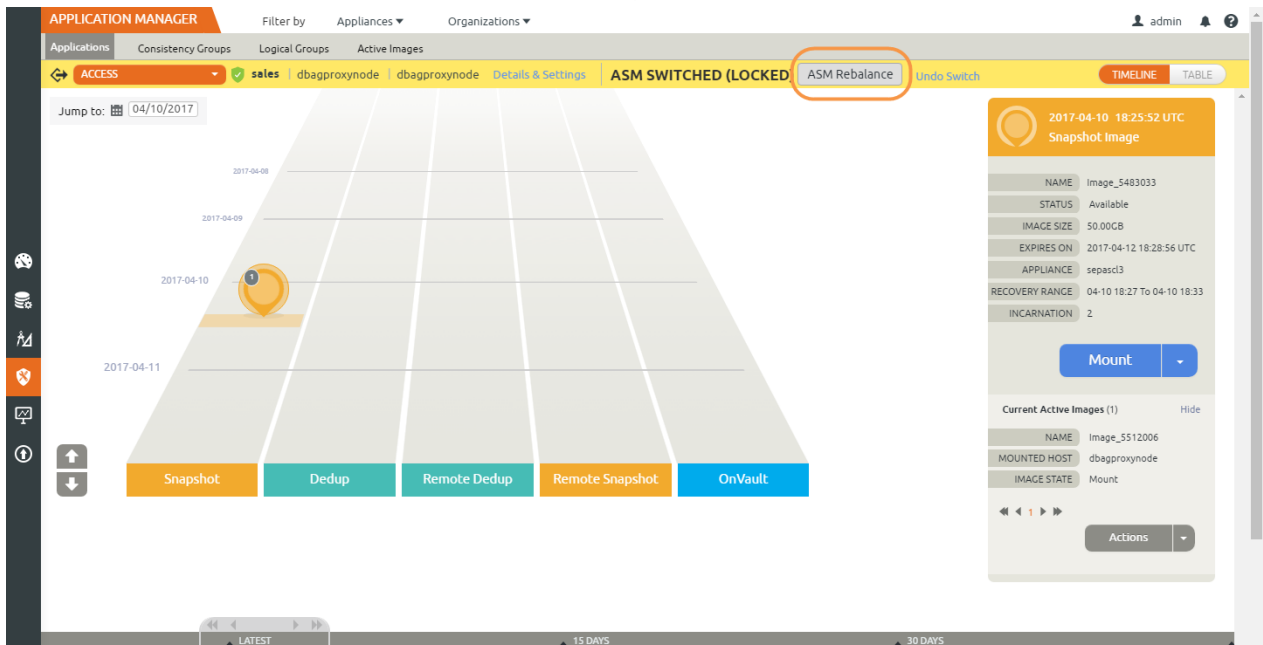
12. You can view progress on the System Monitor in another browser instance. The Job Type is Restore (ASM Switch).

	Status	Host	Application	Started	Ended	Policy	Type
2021	Running: 4%	oravmn2	cchssdb	04-13 23:05:...		Production to...	snapshot (DB...
1785b	Queued	bb6aix7	aixfs	04-13 23:20:...		dbSnap	snapshot
2006	Running: 97%	dbagproxyno...	sales	04-13 23:04:...		Production to...	Restore (ASM...
1782a	Queued	oravmn2	accwfdb	04-13 23:05:...		Production to...	snapshot

Viewing Job Progress in the System Monitor

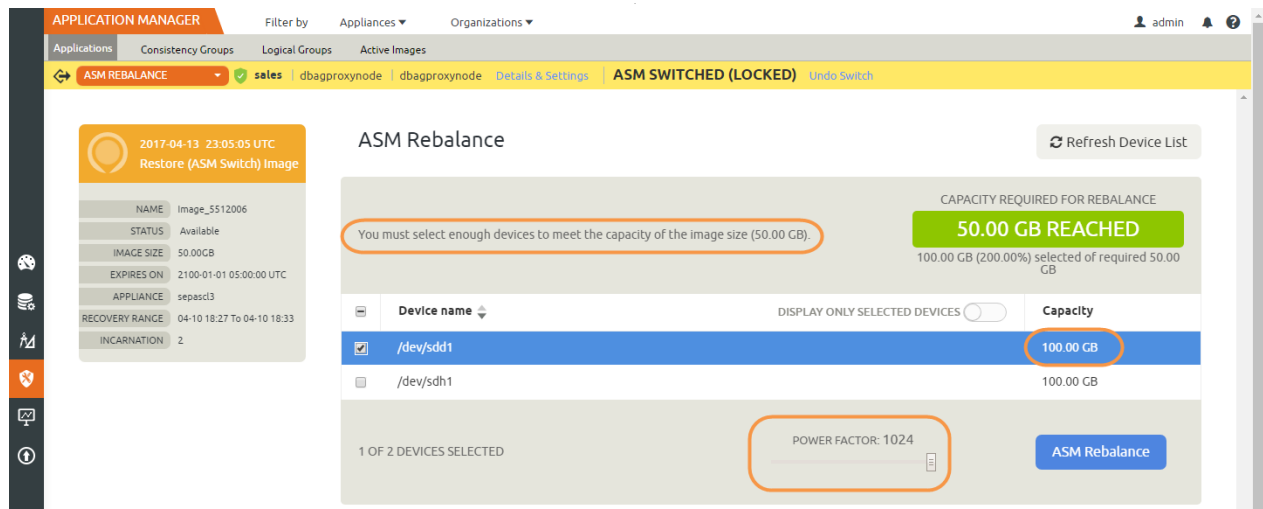
13. When the restore job has finished, go back to the original IVGM browser instance and click **OK**.

14. 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.



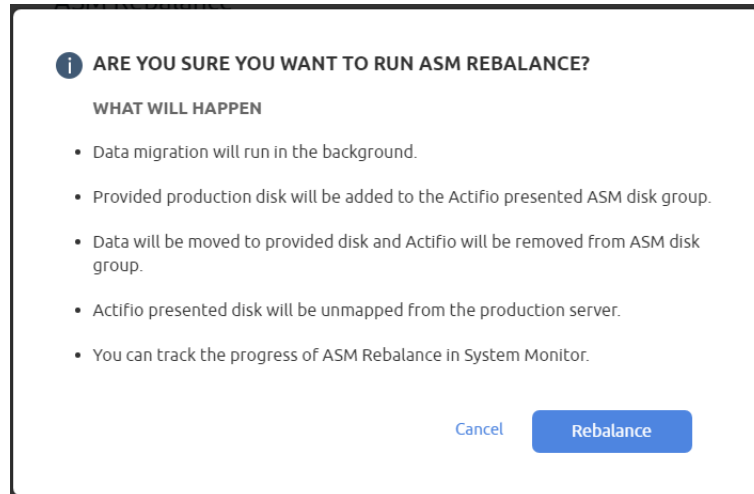
The Database is Switched and Ready to be Rebalanced

15. 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.
16. 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**.



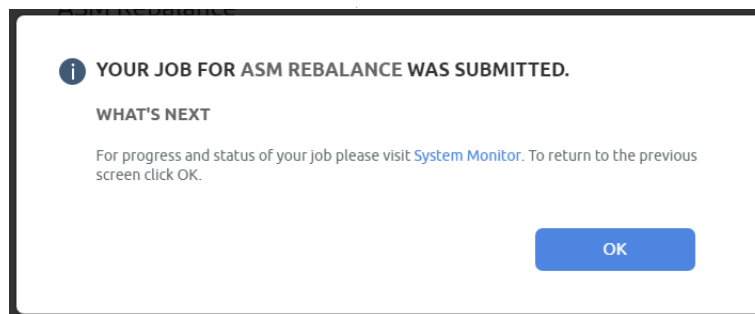
Selecting a Storage Device and a Power Factor

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



Submit the Rebalance Job

18. As with the switch job, you can view progress on the System 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.

12 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 Details and 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 Advanced Settings are detailed in [Application Details and Settings for Oracle Databases](#) on page 33.

This section includes:

[Preparing Oracle Protection in a Windows Environment](#) on page 77

[Identifying Database Instances On Windows](#) on page 78

[Backing Up an Oracle Database in a Windows Environment](#) on page 79

[Watch Script to Watch for Database Volumes Being Mounted](#) on page 80

[Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point if the archivelog is Not Protected through IBM InfoSphere](#) on page 81

[Recovering Oracle Databases in a Windows Environment Manually Using RMAN to a Scheduled Backup Point with Roll-Forward of IBM InfoSphere-Protected archivelog](#) on page 82

Preparing Oracle Protection in a Windows Environment

Before you can discover, protect, and mount 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_TXNTOILIVE 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_SHUTDOWNNTYPE 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

You can create a username and password inside the database and grant it rights to perform the backup:

1. Launch SQLPlus:

```
Set ORACLE_SID=orcl
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$block_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\orcl.bct';
```
4. Check if the log mode is set to Archive Log mode:

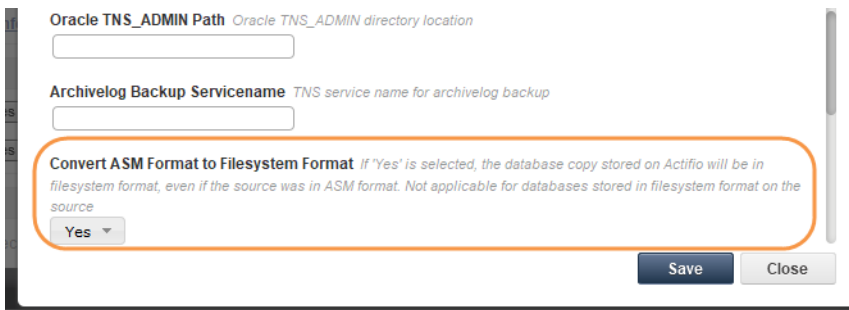
```
sql> archive log list
```

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 archivelog;
sql> alter database open;
```
5. Set the Application Advanced Settings. In particular:
    - o Open the IVGM and enter the database username and password (that we defined above) in the Application Advanced 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 Advanced Settings, see [Application Details and Settings for Oracle Databases](#) on page 33.



The screenshot shows a configuration window with the following fields and options:

- Oracle TNS\_ADMIN Path** (Oracle TNS\_ADMIN directory location): A text input field.
- Archivelog Backup Servicename** (TNS service name for archivelog backup): A text input field.
- Convert ASM Format to Filesystem Format**: A section with a description: "If 'Yes' is selected, the database copy stored on Actifio will be in filesystem format, even if the source was in ASM format. Not applicable for databases stored in filesystem format on the source." Below this is a dropdown menu currently set to "Yes".
- Buttons**: "Save" and "Close" buttons at the bottom right.

6. In IVGM, 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 IBM InfoSphere, and the Oracle processes running. Application-aware mounts are described in [Mounting an Oracle Database as a Virtual Application](#) on page 48.

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 IBM InfoSphere

To recover an entire Oracle database in a Windows environment:

1. Mount the database backup snapshot from IBM InfoSphere back to the Oracle server as detailed in [Chapter 8, Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#). See also [Recovering an Oracle Database Manually Using RMAN](#) on page 55, and [Recovering Tablespace and Data Files](#) on page 67.
2. Set the database environment and start the database in no-mount state using the parameter file from the IBM InfoSphere 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 IBM InfoSphere 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 IBM InfoSphere 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 from the VDP Desktop as detailed in [Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#) on page 41.
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 IBM InfoSphere-Protected archivelog

1. Mount the image from the source database server to recover (see [Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#) on page 41). Mount the backup snapshot from IBM InfoSphere back to the Oracle server as detailed in [Chapter 8, Accessing, Recovering, or Restoring an Oracle Database via the IVGM](#).
2. Set the database environment and start the database in no-mount state using the parameter file from the IBM InfoSphere 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 IBM InfoSphere 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 IBM InfoSphere 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.



# 13 Configurations for Using IBM InfoSphere in an Exadata Environment

InfoSphere VDP Appliances support capture and presentation of Exadata data on Oracle Linux-based database servers with a supported RHEL kernel.

- The InfoSphere VDP Appliance is connected over iSCSI in the network (not in data path). It may be necessary to install the iSCSI RPMs on each Exadata node. Doing so is non-disruptive and fully supported by Oracle.
- RMAN backup uses RMAN to directly write to copy data store presented by IBM InfoSphere as a file system or as an ASM Disk Group.
- IBM InfoSphere incremental-forever backup uses RMAN Incrementally Updated Backups, rolling forward image copy backups.

## IBM InfoSphere Capture of Exadata Data

- Data Capture Formats: Under ASM Disk Group or Under File System
- Backup is fully supported by Oracle for Exadata with HCC or Non-HCC data over iSCSI to IBM InfoSphere VDP Appliance.

---

**Note:** HCC Data cannot be deduplicated.

---

## Presentation of IBM InfoSphere Managed Exadata Data

Traditional recovery of database using RMAN is fully supported by Oracle for Exadata with HCC or Non-HCC data.

Application Aware mounts are also possible:

### Application Aware Mount

| Mounting Data from     | HCC Data                                                                                                                                                                                                                                                                                                                                                                         | Non-HCC Data    |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Exadata to Exadata     | Application Aware mount will be successful but the virtual copy cannot be used until the HCC data is uncompressed.                                                                                                                                                                                                                                                               | Fully supported |
| Exadata to non-Exadata | This can be achieved using a post script during the Application Aware mount.<br><br>This requires proper configuration of the staging disk size during protection setup to accommodate the uncompress of the HCC data during the Application Aware mount<br><br><hr/> <b>Note:</b> Query offload and Smart Scan of Exadata feature will not be available for Virtual copy. <hr/> | Fully supported |



---

# 14 Protecting SAP ECC/BW with an Oracle Database

---

There are four steps to protecting SAP ECC/BW:

1. [Protecting the Oracle Database](#) on page 85
2. [Refreshing the Database](#) on page 85
3. [After the Refresh, on the Target Oracle Database](#) on page 86
4. [After the Refresh, on the Target SAP Application Server](#) on page 86

## Protecting the Oracle Database

To protect the Oracle database, see [Chapter 7, 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 IBM InfoSphere application aware mount to refresh the target Oracle database as detailed in [Mounting an Oracle Database as a Virtual Application](#) on page 48.

## After the Refresh, on the Target Oracle Database

On Target SAP Database server:

1. Check the OPS\$<OS\_USER> in the database: (always enter <os\_user> in uppercase)

```
SELECT * FROM DBA_USERS WHERE USERNAME = 'OPS$<os_user>';
```

2. If the system does not return an entry, create the user:

```
CREATE USER "OPS$<os_user>" DEFAULT TABLESPACE <user_tsp>
TEMPORARY TABLESPACE PSAPTEMP IDENTIFIED EXTERNALLY;
```

---

**Note:** Ensure that the name of the OPS\$ user is specified entirely in uppercase.

---

---

**Note:** The table SAPUSER must occur in the system only once and it must be assigned to the user OPS\$<sid>ADM. Use the following query to check this:

---

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

3. If the system returns an owner <owner> other than OPS\$<sid>ADM, delete the relevant SAPUSER tables:

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

4. If the system does not return OPS\$<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.

# 15 Oracle RMAN Logs

This chapter details:

[Oracle Protection Logs on Linux/Unix](#) on page 87

[Oracle Protection Logs on Windows](#) on page 88

## Oracle Protection Logs on Linux/Unix

These are the logs that you might need to consult:

### 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:

### 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 protected 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 protected 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:

### Oracle Protection Logs: Standard Mounts

| Log                      | Location                                                                                                            | What's In It                                                           | What to Look For                      |
|--------------------------|---------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|---------------------------------------|
| Connector log            | C:\Program Files\Actifio\log<br>If the VDP 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:

### Oracle Protection Logs: Application Aware Mounts

| Log                                     | Location                                                                                                            | What's In It                                          | What to Look For                                                                      |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|---------------------------------------------------------------------------------------|
| Connector log                           | C:\Program Files\Actifio\log<br>If the VDP 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\<br><br>C:\Program Files\Actifio\act_scripts\oracleclone\          | Steps to configure target database post-roll-forward. | Any error with ORA-                                                                   |

---

# 16 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 VDP Connector.

You can also use the IVGM version, detailed in [Chapter 17, Oracle Database Management Using actAGMDBM](#).

This section includes:

[Installing and Configuring actDBM.pl on page 90](#)

[actDBM Commands on page 92](#)

- 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 111](#)

[Perl Examples of actDBM Usage and Results on page 115](#)

[RESTful API Examples of actDBM Usage and Results on page 116](#)

---

**Note:** *actDBM can be used with all InfoSphere VDP Appliances.*

---

## Installing and Configuring actDBM.pl

There are four steps to installing and configuring actDBM.pl:

[Installing actDBM.pl with the VDP Connector](#) on page 90

[Installing and Verifying the Five Required Perl Modules](#) on page 90

[Enabling and Verifying Port 443](#) on page 90

[Storing the Login Credentials for an InfoSphere VDP Appliance \(CDSconfig\)](#) on page 91

### Installing actDBM.pl with the VDP Connector

The actDBM script library is automatically installed on the Oracle host when you install the VDP Connector. It is available on the host under `/act/act_scripts/actdbm`. To install the VDP Connector, see **Connecting Hosts to IBM InfoSphere VDP Appliances** in your IBM InfoSphere Documentation Library.

The main script is a Perl script named `actDBM.pl`. You can see examples of actDBM usage and output in [RESTful API Examples of actDBM Usage and Results](#) on page 116. 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](#) on page 111.

### 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:

#### Installing the Five Required Perl Modules on an AIX System

| PERL Module          | Verify                                          | Install on Red Hat or CentOS Linux               | Install on SUSE or OpenSUSE Linux                   | Install on AIX                                     |
|----------------------|-------------------------------------------------|--------------------------------------------------|-----------------------------------------------------|----------------------------------------------------|
| JSON                 | <code>perl -e "use JSON"</code>                 | <code>yum install perl-JSON</code>               | <code>zypper install perl-JSON</code>               | <code>cpan&gt; install 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>      | <code>cpan&gt; install 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>      | <code>cpan&gt; install 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>         | <code>cpan&gt; install 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> | <code>cpan&gt; install 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 CDS IP> 443
```



```
Trying 172.16.15.200...
Connected to 172.16.15.200.
```

---

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

---

## Storing the Login Credentials for an InfoSphere VDP Appliance (CDSconfig)

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

```
perl actDBM.pl -type cdsconfig
--username <username>
--password <password>
--CDS <appliance ip>
```

### CDSconfig Parameters

| Parameter  | Use                                                                         |
|------------|-----------------------------------------------------------------------------|
| --username | VDP Desktop username to access the appliance. This is a required parameter. |
| --password | Password to access the appliance. This is a required parameter.             |

Repeat as needed for each affected InfoSphere VDP Appliance. After storing the appliance details, you can run actDBM.pl from the Oracle host to the configured InfoSphere VDP 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                                                                                        | <a href="#">listImageDetails</a> on page 93                                     |
| Return a list of applications on a host                                                                                                                                            | <a href="#">listApplication</a> on page 94                                      |
| Return a list of hosts known to an InfoSphere VDP Appliance                                                                                                                        | <a href="#">listDiscoveredHost</a> on page 95                                   |
| Back up a protected (source) database using traditional RMAN option                                                                                                                | <a href="#">backup</a> on page 96                                               |
| Restore and recover the protected (source) database using traditional RMAN option                                                                                                  | <a href="#">restore</a> on page 97                                              |
| Create or Refresh a copy of database with an Application Aware mount provisioning option depending on optype                                                                       | <a href="#">clone</a> on page 98                                                |
| Mount the backup image to specified target server                                                                                                                                  | <a href="#">mount</a> on page 103                                               |
| 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                                  | <a href="#">cleanup (Unmount and Delete an Image)</a> on page 104               |
| Run the pre-created DirectMountWorkflow or LiveCloneWorkflow based on --subtype option of runwf                                                                                    | <a href="#">runwf</a> on page 105                                               |
| Create a new LiveClone using the source image                                                                                                                                      | <a href="#">createliveclone</a> on page 107                                     |
| Refresh an existing LiveClone using the source image                                                                                                                               | <a href="#">refreshliveclone</a> on page 108                                    |
| Instantly switch and start the database to IBM InfoSphere mounted ASM Disk Group. The database will be up and running in its original configuration on your IBM InfoSphere storage | <a href="#">restoreASMswitch (Instant Oracle Database Recovery)</a> on page 109 |
| Migrate the database back to production storage or to new storage for a switched database running out of IBM InfoSphere image                                                      | <a href="#">restoreASMrebalance</a> on page 110                                 |

## 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>
```

### listImageDetails Parameters

| Parameters  | Use                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --dbname    | Name of database. This is a required parameter. The default value is localhostname.                                                                                                 |
| --hostname  | Name of source host. If not specified, host where script runs is used.                                                                                                              |
| --CDS       | The name or IP address of the InfoSphere VDP Appliance                                                                                                                              |
| --wait      | Wait flag to wait until the job has 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 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 116.

## 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/SqlServerWriter>]
[--protection <protectionState all/protected/disabled/notProtected>]
--CDS <CDS name|ip>
[--wait <yes|no>]
```

### listApplication Parameters

| Parameters   | Use                                                                        |
|--------------|----------------------------------------------------------------------------|
| --hostname   | Name of the source host. If not specified, host where script runs is used. |
| --apptype    | Either Oracle or SqlServerWriter                                           |
| --protection | One of four protection states: all, protected, disabled, or notProtected   |
| --CDS        | The name or IP address of the InfoSphere VDP Appliance                     |
| --wait       | 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 InfoSphere VDP Appliance](#) on page 115.

## listDiscoveredHost

To return a list of discovered hosts known to a specific InfoSphere VDP Appliance, use `--type listDiscoveredHost`

### Example

```
perl actDBM.pl --type <listDiscoveredHost>
--CDS <CDS name|ip>
[--wait <yes|no>]
```

### listDiscoveredHost Parameters

| Parameters | Use                                                                       |
|------------|---------------------------------------------------------------------------|
| --CDS      | The name or IP address of the InfoSphere VDP Appliance                    |
| --wait     | 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 InfoSphere VDP Appliance](#) on page 115.

## 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>]
```

### backup Parameters

| Parameter                | Use                                                                                                                  |
|--------------------------|----------------------------------------------------------------------------------------------------------------------|
| --dbname                 | Name of database. This is a required parameter. The default value is localhostname.                                  |
| --hostname<br>(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                    | InfoSphere VDP 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 117.

## 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>]
```

### restore Parameter

| Parameter     | Use                                                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --dbname      | Source database name                                                                                                                                                      |
| --OSusername  | OS Oracle owner                                                                                                                                                           |
| --OSpassword  | OS Oracle owner password. Optional, for backward compatibility with 6.2.x                                                                                                 |
| --image       | 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.                    |
| --backuplabel | 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. |
| --sourceHost  | Source hostname                                                                                                                                                           |
| --pointInTime | Recovery time in the format yyyy-mm-dd hh24:mi:ss.<br>Default: It will recover to all available archive logs.                                                             |
| --CDS         | InfoSphere VDP Appliance IP or name                                                                                                                                       |
| --wait        | Yes or No. The default value is “no”.                                                                                                                                     |
| --openDB      | True or False.<br>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 102.

```
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>]
```

### clone Parameters

| Parameter      | Use                                                                                                                 |
|----------------|---------------------------------------------------------------------------------------------------------------------|
| --optype       | Type of clone operation. This is optional parameter. If not specified, then a newclone operation will be performed. |
| --dbname       | Name of database. This is a required parameter.The default value is localHostname.                                  |
| --TargetDBname | Application aware mount clone database name.                                                                        |
| --OSUsername   | OS Oracle owner name                                                                                                |

## 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.                                                                                                               |

## clone Parameters

| Parameter        | Use                                                                                                                                                                                 |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --DBCharacterSet | Oracle database character set parameter value.                                                                                                                                      |
| --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            | InfoSphere VDP 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 117.

## 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>]
```

### refresh Parameters

| Parameter      | Use                                                                                                                                                                       |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --dbname       | Name of source database. This is a required parameter.                                                                                                                    |
| --TargetDBname | Application aware mount clone database name.                                                                                                                              |
| --image        | 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.  |
| --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.                                                               |
| --pointInTime  | 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. |
| --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.                  |
| --CDS          | InfoSphere VDP Appliance IP address                                                                                                                                       |
| --wait         | 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>]
```

### mount Parameters

| Parameter               | Use                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --dbname                | Name of source database. This is a required parameter.                                                                                                                                                                                                                                                                                                                                   |
| --sourceHost (optional) | Name of source host. If not specified, host where script is running is used.                                                                                                                                                                                                                                                                                                             |
| --targetHost (optional) | Name of target host. If not specified, host where script is running is used.                                                                                                                                                                                                                                                                                                             |
| --image (optional)      | Image name to use for mount the database. If not specified, the latest image will be used for database mount.                                                                                                                                                                                                                                                                            |
| --ASMDiskgroup          | 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.                                                                                                                                                                                                                                            |
| --RACNodeList           | 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. |
| --mountpoint            | This is required for snapshot image backed up on filesystem. This parameter is for databases protected under File System only to provide a mount point to mount a backup image under provided file system mount point.                                                                                                                                                                   |
| --CDS                   | InfoSphere VDP Appliance IP address                                                                                                                                                                                                                                                                                                                                                      |
| --wait                  | 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>]
```

### cleanup Parameters

| Parameter     | Use                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------|
| --dbname      | Name of source database for cleanup. This is a required parameter.                                   |
| --clonedbname | Cloned database name or application aware mount database name which is created using IBM InfoSphere. |
| --sourceHost  | Optional. Name of source host. If not specified, host where script is running is used.               |
| --targetHost  | Optional. Name of target host. If not specified, host where script is running is used.               |
| --CDS         | InfoSphere VDP Appliance IP address                                                                  |
| --wait        | 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 105

[LiveClone Workflow](#) on page 106

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 117.

### 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>]
```

#### 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                 | InfoSphere VDP 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>]
```

### 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.<br>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.<br>If refreshliveclone is YES and provisionliveclone is NO, then running the workflow will refresh only the LiveClone image.<br>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.<br>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                | InfoSphere VDP 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>]
```

### createliveclone Parameters

| Parameter             | Use                                                                                                                                                                                                                                                                                          |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --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.                                                                                                                                                                                                                       |
| --label               | Name of LiveClone image.                                                                                                                                                                                                                                                                     |
| --diskpool            | Name of disk pool for LiveClone image.                                                                                                                                                                                                                                                       |
| --sourceimage         | Name of source image to create a LiveClone image.                                                                                                                                                                                                                                            |
| --backuplabel         | Name of source image label instead of image name.<br>If <i>neither</i> sourceimage nor backuplabel are provided, then the latest image will be used to create the LiveClone.<br>If <i>both</i> sourceimage and backuplabel are provided then sourceimage is used and backuplabel is ignored. |
| --CDS                 | InfoSphere VDP Appliance IP address                                                                                                                                                                                                                                                          |
| --wait                | 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>]
```

### refreshliveclone Parameters

| Parameter             | Use                                                                                                                                                                                                                                                                                    |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| --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.                                                                                                                                                                                                                 |
| --label               | Name of LiveClone image.                                                                                                                                                                                                                                                               |
| --sourceimage         | Name of source image to refresh the LiveClone.                                                                                                                                                                                                                                         |
| --backuplabel         | Name of source image label instead of image name.<br>If <i>neither</i> sourceimage nor backuplabel is provided then latest image will be used to create the liveclone.<br>If <i>both</i> sourceimage and backuplabel are provided then sourceimage is used and backuplabel is ignored. |
| --targetimage         | Target Image Name to be created for LiveClone image. Optional.                                                                                                                                                                                                                         |
| --CDS                 | InfoSphere VDP Appliance IP address                                                                                                                                                                                                                                                    |
| --wait                | 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 IBM InfoSphere-mounted ASM Disk Group and start the database. The database will be up and running in its original configuration on your IBM InfoSphere 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>]
```

### 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              | InfoSphere VDP 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 IBM InfoSphere 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>]
```

### 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            | IBM InfoSphere CDS Appliance or VDP 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 backupup 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 point> --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 ('yyyy-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 InfoSphere VDP 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 InfoSphere VDP 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.

```
[av@av-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 L1 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 images in our deduplication pool.

```
[av@av-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:

```
[av@av-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 VDP Connector where to find ORACLE\_HOME in case we need to use a different folder.

```
[av@av-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 VDP Desktop and just call the workflow instead:

```
[av@av-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.



# 17 Oracle Database Management Using actAGMDBM

DBAs and developers can use actAGMDBM.pl to perform database access tasks using the command line interface. actAGMDBM 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. actAGMDBM.pl is installed on the database server automatically along with the VDP Connector. This section includes:

[Installing and Configuring actAGMDBM.pl](#) on page 120

[actAGMDBM Commands](#) on page 121:

- o [listAppliance](#)
- o [listImageDetails](#)
- o [backup](#)
- o [cleanup](#)
- o [clone](#)
- o [workflow](#)

[actAGMDBM.pl Script Template](#) on page 127

[RESTful API Examples of actAGMDBM Usage and Results](#) on page 128

## actAGMDBM and actDBM

This section is about the IVGM version of actAGMDBM. There is also an VDP Desktop version, detailed in [Chapter 16, Oracle Database Management Using actDBM](#). The two versions are nearly identical, except for where you run them from. In addition, actAGMDBM includes:

- an --AGM parameter used to indicate which IVGM should implement the command
- a listAppliance command used to list information about the InfoSphere VDP Appliances managed by an IVGM

actAGMDBM does not yet support all commands supported by actDBM. More are added with each service pack.

## Running actAGMDBM.pl

To run the actAGMDBM tool, cd to /act/act\_scripts/actAGMDBM folder and invoke `./actAGMDBM.pl`.

To run the script from another directory, include the script directory in the Perl library path. Use the `-I` switch in the argument: `perl -I /act/act_scripts/actAGMDBM/ /act/act_scripts/actAGMDBM/actAGMDBM.pl`

## Usage of actAGMDBM.pl

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

```
actAGMDBM.pl --type backup
```

The type options for actAGMDBM.pl are:

```
[root@zoravmn1 actagm]# perl actAGMDBM.pl
Usage: actAGMDBM
 --type
 <agmconfig>
 <listAppliance>
 <listImageDetails>
 <backup>
```

# Installing and Configuring actAGMDBM.pl

There are two steps to installing and configuring actAGMDBM.pl:

- Installing actAGMDBM.pl with the VDP Connector on page 120
- Storing the Login Credentials for an IVGM (agmconfig) on page 120

## Installing actAGMDBM.pl with the VDP Connector

The actAGMDBM script library is automatically installed on the Oracle host when you install the VDP Connector. It is available on the host under /act/act\_scripts/actagm. To install the VDP Connector, see **Connecting Hosts to IBM InfoSphere VDP Appliances** in your IBM InfoSphere Documentation Library.

The main script is a Perl script named actAGMDBM.pl. You can see examples of actAGMDBM usage and output in [RESTful API Examples of actAGMDBM Usage and Results](#) on page 128. 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 [actAGMDBM.pl Script Template](#) on page 127.

## Storing the Login Credentials for an IVGM (agmconfig)

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

```
perl actAGMDBM.pl --type agmconfig
--username <username>
--password <password>
--AGM <AGM ip>
```

agmconfig Parameters

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

Repeat as needed for each affected IVGM. After storing the appliance details, you can run actAGMDBM.pl from the Oracle host to the configured InfoSphere VDP Appliance.

## actAGMDBM Commands

actAGMDBM.pl permits these operations:

To	See
Store the encrypted IVGM credentials to permit API operations	<a href="#">Storing the Login Credentials for an IVGM (agmconfig)</a> on page 120
Return a list of applications and appliances managed by an IVGM	<a href="#">listAppliance</a> on page 121
Return a list of snapshot images with recovery range for the protected database application	<a href="#">listImageDetails</a> on page 122
Back up a protected (source) database using traditional RMAN option	<a href="#">backup</a> on page 123
Perform an unmount and delete operation on an image.	<a href="#">cleanup</a> on page 124
Create a new clone of a database.	<a href="#">clone</a> on page 125
Create a database workflow.	<a href="#">workflow</a> on page 126

### listAppliance

To return a list of applications and appliances managed by an IVGM, use `--type listAppliance`.

#### Example

```
perl actAGMDBM.pl --type <listAppliance>
[--AGM <AGM details>]
[--appname <appname>]
```

#### listAppliance Parameters

Parameters	Use
--AGM	IP address of the AGM.
--appname	The name of the protected application

## listImageDetails

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

### Example

```
perl actAGMDBM.pl --type <listImageDetails>
--appname <application name>
[--hostname <hostname>]
[--AGM <AGM ip address>]
```

### listImageDetails Parameters

Parameters	Use
--appname	Name of the protected application. This is a required parameter. The default value is localhostname.
--hostname	Name of source host. If not specified, host where script runs is used.
--AGM	IP address of the AGM to execute the command.

For a RESTful example of this command and sample output, see [listImageDetails](#) on page 128.



## backup

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

- db backup
- log backup
- dblog backup

### Example

```
perl actAGMDBM.pl --type backup
--appname <application name>
[--hostname <hostname>]
[--backuptype <db|log|dblog>]
[--backuplabel <backup bookmark label name>]
[--wait <yes|no>]
[--AGM <AGM ip address>]
```

### backup Parameters

Parameter	Use
--appname	Name of the protected application. 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.
--wait	Wait flag to wait until the job is completed. The default value is No.
--AGM	IP address of the AGM.

For a RESTful example of this command and sample output, see [backup](#) on page 128.

## cleanup

To perform an unmount and delete operation on an image, use --type cleanup.

This operation stops and removes any copy of a database running out of a mounted image and removes the file system mount, or drops any ASM disk group as part of cleanup.

### Example

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

### cleanup Parameters

Parameters	Use
--appname	Name of the protected application. This is a required parameter. The default value is localhostname.
--sourceHost	Name of source host. If not specified, host where script runs is used.
--targetHost	Name of target host.
--imagename	Name of the mounted image.
--targetAppName	Name of target cloned database, or mountpoint of target mounted file system.
--AGM	IP address of the AGM to execute the command.
--wait	Wait flag to wait until the job is completed. The default value is No.

## clone

To create a new clone of a database, use `--type clone`

### Example

```
perl actAGMDBM.pl --type clone --optype newclone
[--optype <newclone>]
--dbname <source database name>
[--remoteApp <no|yes>]
--TargetDBname <app-aware mount database name>
--OSusername <OS oracle owner>
--TargetOracleHome <Target Oracle Home>
[--listenerpath <Target Oracle Listener path>]
[--image <Image name>]
[--sourceHost <source hostname>]
[--targetHost <target hostname>]
[--imageType <Type of Image (snapshot, dedup, dedupasync, liveclone)>]
[--pointInTime <Recovery time 'yyyy-mm-dd hh24:mi:ss'>]
[--ASMDiskgroup <ASM diskgroupname>]
[--RACNodeList <list of RAC nodes seperated by colon>]
[--protectNewApplication <yes|no>]
[--totalmemory <Database memory size in MB>]
[--SGApct <SGA %>]
[--SharedPoolSize <shared pool size in MB>]
[--DBcacheSize <DB cache size in MB>]
[--DBrecoveryFileDestSize <DB Recovery file dest size in MB>]
[--Processes <Number of Processes>]
[--OpenCursors <Max Number of Cursors>]
[--DiagnosticDestination <Diagnostic Destination>]
[--tnsip <TNS Listener IP/scan IP>]
[--tnsport <TNS listener port>]
[--tnsdomain <TNS Domain>]
[--redosize <redo size in MB>]
[--recovery <yes|no>]
[--envvar <variable1:variable2>]
[--backuplabel <backup label ID>]
[--standalone <yes|no>]
[--mountpoint <filesystem mountpoint '/act/mnt'>]
[--DoNotUpdateTNSname <flag to update TNSentry yes|no>]
[--DoNotchangeDBid <flag to change DBID yes|no>]
[--template <child db protect template name>]
[--profile <child db protect profile name>]
--AGM <AGM ip|name>
[--wait <yes|no>]
```

## workflow

To create a database workflow, use `--type runwf`

```
perl actAGMDBM.pl --type runwf
```

```
actAGMDBM.pl --type runwf
--subtype [<directmount|liveclone>]
```

### Directmount Workflow

```
perl actAGMDBM.pl --type runwf --subtype directmount
```

#### Usage

```
actAGMDBM.pl --type runwf
--subtype directmount
--dbname <source database name>
--hostname <sourcehostname>
--wfname <workflow name>
--reprovision <yes|no>
[--image <Image name>]
--AGM <AGM name|ip>
[--wait <yes|no>]
```

### LiveClone Workflow

```
perl actAGMDBM.pl --type runwf --subtype liveclone
```

#### Usage

```
actAGMDBM.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>]
--AGM <AGM name|ip>
[--wait <yes|no>]
```

## actAGMDBM.pl Script Template

```
actAGMDBM Script Template:
/*****/
#!/bin/sh
set -x

echo "create AGM config: One time only at for setting up the config"
perl actAGMDBM.pl --type agmconfig --username admin --password password --AGM <AGM IP>

for setting up the backup
echo "***** starting the backup *****"
perl actAGMDBM.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 backuping database *****"
exit 1;
fi
echo "***** end backup *****"

echo "***** start list image details *****"
perl actAGMDBM.pl --type listImageDetails --appname <application name> --hostname <source
hostname> --AGM <AGM name|ip>

if [$? -gt 0]; then
echo "***** error running list image details *****"
exit 1;
fi
echo "***** end list image details *****"

echo "***** start list appliances *****"
perl actAGMDBM.pl --type listAppliance --appname <application name> --AGM <AGM name|ip>

if [$? -gt 0]; then
echo "***** error running list appliance *****"
exit 1;
fi
echo "***** end list appliance *****"

echo "***** completed all operation *****"
exit 0

/*****/
```

## RESTful API Examples of actAGMDBM Usage and Results

```
[root@spsyslinasmstand2 actagm]# perl actAGMDBM.pl --type agmconfig --username admin
--password password --AGM 172.16.200.107
user details are stored successfully
```

### listAppliance

```
[root@spsyslinasmstand2 actagm]# perl actAGMDBM.pl --type listAppliance --AGM
172.16.200.107 --appname hello
Application Name: hello
Appliance Name: spsky1
Appliance IPAddress: 172.27.11.116
Appliance Type: Sky
Host Name: orc122_stan_172.16.157.107
Application Name: hello
Appliance Name: leo
Appliance IPAddress: 172.27.11.51
Appliance Type: CDS
Host Name: ashok_rac1_172.16.157.111
Application Name: hello
Appliance Name: leo
Appliance IPAddress: 172.27.11.51
Appliance Type: CDS
Host Name: ashok redhat
```

### backup

```
[root@spsyslinasmstand2 actagm]# perl actAGMDBM.pl --AGM 172.16.200.107 --type backup
--appname hello --hostname orc122_stan_172.16.157.107
AppID is 939787
slt: ashok_dar1
sltID: 839882
slp: ashok_rem
PolicyID: 839885
839885
Backup JobName: Job_1089384
Backup Job Job_1089384 submitted Successfully
```

### listImageDetails

```
[root@spsyslinasmstand2 actagm]# perl actAGMDBM.pl --type listImageDetails --help
Usage:actAGMDBM --type listImageDetails
--appname <application name>
--hostname <source hostname>
--AGM <AGM name|ip>
```

### listImageDetails

```
[root@spsyslinasmstand2 actagm]# perl actAGMDBM.pl --type listImageDetails --appname hello
--hostname orc122_stan_172.16.157.107 --AGM 172.16.200.107
Database Name: hello
Host Name: orc122_stan_172.16.157.107
Snapshot Image and Recovery Range
Image Name: Image_1076435
RecoveryST: '2018-07-16 09:13:51' RecoveryET: '2018-07-17 09:13:44'
Snapshot Image and Recovery Range
Image Name: Image_1076048
RecoveryST: '2018-07-16 08:27:35' RecoveryET: '2018-07-17 09:13:44'
RecoveryST: '2018-07-15 15:18:09' RecoveryET: '2018-07-17 09:13:44'
DedupAsync Image and Recovery Range
Image Name: spsky1_Image_1080061
```