



ZFS STORAGE
APPLIANCE

January 2016

Oracle Snap Management Utility for Oracle Database v1.3.0 User Guide

Part No. **E39313-03**

Table of Contents	
Preface	vii
Typographical Conventions.....	vii
Access to Oracle Support	viii
Introduction	1
Overview	2
Supported Operations	3
Database Backup (Snap Backup)	4
Database Restore (Snap Restore).....	4
Database Recover (Snap Recover, Point-in-Time).....	5
Database Clone from Snap Backup (Snap Clone).....	6
Database Clone Copy from Snap Backup (Clone Copy)	8
Database Standby Clone (Data Guard Standby Clone).....	9
Database Clone from RMAN Backup (RMAN Clone)	11
Refresh Clone	13
Supported Systems and Configurations.....	14
Supported Storage Systems	14
Special Considerations with Clustered Systems	14
Supported Oracle Databases	14
Supported Database Layouts.....	15
Supported Database File Layouts	16
Database <code>oratab</code> File Entries	16
Database Layout and Types of Snap Backups	17
Using Oracle Intelligent Storage Protocol with Snap Management Utility.....	18
Database Configuration and Storage Types by Host Operating Systems.....	19
Supported Application (Database) Hosts	21
SMU Host Installation Requirements	21
Installing the Oracle Snap Management Utility	22
Installing the Oracle Snap Management Utility on an Oracle Linux Host.....	22

Installing the Oracle Snap Management Utility on Oracle Solaris Hosts	23
Installing the Oracle Snap Management Utility on Microsoft Windows Hosts	23
Configuring the Snap Management Utility Host and Database Host	23
Configuring the Network Port Settings	24
Verifying and/or Configuring Time Settings in the Database and Management Hosts	25
Configuring a Windows HTTPS Connection for the Snap Management Utility	26
Windows Remote Management Protocol and Service Settings for the Windows Database Host.....	27
Configuring Host System Name Services Recognized by the SMU Host.....	28
Starting and Stopping the Snap Management Utility on Host Systems	28
Starting Oracle Snap Management Utility on Linux Hosts	29
Starting Oracle Snap Management Utility on Oracle Solaris Hosts	29
Starting Oracle Snap Management Utility on Windows Hosts	29
Updating the Snap Management Utility on Host Systems	30
Updating Oracle Snap Management Utility on Oracle Solaris Hosts	30
Updating Oracle Snap Management Utility on Oracle Linux Hosts	31
Updating Oracle Snap Management Utility on Windows Hosts	31
Protecting Oracle Snap Management Utility's Data Files.....	31
Restoring a Clone Profile in the SMU Database	32
User Permissions Requirements for Accessing Operations.....	33
Using SMU Delegation Tools	34
Sample sudo policy file configuration format (/etc/sudoers)	38
Accessing the User Interfaces	38
Command-Line Interface.....	38
Accessing the Command-Line Interface Using SSH.....	39

Accessing the Command-Line Interface Using WinRS.....	39
Authenticating with WinRS	39
Accessing and Authenticating Using the Browser User Interface	41
Navigating the Browser User Interface	42
Managing Accounts Using the BUI	45
Adding a new host account	47
Modifying a host account.....	47
Testing an application host account.....	48
Deleting a host account.....	48
Adding a new storage account.....	48
Testing a storage account	49
Modifying a storage account	49
Deleting a storage account.....	49
Managing Applications Using the BUI.....	50
Enrolling a new application for snap backups	51
Testing an application.....	52
Modifying an application.....	52
Browsing application details	52
Removing an application	52
Importing an RMAN backup image	53
Refreshing a clone.....	55
Deprovisioning an application (clone only).....	56
Managing Administration Using the BUI	56
Adding a new user from the Users tab.....	57
Modifying a user entry from the Users tab	57
Deleting (Remove) a user from the Users tab.....	57

Adding a new notification subscription	58
Modifying a notification subscription from the Notification tab	58
Deleting a notification subscription from the Notifications tab.....	59
Managing status polling, and task and activity log displays in the General Settings tab.....	59
Operating and Managing Snap Backups and Clones in the BUI	61
Creating a New Snap Backup	63
Renaming a snap backup.....	64
Cloning a Snap Backup.....	65
Rolling Back (Restoring) a Snap Backup	71
Recovering a Snap Backup (Recovering a Database) from a Point in Time	72
Refreshing a Clone (Updating to Current Source Database)	75
Deleting a Snap Backup.....	75
Running Simultaneous Snap Operations	76
Managing Snap Backup Initiation in the Schedules Tab	76
Adding a schedule	77
Editing a schedule	77
Deleting a schedule	78
Monitoring Snap Operation Tasks.....	78
Canceling currently running tasks	78
Deleting task history	78
Managing Account Settings Using the Account Settings Tab	78
Managing Activity Logs Using the BUI	79
Filtering Activity Logs.....	81
Exporting Activity Logs	82
Purging Activity Logs	83
Using the Command-Line Interface	83
Managing Accounts.....	86
Managing Activities	90

Managing Alerts	91
Managing Backups.....	93
Managing Certificates	94
Managing Keys	94
Managing Schedules.....	96
Managing Tasks.....	98
Managing Users	105
Troubleshooting – General Information.....	108
SMU Restrictions Quick Reference Table.....	109
Appendix A: References	112
Appendix B: Glossary.....	114
Appendix C: Icon Set for the SMU Browser User Interface.....	117
Appendix D: Cloning Wallet Files for an Encrypted Database	120
Appendix E: Troubleshooting Common Issues	123

Preface

This document is written as a basic guide for database administrators to install, configure and administer Oracle's Snap Management Utility for Oracle Database and presumes a working knowledge of database administration concepts. Further helpful references for the components for which this utility have been designed are included in the References section at the end of this document.

Additionally, for the latest updates and helpful information, users are encouraged to visit the My Oracle Support web site at: <http://support.oracle.com>.

Typographical Conventions

The following typographical conventions are used throughout this document.

Convention	Meaning/Use
<i>italics</i> <italics>	Specifies a variable whose value is supplied by the user. The second example shows less than and greater than signs which are often used to indicate a variable but are not typed.
monospace	Indicates commands, filenames, directory paths, and executables, and screen code output.
bold monospace	Indicates a command that the user types.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Oracle Snap Management Utility for Oracle Database and Accessibility

Oracle Snap Management Utility for Oracle Database (referred to as SMU throughout this document) provides operability and documentation features through its browser user interface (BUI) that are designed to enhance access for visually impaired users.

The Online Help function (labeled "Help") of Snap Management Utility that is accessed through the BUI provides an html-based display of contents of this user guide. Labels and icons, and some selectable options in menus, have identifying pop-up labels that display when a mouse pointer is hovered over them.

Next to the Help label, the Accessibility label, when selected, displays an Accessibility Preferences window in which users can select one of three customizations for the screen presentation of SMU's user interface:

- I use a screen reader
- I use high contrast colors
- I use large fonts

Introduction

Oracle Snap Management Utility (SMU) for Oracle Database uses the snapshot technology of the Oracle ZFS Storage Appliance to allow database administrators (DBAs) to back up, restore, clone, recover, and provision from Oracle database backups that are hosted on the appliance. Snapshots are read-only virtual copies of a dataset that can be used for two primary purposes: rolling back the dataset to an earlier point in time and creating a new copy of the dataset. In both cases the snapshots and clones share common data blocks with the original dataset. This means that snapshot and clone creation are extremely time and space efficient; creating copies is instantaneous and does not require actually copying any data. Instead, existing data is referenced from the snapshot.

These kinds of on-disk backups created using Oracle Snap Management Utility facilitate quick restore and cloning that is particularly useful for test and development functions. SMU has extended functionality with its ability to provide a database instance from an Oracle Recovery Manager (RMAN) backup resident on an Oracle ZFS Storage Appliance. SMU is not intended to replace Oracle RMAN or other storage backup processes. It is recommended that users supplement the SMU backups with tape or disk backups of the databases on the Oracle ZFS Storage Appliance.

SMU is designed for installation and use from the DBA's management station (terminal or Web browser). This design removes the need for the DBA to coordinate backup, restore and cloning activities with the storage administrator. SMU combines standard host-side processing with storage operations to create seamless functions that simplify the most common administrative tasks the DBA performs. The Snap Management Utility for Oracle Database allows the DBA to directly harness the powerful features of the Oracle ZFS Storage Appliance.

Overview

Oracle Snap Management Utility for Oracle Database leverages Oracle ZFS Storage Appliance technologies to simplify and automate Oracle Database backup, recovery, restoring, and cloning from backup. SMU benefits the database administrator's task performance with the following features:

- Database and related resource accounts management:
 - Organize and retain account information used to access database resources including the database and storage hosts. SMU operations are tightly integrated so that when the utility executes an operation, it synchronizes its commands among the database, database host, and storage host for zero data loss and consistency.
 - Automatically identify which Oracle ZFS Storage Appliance shares are used by the database, removing the need for the administrator to know the underlying share names.
- Database backup:
 - Create virtually unlimited ZFS snapshot-based backups (limited only by physical system capacity).
 - Schedule backups on a recurring basis.
- Database restore and recover: Reduce the time to restore or recover a database using rollback.
- Database cloning:
 - Quickly create a database clone for testing or development purposes using the Oracle ZFS Storage Appliance clone feature.
 - Create a full clone copy that is independent of the source database for data protection.
 - Efficiently create a standby clone that is synchronized to the primary database for high availability/disaster recovery functions.
- Database provisioning from Oracle Recovery Manager (RMAN) backup: Similar to database cloning, but provision a database instance from an RMAN image copy.

SMU has the following features:

- Supports single instance, Real Application Clusters (RAC), and RAC One Node databases.
- Supports standard databases (non-CDB) as well as multitenant, or container, databases (CDBs) (with Oracle Database 12c), which can hold multiple user-created pluggable databases, for backup, clone, restore, recover, and provisioning operations.

- Supports both filesystem (specifically, Network File System [NFS] or direct NFS [dNFS]) and Automatic Storage Management (ASM) (specifically, iSCSI LUN) storage type databases because SMU leverages the unified nature of the Oracle ZFS Storage Appliance.
- Creates a variety of types of space-efficient backups and clones using ZFS technology.
- Provides the ability to provision new database instances from RMAN backups.

Snap Management Utility for Oracle Database provides robust functionality through both a command-line interface (CLI) and Browser User Interface (BUI). The BUI's navigation tree, operations wizards, menus and icons provide operational details and cues for the existing and required components upon which SMU operates.

Whether using the BUI or CLI, your setup of SMU requires the following basic tasks:

1. Set up accounts. SMU requires a host account for the source Oracle Database instance upon which you may perform operations, and a storage account for the Oracle ZFS Storage Appliance upon which the Oracle Database (host account) resides. Setting up accounts provides the needed permissions as well as configuration to access the host and storage.
2. Enroll applications. Applications are instances of the Oracle databases. Within the accounts, these applications (Oracle databases) are "enrolled" so that the account's administrator can perform tasks on the application (the database instance).
3. Manage administration. Enroll users who may access SMU, and set up their notifications for task events. Also set up monitoring intervals and displays. Likewise, users and notifications can be subsequently deleted or modified, and monitoring settings can be altered.
4. Set up activity logs.

Once these housekeeping activities are accomplished, SMU is set up to perform operations.

Supported Operations

SMU supports the following common database administrative tasks. Each of these tasks uses ZFS technology to eliminate the need to make physical copies of the database blocks, thus saving time and allowing the storage to be used efficiently.

Important: Be sure to check the SMU Restrictions Quick Reference Table (near the end of this document) as well as Supported Systems and Configurations for listings of applicable restrictions for each operation.

Important: Users should not attempt to administer snapshots created by SMU from the Oracle ZFS Storage Appliance browser user interface or CLI. SMU will not use snapshots created from the native appliance interfaces. SMU will only use the snapshots created by its own snap backup operations.

Database Backup (Snap Backup)

SMU can be used to create on-disk backups of databases. The backups are based on ZFS snapshots. This means the backups can be taken quickly and are space efficient. Backups can be created manually or automatically. Automatic backups can be scheduled to occur on a recurring basis. Automatic backups can also employ a retention policy based on the number of backups to retain at a time.

Three types of backups, offline (also referred to as cold backup), online (also called hot backup), and standby are supported.* Offline backups are backups taken when the database is shut down. The software will shut down the database temporarily and then restart it after taking the snap backups. Online backups are taken when the database is placed into backup mode while remaining online. Standby backups are database backups that can be used to create standby clones, which serve as copies of the database that stay in sync with that primary database.

Online backups take snap backups of the database shares in a particular order and in between changing the database mode and archiving the current logs. The general steps taken during an online backup are:

1. Place the database into backup mode.
2. Snapshot the datafile shares.
3. Take the database out of backup mode.
4. Archive the current logs.
5. Snapshot the archived log shares.
6. Snapshot any other database shares.

When creating an offline backup of a clustered database, any database nodes that have been stopped but not disabled will be restarted at the end of the backup task when SMU restarts the database.

Database Restore (Snap Restore)

Using the rollback feature, SMU can restore a database from an on-disk backup, whether it is an offline, online or standby snap backup. Rollbacks allow you to revert a dataset back to a point in time without having to copy or delete any data.

However, because of the way rollback works, when you perform a restore from a snap backup, any snap backups that were created after the snap backup you are restoring to are deleted. Additionally, if any of the snap backups have been cloned, the restore and rename operation cannot be performed and will fail. SMU will block the rename operation on the snap backup that has clones. These newer clones should first be deprovisioned to avoid either blocked operations or deletions.

To restore from an offline snap backup, SMU will shut down (abort) the database if it is running, roll back the shares to the specified snapshot, and then restart the database.

To restore from an online snap backup, SMU will shut down (abort) the database if it is running, roll back the database shares, start up (mount) the database, create a new controlfile, recover the database, and then open the database and reset the logs.

Database Recover (Snap Recover, Point-in-Time)

While a Database Restore operation restores the database to the point-in-time of the backup, Database Recover extends SMU functionality to restore the database to a point in time in between backups. This capability is built upon the Oracle Database incomplete recovery feature. Incomplete recovery, also known as point-in-time recovery, uses a backup created prior to the time to which you want to recover and then applies a series of archived logs in the correct order from the backup point to the recovery point to bring the database to the desired state.

For more information on Oracle Database's incomplete recovery operations, please refer to:

https://docs.oracle.com/cd/B19306_01/server.102/b14220/backrec.htm#i1006524

Recover offers three options for specifying the recovery point:

- Time-based recovery – Recovers the data up to a specified point in time.
- Change-based recovery – Recovers until the specified system change number (SCN).
- Log sequence recovery – Recovers until the specified log sequence number.

Based on the specified recovery point, the software will automatically select the nearest suitable backup for that recovery point. The selected backup will be at some point prior to the desired recovery point and can be either an online or offline backup, as long as the archive logs and data files are in separate shares and the database is in archivelog mode.

Recovering to the user-specified point in time requires that all needed archived logs from the backup point to the recovery point be present and available in the archive log destination (usually the fast recovery area, or FRA). The recover operation only rolls back the data file shares of the database. It does not roll back any of the other database shares such as the FRA or archived log shares.

Recovery is essentially a combination of media restore and media recovery. All backups created after the backup selected for recovery are deleted during the recover operation; just as when rolling a share back to a snapshot, all subsequent snapshots are automatically destroyed.

Time-based recovery requires that the database node(s) system clock be set to the correct time and node(s) operating system set to the correct timezone. When a log is archived, the current local time is written to the log. RMAN will use this time when processing archive logs and deciding which archive logs to apply. If the operating system clock or timezone is set incorrectly, then time-based recovery will fail, resulting in either an RMAN-06555 error (must restore from earlier backup) or a recovery point after the desired point in time. This time syncing

requirement also holds true for clustered databases. NTP/time server should be configured on each database host and also the management host where SMU runs.

For more detailed information on how Oracle Database works with times, timestamps and timezones reference the Doc ID 340512.1, in the My Oracle Support (MOS) Knowledge Center at <http://support.oracle.com>.

Database Clone from Snap Backup (Snap Clone)

SMU can be used to clone a database, creating a new primary database. This database is created using ZFS clone technology. ZFS clones are thin-cloned datasets that share common data blocks with the original dataset. This allows the clones to be space efficient and created very quickly.

The clone databases that SMU creates are sized and configured based on the source database. All of the metadata stored with the backup being cloned is used to construct an identical clone of that database.

The following process is used to create a snap clone:

1. Snap backups (share snapshots) are cloned.
2. Clone shares are mounted (if filesystems) or mapped as SCSI disks (if LUNs) on the target database host(s).
3. SMU creates a parameter file for the clone database using input from the user and source parameters stored with the snap backup.
4. SMU starts the clone database.
5. SMU creates a new controlfile for the clone database.
6. SMU recovers the database if necessary.
7. SMU opens the database and resets the logs.
8. SMU recompiles the schema objects. This step allows for creating clones in an environment with a different operating system type from the source database, as long as the supported systems are of the same endian* architecture. (*See Glossary for further information.)
9. SMU adds new temp files.
10. SMU adds or removes undo tablespaces, depending on the number of nodes the clone database is using.

Clone filesystem shares are placed in the same project as the source filesystem shares, and clone LUNs are placed in the same initiator and target groups as the source database shares.

When creating a clone database that uses filesystem storage type, SMU will mount the clone shares on the target database host/nodes after they are created on the Oracle ZFS Storage Appliance. In order to mount the shares, SMU must decide which network address to use based on what network interfaces and addresses are configured on the Oracle ZFS Storage Appliance. SMU uses the following algorithm to determine which network address should be used to mount the shares:

1. Look for non-administrative Oracle ZFS Storage Appliance network addresses that are on the same subnet as the database host.
2. Look for non-administrative storage appliance network addresses that are reachable.
3. Look for any storage appliance network address that is on the same subnet as the database host.
4. Look for any storage appliance network address that is reachable.

When cloning a snap backup of an ASM (over iSCSI LUNs) database to a new host, initiators of the new host must be configured in the appropriate iSCSI initiator group and the new host must be already logged in to the appropriate iSCSI target group. When the clone LUNs are created, they are mapped in the same iSCSI initiator and target groups as the original iSCSI LUNs. After the clone LUNs are created, SMU makes the target host scan its SCSI bus to discover the clone LUNs. The clone LUNs will only be visible when the host is logged in to the iSCSI target group.

Linux-based hosts must have the `sg3_utils` package installed. SMU will use the `scsi-rescan` command contained in this package to add and remove clone disks.

SMU provides several options when creating a snap clone:

- Create a database clone on the same host as the original or source database.
- Create the clone on another host (as long as it has the same architecture and Oracle Database software as the source database host). The target host(s) must be using the same type of operating system (endian order; for example, Oracle Linux x86 and Oracle Solaris x86), clusterware (cloud/grid control) version and database software and patch levels as the source database node(s).
- Create a clone that is single instance or clustered. Note that creating a clustered database clone requires that the target host be a node in an existing RAC configuration. SMU will perform the necessary conversion steps on the backup used for cloning to make it suitable for a single instance or RAC (or RAC One Node) environment.

Important: Only deprovision (delete) a clone database that was created using the SMU from the SMU. Do *not* delete the database outside of SMU using DBCA (Database Configuration Assistant, a database software tool), for example. SMU attempts to connect to and query the clone database in order to deprovision it. When SMU cannot connect to the clone database, SMU uses information on the clone profile file that is captured and stored during

clone creation. Do not manually delete clone profile files in order for clone deprovisioning to work when the clone database is not running.

Important: For security reasons, the newly created clone does not automatically inherit the sysdba account from the source database. Only the sys and system users will have automatic access to the new clone.

Important: Before creating a clone that is dNFS file type, the target Oracle home must first be configured for that type. Note that SMU does not update the `oranfstab` file on Linux or Oracle Solaris hosts to allow access to the cloned database using dNFS with specific parameters such as multiple paths or NFSv4 access. If the cloned database is to be accessed using dNFS, the `oranfstab` file should be manually updated. For more information about `oranfstab` and dNFS with Oracle Databases, refer to the appropriate Oracle Database Installation Guide for your release and operating system.

Database Clone Copy from Snap Backup (Clone Copy)

SMU can duplicate a database, using the clone copy feature, based on the bit-by-bit copies of data shares from a source database. Unlike a thin clone, which does not contain its own data files but rather refers to them in the source database, a clone copy holds its own data shares independent of the original database. Consequently, a clone copy operation can also create a clone on pools of different storage from the source database while thin clone only permits creation of a new clone on the same storage pool.

A clone copy can be created from either online snap backup or offline snap backup. All snap operations except refresh clone can be performed on a clone copy, including snap backup, restore, recover, and deprovision. Refresh clone cannot be performed, as a clone copy becomes independent of the source database once it is created.

Clone copy is based on the remote replication features of the Oracle ZFS Storage Appliance. Remote replication requires the setup of a peer storage replication target that can receive and store data replicated from the source storage holding the source database. This replication target configuration must be properly performed prior to a clone copy operation or the operation will fail. Note that replication target configuration on the Oracle ZFS Storage Appliance requires a root password; therefore, a storage administrator may be required for this configuration if the SMU user does not have root privileges. Remote Replication service must be enabled on the affected Oracle ZFS Storage Appliances. Remote replication may be set up to another Oracle ZFS Storage Appliance, another storage pool on the same Oracle ZFS Storage Appliance, or to the same pool on the Oracle ZFS Storage Appliance.

Clone copy functions on both single head and clustered Oracle ZFS Storage Appliance configurations, and supports databases with either NFS or ASM/iSCSI LUN storage configuration. Clone copy can create a clone database for Single Instance Database, Real Application Cluster (RAC) Database or RAC One Node configurations.

Important: For ASM/iSCSI LUN databases, the iSCSI initiators of the clone target database host must be configured in the same iSCSI Initiator Group to which the clone source database host is mapped. This configuration must be performed prior to clone copy operations. Clone copy of an ASM/iSCSI LUN database to a different storage appliance is not supported.

The clone copy process overview is as follows, with actions 1 and 2 as prerequisites to the clone copy initiation:

1. Configure replication targets on the source storage to enable replication to the target storage (may require a storage administrator because of the need for root access to storage).
2. In SMU, create a backup snapshot of the database shares from which a clone copy is to be created.
3. Create a replication action for each project that contains database shares using a configured replication target pointing to the target (receiving) storage. The replication action is a configuration object on a source storage specifying a project or share, target storage, or perhaps a synchronization policy. Its counterpart on the target side is the replication package.
4. Initiate a send update operation for the replication action. This send update is also called a synchronization update, in which the source data is transmitted to the replication target.
5. Wait for the update operation to complete.
6. Sever the replica (replication package on the target storage). Severing the replica requires changing the project name, mountpoint, and so on to resolve name conflicts.
7. Roll back the replica to the snapshot from which the clone database shares are originated from.

Since clone copy produces a complete and independent replica of the source database, the resulting clone database consumes the same amount of the storage capacity as the source storage and will require adequately sized pools on the target storage.

For more detailed information on the remote replication service of the Oracle ZFS Storage Appliance, see the *Oracle ZFS Storage Appliance Administration Guide* listed in the References section of this document.

Database Standby Clone (Data Guard Standby Clone)

SMU can be used to set up a data guard configuration. Data Guard is the database feature that provides high availability and protection through a simple, fast and reliable one-way replication of the database. A data guard configuration consists of a group of databases. One database acts as the primary database and the other databases, which are physical copies of the primary, as standbys. As changes are made to the primary database, the redo is streamed to each standby in the group and archived with and applied to their copy of the database. If the primary database goes down, a transition occurs where one of the standbys in the group assumes the primary role to allow all database clients to continue their work.

The process of creating a standby database involves first seeding the database from a backup of the primary database. This seeding is very similar to the process for creating a clone database. In fact, a standby database starts out as a clone database. Once the clone is created, then the data guard broker is used to add the clone database to a data guard configuration. During this process the database is transformed into a physical standby of the primary database and starts receiving the changes from the primary database and archiving and applying them to its copy of the database files.

Normally seeding a standby database requires many manual steps of duplicating the primary database, copying the additional required files, setting up the standby database instance from a new Oracle home, configuring the data guard parameters, making the duplicated files available in the target environment, and so on. SMU performs all of these steps to greatly simplify the whole process.

SMU can create a database clone that is designated as a Data Guard standby database, so that it can stay in sync with a primary database in real time using the Oracle Database Data Guard feature. The standby candidate database is first cloned. Then this clone database, rather than being opened, is placed in a special mode where it remains in a recovery state, and receives updates in the form of redo logs that are streamed over the network from the primary database.

The standby clone process overview is as follows. Steps 1 to 4 are prerequisites to standby clone initiation and are performed by the user:

1. On the storage containing the database that is to become the data guard primary, configure the Oracle ZFS Storage Appliance that will contain the data guard standby as a remote replication target. Enroll the standby storage with a storage account in SMU. Enroll the database that will become the data guard primary as an application in SMU. Enroll both the primary and standby hosts with host profiles.
2. In SMU, navigate to the application representing the database that will become the data guard primary. Then, create a database backup of that database with the type Standby.
3. Select the backup that you created in step 2, and then select the action Clone. Because the backup was created with the type of Standby, the clone operation will automatically be set to a Data Guard Standby Clone.
4. Continue through the clone process as you would for a Clone Copy operation. Specify the data guard protection mode desired. When the setup is completed, click Finish.

SMU will now perform the following operations:

5. SMU will create a replication action for each project that contains database shares using a configured replication target pointing to the target (standby) storage.

6. SMU will initiate a send update operation for the replication action. This send update transmits the source data to the replication target; thus, the shares containing the data guard primary will be sent to the standby storage.
7. SMU will wait for the update operation to complete on the standby storage. This may take considerable time.
8. SMU will sever the replica (replication package on the target storage). New names specified by the user are given to project and mount points on the standby storage.
9. SMU will then roll back the standby shares to the snapshot level created when the standby backup was created.
10. SMU will prepare the primary and standby databases to set up a data guard configuration using Data Guard Broker.
11. SMU will enable the data guard configuration to begin log shipping and application to the standby database.

SMU utilizes Data Guard Broker to configure the primary database and the standby clone database as a data guard configuration. SMU will enable Data Guard Broker on both the primary and the standby databases as a part of standby clone creation unless data broker is already enabled. For Oracle Database 11g in ASM/iSCSI configuration, SMU requires SSH File Transfer Protocol (SFTP) enabled on the primary database host and the standby database host.

SMU also requires Oracle Net listeners specified in the `listener.ora` file. At least one listener specified in the `listener.ora` file must be running if the `listener.ora` file is not located in the default location (`$ORACLE_HOME/network/admin`). In order to avoid listener conflicts, do not specify multiple listener configuration files.

SMU can also create standby clones for container databases, or CDBs, which can hold multiple user-created pluggable databases, in support of Oracle Database 12c's multitenant capabilities.

Database Clone from RMAN Backup (RMAN Clone)

SMU can be used to create a new primary database from an RMAN backup for development or testing purposes. The Oracle ZFS Storage Appliance can be used as a database backup device, especially within Oracle Exadata environments. Database backups on the Oracle ZFS Storage Appliance, when they are in image copy format, can be the basis for creating database clones by using ZFS snapshot and clone technology.

You can only create filesystem storage type clone databases using this cloning method.

Requirements for the RMAN backup include:

1. The RMAN backup must be in image copy format.
2. Backup filenames must be in the %U format specification.
3. Backup files must include datafiles, archived log, and controlfile.
4. The selected share(s) must contain a single RMAN backup set to be used in the clone operation. No other RMAN files or Oracle files may reside on the share(s).
5. The target host(s) for the clone operation must not be the source database node(s).
6. The target host(s) must have datapath connectivity to the Oracle ZFS Storage Appliance.
7. The target host(s) must be using the same type of operating system (endian order; for example, Oracle Linux x86 and Oracle Solaris x86), clusterware (cloud/grid control) version and database software and patch levels as the source database node(s).

Note: For detailed information on cross-platform database cloning, see the *Oracle Database Backup and Recovery User's Guide* for the release you are using, the section titled "Transporting Data Across Platforms." Refer to the References section at the end of this guide for location information.

RMAN cloning works in the following way:

1. Backup shares are snapped and cloned.
2. Clone shares are mounted on the target host.
3. SMU scans the backup shares to identify the various database files.
4. SMU starts up a temporary database instance so that it can mount the backed-up controlfile. Note that this requires that no database with the same name as the backed-up database already exist on the target host.
5. SMU performs a set of queries against the backup controlfile to get information about the backup. This includes calculating the point to recover the database to (max SCN) and the size the flash recovery area (FRA) should be.
6. After gathering information about the backup, SMU shuts down the temporary instance.
7. SMU creates a parameter file for the clone database.
8. SMU starts the clone database.
9. SMU creates a new controlfile for the clone database.
10. SMU recovers the database.

11. SMU opens the database and resets the logs.

12. SMU re-compiles the schema objects. This step allows for creating clones in an environment with a different operating system type from the source database, as long as the supported systems are of the same endian architecture.

Important: Only deprovision (delete) the database clone (from RMAN backup) that was created using the SMU from the SMU. Do *not* delete the database clone outside of SMU using DBCA (Database Configuration Assistant, a database software tool), for example. SMU needs to connect to and query the clone database before it is deprovisioned.

Refresh Clone

The refresh clone operation simplifies the process of updating, or refreshing, data contained in an existing clone to bring it up to date to reflect the latest state of the source database. Currently, maintaining a current database clone, such as for ongoing dev/test environments, requires a two-step process of deprovisioning the old clone, then creating a new one. The refresh operation combines those processes, streamlining them into one easy, efficient user step.

The basic steps for the refresh clone operation are:

1. The source backup and target application are specified for the refresh command.
2. The old clone is deprovisioned.
3. A new clone is created from the user-selected backup.
4. The `log_mode` and `sga_size` database parameter values of the source database when the selected backup was taken are used for the new clone.
5. The `open_mode` database parameter value of the old clone prior to the refresh clone is retained for the new clone.
6. The rest of the database parameter values when the old clone was created are used to create the new clone.

Supported Systems and Configurations

While Snap Management Utility supports some common administrative tasks performed on Oracle databases utilizing Oracle ZFS Storage Appliances, there are important restrictions that must be followed on these systems. Carefully review the following supported systems, configurations, and database file types.

Supported Storage Systems

Snap Management Utility supports all models of the Oracle ZFS Storage Appliance, which includes single and clustered systems. SMU does not support third-party storage systems.

The Oracle ZFS Storage Appliance must be running Oracle ZFS Storage Appliance Software OS 8.4 (2013.06.05) or later.

Special Considerations with Clustered Systems

SMU was originally designed to access only one head of a clustered Oracle ZFS Storage Appliance during operations. SMU accesses the head specified in the storage account that is set up in the SMU, and the appropriate pools and shares must be available from that head.

SMU can administer databases stored on a clustered Oracle ZFS Storage Appliance and which span both heads in an active-active configuration. This feature better supports best practices for achieving high-performance backup operations with systems such as Oracle Exadata.

When working with a clustered database, the grid infrastructure/clusterware software must be configured to use the correct timezone, and the system clocks of each node in a cluster must be in sync by running NTP (network time protocol).

Supported Oracle Databases

Please note that these following supported versions are minimum versions, and users should always update to the latest releases and patches beyond the minimum supported versions.

The Oracle Snap Management Utility supports the following application software:

- Oracle Database 11g R2 – Enterprise Edition
- Oracle Database 12c R1 (12.1.0.2 or later) – Enterprise Edition

Note that these editions are *not* supported:

- Express Edition

- Standard Edition One
- Standard Edition (Note that this edition is limited; certain SMU features such as Clone Copy require Enterprise Edition)

Refer to Table 2, Supported Database Configurations, for further Oracle Database version information. This Oracle Database application software must be installed and configured separately.

Snap Management Utility supports both single instance and Oracle Real Application Cluster (Oracle RAC) database configuration types. Oracle RAC One Node database is supported as well.

In Oracle RAC configurations, SMU does not allow selection of which nodes in a cluster to run the clone databases on. SMU will use all enabled nodes in the cluster for the clone database.

Note: Common Internet File System (CIFS), used primarily in Microsoft Windows network environments, is not supported by Oracle Database. For more information, refer to http://docs.oracle.com/cd/E11882_01/win.112/e10845/architec.htm.

Supported Database Layouts

The types of backups that can be performed on a database depend on the database layout. SMU makes no restrictions on the number of shares a database can use. The shares can span projects, pools and, in the case of clustered storage systems, heads or controllers. There are, however, some restrictions on how files can be laid out within those shares. SMU requires that the datafiles and archived logs be in separate shares in order to take online, or hot, backups.

During an online backup, snapshots are taken of the datafile shares first, then a log switch is performed, and then snapshots are taken of the archived logs shares. Because snap backups occur at the share level, the database files must be in different shares in order for this sequencing to work. There are no file restrictions with offline (cold) backups.

Database layout is also an important consideration with database clones created by importing RMAN backups. Clone databases created by importing an RMAN backup will have the same layout as the backup shares: the clone database datafiles will be in clones of the backed-up datafiles' shares and the clone database archived logs will be in the clones of the backed-up archived log shares. If only a single backup share is used, the datafiles and archived logs will be in the same share.

Consequently, the clone database operation will not support online backups because the files are in the same share. In order to create a clone database that can have online backups taken, there must be at a minimum two backup shares with the backed-up datafiles in one share and the backed-up archived logs in the other share.

Refer to table 1 for the database layouts supported by SMU. SMU requires that each database instance be in separate Oracle ZFS Storage Appliance shares, whether they are NFS/dNFS filesystem or iSCSI LUN shares, and each archiving area must be on separate shares or LUNs from the database. So each application requires two shares, one for the archive and one for the database, and nothing else can reside on those Oracle ZFS Storage Appliance shares.

When SMU performs operations on an Oracle Database, it uses Oracle ZFS Storage Appliance snapshot or clone operations; this is why each database must exist on its own separate set of Oracle ZFS Storage Appliance shares/LUNs. On clone operations, SMU mounts the cloned Oracle ZFS Storage Appliance share/LUN on the targeted host and then logs into this server to manipulate the cloned database files to complete the clone operation.

Supported Database File Layouts

As detailed in Table 1, SMU requires that these database files be contained in one or more appliance shares:

- datafiles
- controlfiles
- online redo logs
- archived redo logs

Additionally, the datafiles and archived redo log files must be in separate shares in order to support the taking of online snap backups for NFS/dNFS shares. During an online backup, snapshots of the datafile shares are taken first and the archived redo log shares are taken next, along with any other database shares.

Each database must use its own separate set of shares. No two databases can use the same set of shares. SMU performs operations at the appliance share-level using ZFS operations. SMU will not copy, modify or remove individual files in a share. It will only use ZFS operations to clone or snapshot a share.

Database `oratab` File Entries

SMU requires that there be an entry in the `oratab` file for each database it will administer. SMU uses the information in the `oratab` file to determine the Oracle home of the database. During each SMU operation, SMU will first look up the Oracle home for the database to be operated on by searching the `oratab` file on the database host. If no entry is found, then the operation will end in an error. Also, an `oratab` entry may be erroneously removed or modified after an instance restart. In cases such as these, the entry may need to be manually added to the `oratab` file by applying an Oracle patch. For further information on this error condition and to download the patch, please see the Oracle Support Document (MOS note) 1922908.1 (12.1.0.2).

During a cloning operation, SMU will add an entry for the clone database to the local `oratab` file. The `oratab` file is located in different locations based on the host operating system:

Oracle Solaris – /var/opt/oracle/oratab

Oracle Linux – /etc/oratab

Microsoft Windows – N/A (Oracle home information is stored in the Windows registry.)

Database Layout and Types of Snap Backups

Generally, database backups can be performed when a database is either offline or online. When a database is offline it cannot be updated by users, so there is no potential for new data loss. However, with online (hot) backups, data may be concurrently accessed by users, so a successful backup scenario must accommodate new or changing data capture.

Oracle classifies backups as either consistent or inconsistent. In consistent backups, all files contain the same set of changes. Conversely, inconsistent backups do not contain the same set of changes. Consistent backups are accomplished by shutting down the database and making the backup while the database is closed. This corresponds to an offline or cold backup. Consistent backups do not require recovery when they are used later. Consistent backups are the only valid backup option for databases operating in the NOARCHIVELOG mode.

Inconsistent backups correspond to online or hot backups. These types of backups require recovery – which is the process of making the files consistent – for later use.

For the Snap Management Utility, additional restrictions exist with the database file layouts, depending on which type of snap backups you want to create. The following table describes these restrictions for both offline (cold) and online (hot) backups. Also refer to the subsection “Database Backup (Snap Backup)” in the section “Supported Operations” for more information.

TABLE 1. SUPPORTED DATABASE LAYOUTS

DATABASE STORAGE TYPE	OFFLINE (COLD) BACKUP AND STANDBY BACKUP	ONLINE (HOT) BACKUP
Filesystem (NFS/dNFS)	Database files can be in one or more shares. Database files may span shares and pools within the same storage head.	Database datafiles must be in separate shares from the database archived logs. Database files may span shares and pools within the same storage head.

TABLE 1. SUPPORTED DATABASE LAYOUTS

DATABASE STORAGE TYPE	OFFLINE (COLD) BACKUP AND STANDBY BACKUP	ONLINE (HOT) BACKUP
ASM	Database files are in a diskgroup that uses external redundancy and consists of one or more iSCSI LUNs. LUNs must not be partitioned. Cloning with ASM requires Oracle Database version 11.2.0.2 or later.	Not supported.
FOR BOTH ASM and NFS/dNFS	Control files, datafiles, and logs (online and archive) must reside on the Oracle ZFS Storage Appliance.	

Using Oracle Intelligent Storage Protocol with Snap Management Utility

A feature of Oracle Database 12c called Oracle Intelligent Storage Protocol (OISP) provides beneficial autotuning and auto-provisioning features when paired with the Oracle ZFS Storage Appliance as host storage. Oracle Snap Management Utility does not currently provide configuration of clone databases it creates to use OISP.

However, users can perform this additional configuration on these SMU-generated clone databases by following the steps outlined in the Oracle Support Document 1943618.1 titled "Oracle ZFS Storage Appliance: How to Enable Oracle Intelligent Storage Protocol (OISP)" in My Oracle Support (MOS). The MOS link is provided in Appendix A: References at the back of this document. Users will need to perform steps 1-2 and 5-7; SMU performs steps 3 and 4. The documented steps 1 and 2 involve configuring SNMP on the Oracle ZFS Storage Appliance, which SMU does not perform.

Database Configuration and Storage Types by Host Operating Systems

The following table shows the supported configurations by host operating system.

TABLE 2. SUPPORTED DATABASE CONFIGURATIONS BY STORAGE TYPE AND HOST OS

STORAGE TYPE	ORACLE LINUX/ RED HAT LINUX (6.X, 7.X)	ORACLE SOLARIS (SOLARIS 10 UPDATE 7 OR LATER)	MICROSOFT WINDOWS (WINDOWS 2008 R2)
Filesystem (Kernel NFS)	Supported	Supported	Not supported
Filesystem (Direct NFS)	Requires Oracle Database 11g R2 or later. For clone operation in this OS, SMU will not create or modify the <code>oranfstab</code> file, but will only update the system mount table with entries for the clone shares.	Requires Oracle Database 11g R2 or later. For clone operation in this OS, SMU will not create or modify the <code>oranfstab</code> file, but will only update the system mount table with entries for the clone shares.	With Oracle Database 11g R2 only, note a static location layout requirement due to bug 13571798. See My Oracle Support (MOS) at http://support.oracle.com , metalink doc 1452760.1 for more information. This release only supports adding single network paths to the filesystem shares in the <code>oranfstab</code> file.

TABLE 2. SUPPORTED DATABASE CONFIGURATIONS BY STORAGE TYPE AND HOST OS

STORAGE TYPE	ORACLE LINUX/ RED HAT LINUX (6.X, 7.X)	ORACLE SOLARIS (SOLARIS 10 UPDATE 7 OR LATER)	MICROSOFT WINDOWS (WINDOWS 2008 R2)
ASM	<p>ASM cloning is not supported on Linux version 5.5. External redundancy only; cloning requires Oracle Database 11.2.0.2 or later due to bug 9316059.</p> <p>Diskgroup members must be physical or raw disks. Virtual Pseudo disks which are used with advanced I/O such as ASMLib, MPIO, logical volume groups, device-mapper, are not supported.</p>	<p>External redundancy only; cloning requires Oracle Database 11.2.0.2 or later due to bug 9316059.</p> <p>Diskgroup members must be physical or raw disks. Virtual Pseudo disks which are used with advanced I/O such as ASMLib, MPIO, logical volume groups, device-mapper, are not supported.</p>	<p>External redundancy only; cloning requires Oracle Database 11.2.0.2 or later due to bug 9316059.</p> <p>Diskgroup members must be physical or raw disks. Virtual Pseudo disks which are used with advanced I/O such as ASMLib, MPIO, logical volume groups, device-mapper, are not supported.</p>
*****	*****	*****	*****
<p>Note: ASM_DISKSTRING parameter settings. (Use of ASMLib is not supported.)</p>	<p>Parameter default is /dev/raw/* . Must append ASM_DISKSTRING parameter with the value /dev/smuasm/* . See My Oracle Support (MOS) at http://support.oracle.com metalink doc 9316059.8 for more information on this bug and other SMU issues.</p>	<p>Parameter default is /dev/rdisk/* . Leave as is. See My Oracle Support (MOS) at http://support.oracle.com metalink doc 9316059.8 for more information on this bug and other SMU issues.</p>	<p>Parameter default is *\\.\ORCLDISK* Leave as is. See My Oracle Support (MOS) at http://support.oracle.com metalink doc 9316059.8 for more information on this bug and other SMU issues.</p>

Supported Application (Database) Hosts

The Oracle Snap Management Utility supports the following application hosts:

- Oracle Linux/Red Hat Linux (6.x, 7.x)
- Oracle Solaris (Solaris 10 Update 7 or later)
- Microsoft Windows (Windows Server 2008 R2)

One of the following must be installed and configured on each application host:

- Windows Remote Shell 2.0 (for Microsoft Windows only)
- Secure Shell 2.x (SSH2) (for Oracle Solaris or Oracle Linux only)

These components are used by the utility to perform remote command execution on the host. They are usually a part of the core operating system and are installed by default. However, if they are not present on the host, you may have to install these components manually.

Oracle Linux and Oracle Solaris hosts do not require any additional configuration.

Microsoft Windows hosts may require additional configuration of the Windows Remote Management Service. The Windows Remote Management (WinRM) Service is started automatically on Windows Server 2008 hosts; however, no WinRM listener is configured by default. A listener must be configured to enable remote management of the host. SMU supports encrypted (HTTPS) listeners only. To enable a listener, follow the steps outlined later in this document under “Configuring a Windows HTTPS Connection.”

For additional information, consult the article "Installation and Configuration for Windows Remote Management" available at: [http://msdn.microsoft.com/en-us/library/aa384372\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/aa384372(v=vs.85).aspx).

In addition to enabling a listener, Windows Remote Management protocol and service settings must be modified for correct SMU operation. These configurations are detailed in the section “Configuring Host Systems to Use the Snap Management Utility” for detailed Microsoft Windows host settings.

SMU Host Installation Requirements

The host system (either application host or management host) on which you install the Snap Management Utility will require the following:

- Oracle Solaris: Solaris 10 (SPARC and x86 editions) and Solaris 11
- Linux: Oracle Linux 6 or later

- Microsoft Windows: Windows Server 2008 R2
- 4 GB RAM minimum; 8 GB RAM recommended
- 100 MB free disk space minimum
- Java Runtime Environment (JRE) 6 or later

As just noted, the Oracle Snap Management Utility requires JRE 6 or later. You can use the Sun Java Development Kit (JDK), the Oracle JRE, or the OpenJDK version of the runtime environment.

Installing the Oracle Snap Management Utility

The Oracle Snap Management Utility is distributed as a set of packages, one package for each supported host environment. The packages are installed using the standard package install command for each environment. You can install the package on the application host or on a separate management host.

Before installing the utility, you must download and extract the software distribution onto the target host or a network location accessible from the target host.

Installing the Oracle Snap Management Utility on an Oracle Linux Host

To install the Oracle Snap Management Utility on a Linux host, perform the following steps.

1. Log in as root (or other privileged user).

```
$ su
Password:
#
```

2. Change directory to where the Linux package is located. The same package is used on either 32-bit or 64-bit Linux hosts.

```
# cd <extract-dir>/Linux/noarch
```

3. Install the package. In the following command, *n.n.n-nn* represents the latest version number of the utility.

```
# rpm -i oracle-smu-<n.n.n-nn>.noarch.rpm
```

By default, the utility is installed under `/opt/oracle/smu`.

Installing the Oracle Snap Management Utility on Oracle Solaris Hosts

To install the Oracle Snap Management Utility on Oracle Solaris hosts, perform the following steps:

1. Become root (or other privileged user).

```
$ su
    Password:
    #
```

2. Change directory to where the Oracle Solaris package is located.

```
# cd <extract-dir>/SunOS/all
```

3. Install the Oracle Solaris package.

```
# pkgadd -d ORCLsmu-n.n.n-nn.all.pkg all
```

By default, the utility is installed under `/opt/ORCLsmu`.

Installing the Oracle Snap Management Utility on Microsoft Windows Hosts

To install the Oracle Snap Management Utility on Windows hosts, perform the following steps:

1. Change to the folder where the Windows package is located.
2. Double click the setup program to launch a graphical installation.

If a compatible version of Java is not found on the system, the install program will install the Sun JRE 6 before installing the utility.

By default, the software is installed in `%ProgramFiles%\Oracle\Oracle Snap Management Utility`.

Configuring the Snap Management Utility Host and Database Host

There are three types of systems that make up an SMU environment: storage appliance, database host, and management (SMU) host. The storage appliance contains the database files, the database host is where the database instances run, and the management host is where SMU runs. It is possible for the database host and management host to be the same.

Default network port settings for all hosts can be reconfigured with the information provided in the following table.

Additionally, the systems' host time clock and timezone should be properly set and, in the case of clustered database systems, these settings should be in sync for each node of the cluster on both data host and management host.

Snap Management Utility requires no further configuration on Oracle Solaris or Oracle Linux database hosts. Microsoft Windows hosts require further configuration, including an HTTPS connection and its required certificate, and Windows Remote Management Service protocol and configuration settings.

IMPORTANT: Only one instance of the Snap Management Utility software should be used to administer a database. Each instance of SMU software stores information about the backups that it creates. This information is not available to other SMU software instances.

Configuring the Network Port Settings

After startup, SMU opens a set of network ports to accept incoming connections from SSH clients, Windows Remote Shell (WinRS) clients, and web browsers. The ports used are user configurable by editing the `smu.conf` file which is located in the `installation/etc` directory or folder. The following table shows the default port settings.

TABLE 3. DEFAULT SMU SERVER PORT SETTINGS

SERVICE	DEFAULT PORT	CLIENT USAGE NOTES
Windows Remote Shell (HTTPS/encrypted)	8001	<p>Use the <code>winrs.exe</code> command with either <code>-ssl</code> to specify secure socket layer or <code>-r</code> to specify https.</p> <p>The SMU server uses a self-signed certificate. This certificate must be installed in the trusted certificate store on the Windows client host before you can connect to the SMU server.</p> <p>Command examples:</p> <pre>winrs -r:<SMU host FQDN>:8001 - u:<LOCAL user> -ssl smu winrs -r:https://<SMU host FQDN>:8001 -u:<LOCAL user> smu</pre>

TABLE 3. DEFAULT SMU SERVER PORT SETTINGS

SERVICE	DEFAULT PORT	CLIENT USAGE NOTES
SSHD Version 2	8002	Use the <code>-o "HostKeyAlias <alias>"</code> option to specify an alias to use. EXAMPLE: <code>ssh -l <LOCAL or LDAP user> -p 8002 -o "HostKeyAlias smu" <SMU host></code>
Web Server (HTTPS/encrypted)	8443	URL: <code>https://<SMU host>:8443/smu</code> The SMU server uses a self-signed certificate. You must accept this certificate before the connection will be allowed.

Verifying and/or Configuring Time Settings in the Database and Management Hosts

Some SMU operations, particularly recovery, depend upon the processing of archive logs, for which individual logs are timestamped with the current local time as they are written and archived. If the host operating system clock or the timezone is set incorrectly, the operation can fail.

Network time protocol (NTP)/time server should be configured on each database host and also the management host where SMU runs. For clustered database systems, grid infrastructure mandates that the system clocks of each node in a cluster be running in sync with NTP.

To verify that the archive logs are stamped using the correct time and timezone, you can run the following query on the database:

```
SQL> select sequence#, first_change#, to_char(first_time, 'yyyy-mm-dd hh24:mi:ss'),
next_change#, to_char(next_time, 'yyyy-mm-dd hh24:mi:ss') from v$archived_log order by
first_change#;
```

If needed, change the timezone setting in the following file for each node of the cluster:

```
$GRID_HOME/crs/install/s_crsconfig_<nodename>_env.
```

If you change the timezone, you must restart CRS for the changes to take effect.

The following documentation can be referenced to understand how the Oracle database works with times, timestamps and timezones:

Oracle Support Document 340512.1 (Timestamps & time zones - Frequently Asked Questions) can be found at: <https://mosemp.us.oracle.com/epmos/faces/DocumentDisplay?id=340512.1>

Oracle Support Document 1627439.1 (How to Diagnose Wrong Time (SYSDATE and SYSTIMESTAMP) After DST Change , Server Reboot , Database Restart or Installation When Connecting to a Database on an Unix Server) can be found at: <https://mosemp.us.oracle.com/epmos/faces/DocumentDisplay?id=1627439.1>

Oracle Support Document 1209444.1 (How To Change Timezone for 11gR2 Grid Infrastructure) can be found at: <https://mosemp.us.oracle.com/epmos/faces/DocumentDisplay?id=1209444.1>

Configuring a Windows HTTPS Connection for the Snap Management Utility

In order to use HTTPS connections with the Windows-based host, you must configure the HTTPS listener on the Windows host. An HTTPS listener will require an SSL certificate. The certificate may already be installed on the host or you can install a certificate authority (CA)-signed certificate or generate a self-signed certificate to use.

The following instruction uses the Java 7 `keytool` command to generate a self-signed certificate for the host.

1. Execute the following command to generate a self-signed certificate:

```
keytool -genkey -keyalg rsa -keysize 1024 -sigalg SHA1withRSA -alias cert -
keystore keystore.jks -storepass changeit -keypass changeit -validity 360 -
dname "CN=<fully qualified domain name of the host>" -ext
KU:true=dS,keyE,dataE -ext EKU:true=serverAuth
```

2. Convert the certificate to the correct format:

```
keytool -importkeystore -srckeystore keystore.jks -srcstorepass changeit
-destkeystore cert.p12 -deststoretype PKCS12 -deststorepass changeit
```

3. Import `cert.p12` into local computer account of "Trusted Root Certification Authorities" and "Personal" trust store on the Windows host.

4. Configure the WinRM HTTPS listener:

```
winrm create winrm/config/listener?Address=*&Transport=HTTPS
@{Hostname="<fully qualified domain name of the
host>";CertificateThumbprint="<hex thumbprint digits>"}
```

5. View the certificate thumbprint by double-clicking the `host.p12` file and looking at the thumbprint field of the certificate. A thumbprint is 40 hex digits.

6. Verify the listener is configured:

```
winrm enumerate winrm/config/listener
```

Windows Remote Management Protocol and Service Settings for the Windows Database Host

In addition to enabling a listener for Windows hosts, the following Windows Remote Management protocol and service settings must be modified for correct SMU operation:

TABLE 4. RECOMMENDED WINDOWS REMOTE MANAGEMENT PROTOCOL SETTINGS

TERM	DEFAULT VALUE	SMU RECOMMENDED VALUE	COMMENTS
MaxTimeoutms	60000 (1 minute)	7200000 (2 hours)	Some SMU operations can take a long time to complete. For example, when cloning a database to a different host, SMU may need to recompile schema objects if the host platform changes. This can take an hour or more.
IdleTimeout	180000 (3 minutes)	7200000 (2 hours)	Increasing the idle timeout for the remote shell on the Windows host allows it to remain open when SMU executes long-running commands. Because SMU alternates issuing commands among the database, storage, and the host, the host can be idle but must remain open.

TABLE 5. RECOMMENDED WINDOWS REMOTE MANAGEMENT SERVICE CONFIGURATION SETTINGS

TERM	DEFAULT VALUE	SMU RECOMMENDED VALUE	COMMENTS
MaxConcurrentOperationsPerUser	15	1500	While performing operations, SMU executes a greater number of SQL Plus, RMAN, and system commands than the number of commands allowed by default.

TABLE 5. RECOMMENDED WINDOWS REMOTE MANAGEMENT SERVICE CONFIGURATION SETTINGS

TERM	DEFAULT VALUE	SMU RECOMMENDED VALUE	COMMENTS
Basic	FALSE	TRUE	This SMU release only supports Basic access authentication.

The following example modifies the Windows Remote Management configuration to meet SMU requirements:

```
C:\>winrm set winrm/config @{MaxTimeoutms="7200000"}
C:\>winrm set winrm/config/winrs @{IdleTimeoutms="7200000"}
C:\>winrm set winrm/config/service @{MaxConcurrentOperationsPerUser="1500"}
C:\>winrm set winrm/config/service/auth @{Basic="True"}
```

Configuring Host System Name Services Recognized by the SMU Host

While performing its tasks, Snap Management Utility (SMU) logs in to various database hosts and Oracle ZFS Storage Appliances. Which systems are used is specified by the accounts that have been added to SMU. Each storage and host account has a hostname property that identifies the name of the system to access. The hostnames or IP addresses used for this property must be resolvable from the management host where SMU is running.

In addition, on database hosts, depending on the database storage type, SMU will examine the system mount table and will resolve the hostnames specified in the resource string for NFS entries. Any hostnames used by the database host or storage appliance must be resolvable from the management host where SMU runs. It may be necessary to add entries to the local hosts file on the management host if there are any hostnames that do not have entries in the system name service.

Starting and Stopping the Snap Management Utility on Host Systems

The first step in using Snap Management Utility (SMU) is to start the SMU server. SMU has been designed to run continuously on the host upon which it is installed, much like a daemon or network service.

To start the SMU server, execute the SMU launcher script. This script is located in the install bin directory or folder.

Note: To start or stop the SMU server, the user must be root. If your server reboots, you will need to restart the Snap Management Utility.

Starting Oracle Snap Management Utility on Linux Hosts

SMU is controlled by a System V init script. This replaces the control script SMU employed in previous SMU releases and allows SMU to be automatically shut down and restarted during system boot.

To start Oracle Snap Management Utility on Oracle Linux hosts, use the following command that will run the SMU server in the background:

```
# /sbin/service smud start
```

To stop the SMU, use the following command:

```
# /sbin/service smud stop
```

To check the running status of the service, use the following command:

```
# /sbin/service smud status
```

Starting Oracle Snap Management Utility on Oracle Solaris Hosts

For Oracle Solaris hosts only, SMU uses Oracle Solaris' Service Management Facility (SMF) service upon startup to run the SMU server in the background.

To start Oracle Snap Management Utility on Oracle Solaris hosts, use the following command.

```
$ svcadm enable svc:/applicaton/management/smu
```

To stop the SMU, use the following command:

```
$ svcadm disable svc:/applicaton/management/smu
```

To verify that the server is running, use the following command:

```
$ svcs applicaton/management/smu
```

Starting Oracle Snap Management Utility on Windows Hosts

To start Oracle Snap Management Utility on Windows hosts, use the following commands.

To start the service (only need to do this once after installation):

```
C:\>sc start "Oracle SMU Service"
```

To stop the service:

```
C:\>sc stop "Oracle SMU Service"
```

To check the status of the service:

```
C:\>sc query "Oracle SMU Service"
```

After the SMU server has successfully started, it creates its files in the SMU server user home directory. Depending on the operating system, the following directory/folder is created:

For Oracle Linux: `/var/opt/oracle/smu`

For Oracle Solaris: `/var/opt/ORCLsmu`

For Windows: `C:\ProgramData\Oracle\Oracle Snap Management Utility`

Updating the Snap Management Utility on Host Systems

Keep apprised of updates to Snap Management Utility for Oracle Database by visiting My Oracle Support. New releases and minor updates offer the benefits and features of ongoing enhancements to SMU. Instructions for installing these updates are always contained in Read Me files that are included with the update.

For Oracle Solaris hosts, you must uninstall the previous version of SMU before installing the update.

Updating Oracle Snap Management Utility on Oracle Solaris Hosts

Use the following commands:

```
# pkgrm ORCLsmu (removes the existing/old package version)
```

```
# pkgadd -d ORCLsmu-$VERSION-$RELEASE.all.pkg ORCLsmu (installs the new package version $VERSION, release $RELEASE)
```

Updating Oracle Snap Management Utility on Oracle Linux Hosts

Use the following command:

```
# rpm -U oracle-smu-$VERSION-$RELEASE.noarch.rpm (installer upgrade existing software to version $VERSION and release $RELEASE)
```

Updating Oracle Snap Management Utility on Windows Hosts

Use the same operation that you use to install the Oracle Snap Management Utility on Windows hosts: double click the setup program.

Protecting Oracle Snap Management Utility's Data Files

Protecting the integrity of Snap Management Utility's own data files is a key best practice for ensuring SMU's successful operation. Use the following procedures to keep these files safe. The SMU restore utility is also useful in cases of detected corruption.

Depending on the host operating system, SMU uses the following associated directories/folders to store its own data files:

Oracle Solaris – `/var/opt/ORCLsmu`

Linux – `/var/opt/oracle/smu`

Windows – `%PROGRAMDATA%\Oracle\Oracle Snap Management Utility`

This directory/folder should be backed up in order to protect this data. Either manually back up the directory/folder or use standard operating system tools or vendor software to make the backups. To ensure a consistent backup of the files, the SMU service/daemon should be temporarily stopped for the backup process.

It is also possible to migrate SMU from one management host to another by taking a backup of the SMU data on the current management host and restoring the SMU data on a new management host. SMU employs a Java database to store configuration information of resources that it manages. To protect this core SMU working data, SMU runtime performs a daily online backup of the Java database. By default, the online backup of the Java database is scheduled for 23:00 pm. The daily backup schedule can be changed by modifying the `backup.schedule` parameter in the `$SMU_HOME/etc/smu.conf` file. The backup files are located in the `.../backups/mm-dd-yy` directory. Use standard operation system commands to delete the backup files beyond the retention expiration date.

The following SMU restore utility that follows is also useful in cases where corruption has already occurred.

To restore the SMU database, use the following commands.

For Linux/Oracle Solaris:

```
restore [ -d <backup_directory> ]
```

Note that this SMU command, which is separate from the Linux OS-level restore command, must be run from within SMU at `$SMU_HOME/bin/restore`.

For Windows:

```
restore.bat [ -d <backup_directory> ]
```

where

-d backup directory = the location of database directory to restore from

Backups are stored in the following directories:

Linux: `/var/opt/oracle/smu/backups/<date>/db`

Oracle Solaris: `/var/opt/ORCLsmu/backups/<date>/db`

Windows: `%ProgramData%\Oracle\Oracle Snap Management Utility\backups\<date>\db`

Restoring the SMU database must be performed while the SMU service is not running. Stop the SMU service prior to performing the restore. After performing the restore, references to the backups and clones on the storage created after the backup time may be lost.

Restoring a Clone Profile in the SMU Database

Once a database clone is created, its associated share, storage, database and host information are contained in xml file form in its clone profile, located in the `$SMU_WORKING_DIR/store` directory. This clone profile is used to deprovision the clone database whenever the database instance is unavailable. If the clone profile is missing or corrupted, or if the clone database information has changed, the SMU command utility `reprofile` can be used to recreate the latest information for the clone profile in the SMU database.

For Linux/Oracle Solaris, use the command `reprofile`.

For Windows, use the command `reprofile.bat`.

The `reprofile` operation must be executed while the SMU service is not running. Stop the SMU service prior to performing `reprofile`.

User Permissions Requirements for Accessing Operations

During operations the SMU software accesses both the Oracle Database (database host) and the Oracle ZFS Storage Appliance (storage system). In order for SMU to access these systems, the software requires that user accounts be added for these systems. The users specified in these accounts must have the appropriate permissions or authorizations to perform some privileged commands.

SMU supports three types of host account users: root, Oracle user and delegated user. The following outlines the permissions that are required by the host account and storage account users.

Host Account User – Must be the root user, or the Oracle user if the Oracle user can do the following:

- Modify the filesystem table (`/etc/fstab` or `/etc/vfstab`)
- Create mountpoints under the root directory (`/`)
- Mount NFS filesystems
- Unmount NFS filesystems
- Scan the SCSI bus for new LUNs

Storage Account User – For detailed information on storage appliance configuration, consult the *Oracle ZFS Storage Appliance Administration Guide* listed in the References section of this document or access the online help function within the Oracle ZFS Storage Appliance BUI. Users must have the following configuration on the Oracle ZFS Storage Appliance:

- Authorization Levels:
 - Scope: Project and shares
 - Pool: * or specific pool
 - Project: * or specific project
 - Share: * or specific share
- Required Action/Operation Authorizations:
 - `changeAccessProps`
 - `changeGeneralProps`
 - `clone`

- createShare
- delete
- rename
- rollback
- takeSnap

The SMU delegation feature allows designation and use of the sudo tool on Linux and Oracle Solaris database hosts. This delegation to sudo enables a third type of user, called delegated user, who is not root user and not Oracle user. This delegated user can be established with credentials that provide a restricted access for only certain operations that are part of a root user's or Oracle user's permissions. The root user establishes who has sudo privileges.

The term sudo literally means "super user" and "do." See the Glossary for a fuller definition of the sudo tool.

Using SMU Delegation Tools

The SMU delegation tools feature supports the use of sudo on Oracle Linux and Oracle Solaris database hosts. (This feature is not supported on Windows.) For Oracle users and delegated users, specific database host permissions and authorizations must be configured by the administrator. Using this delegation feature, which is configured as part of the initial application host account setup, the previously listed permissions that are required of root users and must be specifically designated for an Oracle user can be assigned, or delegated, through sudo.

The following host account setup table summarizes permissions and authorizations configurations for Oracle users and delegated users in both single instance and Oracle RAC host accounts.

TABLE 6. RECOMMENDED HOST ACCOUNT USER PERMISSIONS AND AUTHORIZATIONS CONFIGURATIONS

Host Account Type	Host Account User	Permissions and Authorizations Configuration on Linux	Permissions and Authorizations Configuration on Oracle Solaris
Single Instance	Oracle User	<p>Set !requiretty to sudo user.</p> <p>Sudo privileges include:</p> <ul style="list-style-type: none"> Using /usr/bin/find command for RMAN import feature For NFS: <ul style="list-style-type: none"> mount/umount NFS filesystems Create/delete mountpoints under the root directory (/) (Sudoedit privilege) Modify the filesystem table (/etc/fstab) For ASM: <ul style="list-style-type: none"> Scan the SCSI bus for new LUNS (/usr/bin/scsi-rescan -r, /sbin/iscsiadm) Trigger and settle udev events (Sudoedit privilege) Modify the udev rules (/etc/udev/rules/*.rules) 	<p>Set !requiretty to sudo user.</p> <p>Sudo privileges include:</p> <ul style="list-style-type: none"> Using /usr/bin/find command for RMAN import feature For NFS: <ul style="list-style-type: none"> mount/umount NFS filesystems Create/delete mountpoints under the root directory (/) (Sudoedit privilege) Modify the filesystem table (/etc/vfstab) For ASM: <ul style="list-style-type: none"> Scan the SCSI bus for new LUNS (/usr/sbin/disks -C)
Single Instance	Delegated User	<p>Privileges same as Oracle User's, plus:</p> <ul style="list-style-type: none"> Group Memberships (typically, same as Oracle User's): <ul style="list-style-type: none"> primary group = OSDBA (default dba) secondary group = Oracle Install Group (default oinstall) if ASM is used, also add ASMADM and ASMDBA groups. <p>Requires user home directory.</p> <p>Sudo privilege to run \$ORACLE_HOME/bin/orapwd as the Oracle User.</p> <p>For Standby Clone:</p> <ul style="list-style-type: none"> Sudoedit privilege to modify the listener.ora and tnsnames.ora files (\$TNS_ADMIN/listener.ora, \$TNS_ADMIN/tnsnames.ora) 	<p>Privileges same as Oracle User's, plus:</p> <ul style="list-style-type: none"> Group Memberships (typically, same as Oracle User's): <ul style="list-style-type: none"> primary group = OSDBA (default dba) secondary group = Oracle Install Group (default oinstall) if ASM is used, also add ASMADM and ASMDBA groups. <p>Requires user home directory.</p> <p>Sudo privilege to run \$ORACLE_HOME/bin/orapwd as the Oracle User.</p> <p>For Standby Clone:</p> <ul style="list-style-type: none"> Sudoedit privilege to modify the listener.ora and tnsnames.ora files (\$TNS_ADMIN/listener.ora, \$TNS_ADMIN/tnsnames.ora)

Host Account Type	Host Account User	Permissions and Authorizations Configuration on Linux	Permissions and Authorizations Configuration on Oracle Solaris
		<ul style="list-style-type: none"> Sudo privilege to run data guard and listener command line interfaces (<code>\$ORACLE_HOME/bin/dgmgrrl</code>, <code>\$ORACLE_HOME/bin/lsnrctl</code>) Sudo privilege to copy oracle owned files (password file, pfile, listener.ora, tnsnames.ora) Sudo privilege to run <code>trcroute</code> command to test service (<code>\$ORACLE_HOME/bin/trcroute</code>) 	<ul style="list-style-type: none"> Sudo privilege to run data guard and listener command line interfaces (<code>\$ORACLE_HOME/bin/dgmgrrl</code>, <code>\$ORACLE_HOME/bin/lsnrctl</code>) Sudo privilege to copy oracle owned files (password file, pfile, listener.ora, tnsnames.ora) Sudo privilege to run <code>trcroute</code> command to test service (<code>\$ORACLE_HOME/bin/trcroute</code>)
Oracle RAC	Delegated User	<p>User's single instance host requirements must be satisfied plus:</p> <p>User must exist on all RAC nodes.</p> <p>User's configuration such as user id and group must be the same on all RAC nodes.</p> <p>Sudo privilege to run <code>\$GRID_HOME/bin/srvctl</code></p>	<p>User's single instance host requirements must be satisfied plus:</p> <p>User must exist on all RAC nodes.</p> <p>User's configuration such as user id and group must be the same on all RAC nodes.</p> <p>Sudo privilege to run <code>\$GRID_HOME/bin/srvctl</code></p>

The following table lists the SMU command summary for sudo policy files for the Linux and Oracle Solaris operating systems:

TABLE 7. SUDO POLICY FILE SMU COMMAND SUMMARY

Operating System	Standby Backup	NFS Clone/ Deprovision	ASM Clone/ Deprovision	Import RMAN/ Deprovision
Linux	/bin/cp, /bin/mkdir	/bin/mount, /bin/ln, /bin/umount, /bin/mkdir, /bin/rmdir, sudoedit /etc/fstab, /bin/rm, /usr/bin/install, /bin/env, \$ORACLE_HOME/bin/orapwd (delegated user), /bin/cp (standby clone), \$ORACLE_HOME/bin/dgmgml (delegated user and standby clone), \$ORACLE_HOME/bin/lsnrctl (delegated user and standby clone), sudoedit \$TNS_ADMIN/listener.ora, sudoedit \$TNS_ADMIN/tnsnames.ora (delegated user and standby clone), \$ORACLE_HOME/bin/trcroute (delegated user and standby clone)	/usr/bin/scsi-rescan, /sbin/scsi_id, /bin/mount, /bin/umount, /bin/mkdir, /sbin/udevadm settle, \$ORACLE_HOME/bin/srvctl, /bin/rmdir, sudoedit /etc/udev/rules.d/99- smu.rules, /sbin/scsi_id (for Oracle Linux 6), /usr/lib/udev/scsi_id (for Oracle Linux 7), /bin/rm \$ORACLE_HOME/bin/orapwd (delegated user), \$ORACLE_HOME/bin/dgmgml (delegated user and standby clone), \$ORACLE_HOME/bin/lsnrctl (delegated user and standby clone), sudoedit \$TNS_ADMIN/listener.ora, sudoedit \$TNS_ADMIN/tnsnames.ora (delegated user and standby clone), \$ORACLE_HOME/bin/trcroute (delegated user and standby clone)	/bin/mount, /bin/umount, /bin/mkdir, /bin/rmdir, sudoedit /etc/fstab, /usr/bin/find, /bin/rm, /usr/bin/install, \$ORACLE_HOME/bin/ orapwd (delegated user)
Oracle Solaris	/bin/cp, /bin/mkdir	/usr/sbin/mount, /usr/sbin/umount, /bin/mkdir, /bin/rmdir, sudoedit /etc/vfstab, /bin/ln, /bin/env, /bin/rm, /usr/sbin/install, \$ORACLE_HOME/bin/orapwd (delegated user), /bin/cp (standby clone), \$ORACLE_HOME/bin/dgmgml (delegated user and standby clone), \$ORACLE_HOME/bin/lsnrctl (delegated user and standby clone),	/usr/sbin/disks, /usr/sbin/ mount, /usr/sbin/umount, /bin/mkdir, /bin/rmdir, /bin/chmod, /bin/chown, /bin/rm, \$ORACLE_HOME/bin/orapwd (delegated user), \$ORACLE_HOME/bin/dgmgml (delegated user and standby clone), \$ORACLE_HOME/bin/lsnrctl (delegated user and standby clone), sudoedit \$TNS_ADMIN/listener.ora, sudoedit \$TNS_ADMIN/tnsnames.ora (delegated user and standby clone),	/usr/sbin/mount, /usr/sbin/umount, /bin/mkdir, /bin/rmdir, sudoedit /etc/vfstab, /usr/bin/find, /bin/rm, /usr/sbin/install, \$ORACLE_HOME/bin/ orapwd (delegated user)

Operating System	Standby Backup	NFS Clone/ Deprovision	ASM Clone/ Deprovision	Import RMAN/ Deprovision
		sudoedit \$TNS_ADMIN/listener.ora, sudoedit \$TNS_ADMIN/tnsnames.ora (delegated user and standby clone), \$ORACLE_HOME/bin/trcroute (delegated user and standby clone)	\$ORACLE_HOME/bin/trcroute (delegated user and standby clone)	

Sample sudo policy file configuration format (/etc/sudoers)

```
Defaults <user> !requiretty
Cmd_Alias SMU = /bin/mount, /bin/umount, /bin/mkdir, /bin/rmdir,
/bin/rm, sudoedit /etc/fstab, /usr/bin/install,
/u01/app/oracle/product/11.2.0/dbhome_1/bin/orapwd ...
<user> ALL=(root, oracle) NOPASSWD:SMU
```

Be sure to consider your organization's security policies or standard operating procedures when assigning permissions as described in these requirements.

IMPORTANT: The sudo command/utility is standard on Linux and Oracle Solaris 11. It is not standard, however, for Oracle Solaris 10 clients, and administrators must obtain it from a third-party location. Recommended locations are the SUDO web site at www.sudo.ws or the Solaris Companion CD at sunfreeware.com.

Accessing the User Interfaces

SMU provides two user interfaces: a browser user interface (BUI) and a command-line interface (CLI). The command-line interface can be accessed using SSH or WinRS clients. The browser user interface can be accessed using any standard web browser.

Command-Line Interface

The SMU command-line interface can be accessed using standard remote shell commands: the `ssh` command on UNIX hosts, and the `winrs.exe` command on Windows hosts.

Accessing the Command-Line Interface Using SSH

Secure Shell is the standard for remote shell access on UNIX hosts. SMU embeds an SSHDv2 server to allow easy access to the SMU CLI from UNIX environments.

To access the command-line interface using SSH, type the following command:

```
$ ssh -l <username> -p <port> -o "HostKeyAlias <alias>" <hostname>
```

You can specify a local or LDAP user that exists in the SMU database. For *port*, use the port number that is specified in the `smu.conf` file. The `HostKeyAlias` option is used to create and/or reference a separate record in the `SSH $HOME/.ssh/known_hosts` file. By default the `ssh` command uses the hostname as the alias. By specifying a separate alias, you are able to have two separate records in the client file: one for the host and one for the SMU server.

Accessing the Command-Line Interface Using WinRS

Windows Remote Shell (WinRS) is a standard command for accessing other Windows systems remotely. For more information on Windows Remote Shell, please consult the Microsoft Windows documentation at: <http://technet.microsoft.com>. SMU embeds a WinRS server for easy access to the SMU CLI from Windows environments.

To access the command-line interface using encrypted WinRS, type the following command:

```
C:\>winrs -r:<hostname>:<port> -u:<username> -ssl smu
```

Only local SMU users are supported using this access method. The port should be the port number that is specified in the `smu.conf` file for WinRS. The command argument must be `smu`. This indicates that the SMU command shell should be executed.

Authenticating with WinRS

The SMU WinRS server only supports Basic authentication. You must configure the following `winrm/config/client` properties to access the SMU server:

```
C:\>winrm set winrm/config/client {@TrustedHosts="<hostname>" }
```

You must also ensure that the LAN Manager authentication level will support the Microsoft security protocols LAN Manager and NT LAN Manager (designated as LM & NTLM), and use NTLMv2 session security if negotiated.

1. Open the Local Security Policy.
2. From a command prompt, enter `secpol.msc`.

3. Double-click Local Policies.
4. Double-click Security Options.
5. Double-click Network security: LAN Manager authentication level.
6. Select Send LM & NTLM – Use NTLMv2 session security if negotiated from the pull-down menu.
7. Click OK.

Note that you will have to add the SMU service self-signed certificate to the trusted certificate store on your Windows host before connecting. To obtain the SMU service self-signed certificate, perform the following steps. Instructions for both SSH and web browser sessions are shown.

Using SSH:

1. Log in to SMU.
2. Enter **certs get all**.
3. Copy the encoded certificate (encoded text starts with "----BEGIN CERTIFICATE---" and ends with "---END CERTIFICATE---") and paste it into a file.
4. Save the file, using the suggested file extension `.cer`, which is the standard file extension for X.509 certificates.
5. Right-click on the certificate file from Internet Explorer and select "Install Certificate".
6. Make sure the certificate is installed into the "Trusted Root Certificate Authorities" certificate store.

Using your web browser:

1. Point your browser to URL `https://<SMU Host>:8443`.
2. You should be presented with the "Untrusted Connection" page, where you will select the following:
 - Click the Add Exception... button.
 - Click the View button.
 - Click the Details tab.
 - Click the Export button.
3. Save the certificate to a file. The file extension `.cer` is automatically used.
4. Right-click on the certificate file from Explorer and select "Install Certificate".

5. Make sure the certificate is installed into the local computer account of the "Trusted Root Certificate Authorities" certificate store and personal trust store of the Windows host.

Accessing and Authenticating Using the Browser User Interface

The SMU browser user interface is accessed using a standard web browser.

The following table lists standard supported web browsers.

TABLE 8. SUPPORTED WEB BROWSERS

WEB BROWSER	SUPPORTED VERSION
Firefox	2.x through 13+
Internet Explorer	7 through 10
Safari	3.1.2, 4, 5
Chrome	1 through 20+
Safari iOS	4.3.3, 5

To access the encrypted SMU web application, use the following URL:

```
https://<hostname>:<port>/smu
```

The *hostname* is the hostname or IP address of the host upon which SMU is running. The *port* number is the HTTPS port (encrypted application) specified in the `smu.conf` file. The default port is 8443.

The SMU web server uses a self-signed certificate. You must accept the certificate the first time you access the SMU web server using HTTPS.

Once you are connected to the SMU BUI, you will be presented with the SMU login page. To log in to the SMU BUI, enter your SMU user and password. For initial access to the SMU BUI, you can use the user name `admin`, then the default password `changeit`. Once you log in to the SMU BUI, you should be able to change the password on the Users tab of the Administration panel.

Navigating the Browser User Interface

The Snap Management Utility for Oracle Database provides an easy-to-use Browser User Interface (BUI) for managing its operations. Figure 1 shows the basic BUI layout, in which the main window is divided into a Workgroup navigation tree at left, a main or center detailed display for the selected Workgroup item and, at bottom, a Tasks screen which displays details of tasks in process or previously executed.

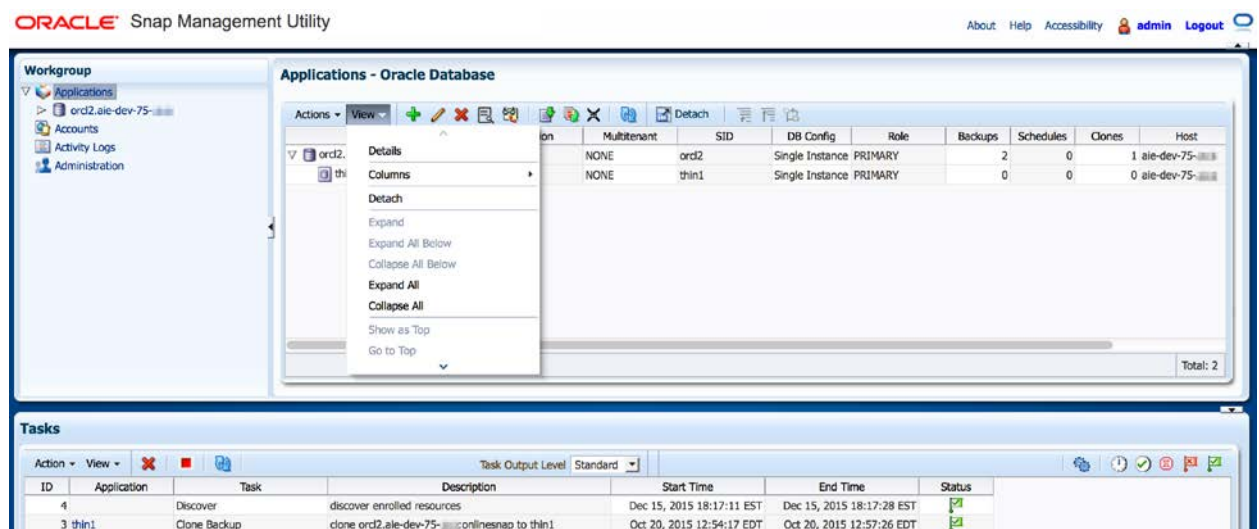


Figure 1. Snap Management Utility BUI showing view options menu under Applications

On the upper right side of the BUI, you can access links for general information about the installed Snap Management Utility (the link labeled [About](#), which provides the SMU version), the online help ([Help](#)), which contains a quick-reference form of the user guide information contained in this document, the accessibility preferences (the link labeled [Accessibility](#)), the logged-in user (“admin” in this example), and the SMU logout.

Clicking on the [Accessibility](#) link generates the Accessibility Preferences dialog window, as seen in figure 2. From this window, you can choose display preferences so that the Snap Management Utility BUI will optimize for the selected accessibility preference.

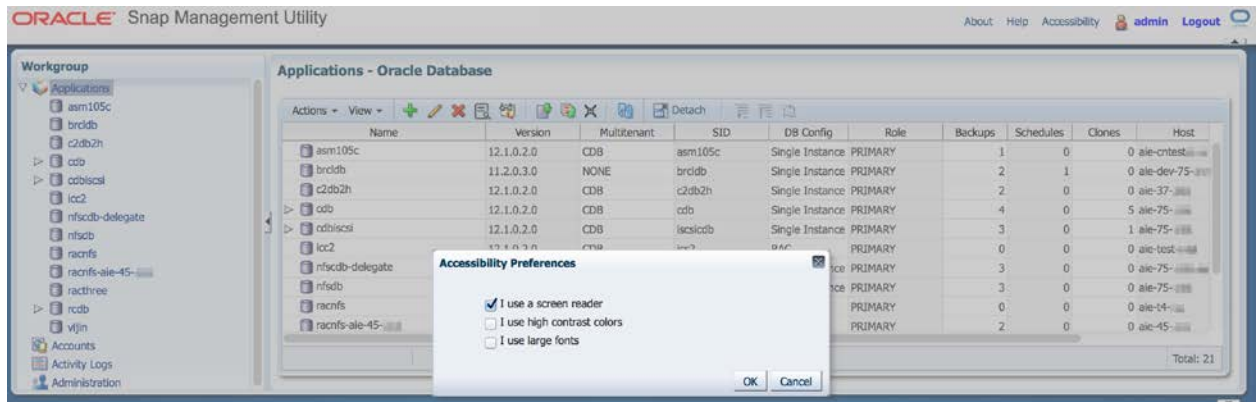


Figure 2. BUI links for online help and accessibility settings, with Accessibility Preferences dialog window displayed

The Tasks screen provides important details for any operations initiated in SMU. Actions such as deleting or canceling a task, or changing or refreshing the display are managed with pull-down menus and selectable icons. On the right side of the Tasks display, icons provide sort options for the Tasks display. Mouse over each of the icons for a text descriptor of the icon's sort parameter. (Note that any task failures have purposefully been shown in this panel for examples.)

The Workgroup panel's navigation tree is subdivided into Applications, Accounts, Activity Logs, and Administration. For each of the left-sided categories under the label Workgroup, the Actions and View pull-down menus, as well as the corresponding icons, will change for the permissible actions within that category.

Both the Applications and the Accounts view screens share an important icon for validation testing of a selected application or account, whether the account is a storage or an application host account. The Test Connection icon is seen in the following figure. The Test Connection action can also be selected from the corresponding Actions menu for either the Applications or Accounts view. These validation tests will test the connection for a particular host or account, providing current status and functionality, availability, and other basic information, and refresh the basic information for the selected storage or host account, or application.

Note: This Test Connection icon appears when either the top-level (parent) label for either Application or Account is selected in the Workgroup navigation tree. You cannot select a specific child (specific application or account) in the Workgroup navigation tree because you are no longer at the Application-level view or Account-level view by doing so.

ORACLE Snap Management Utility

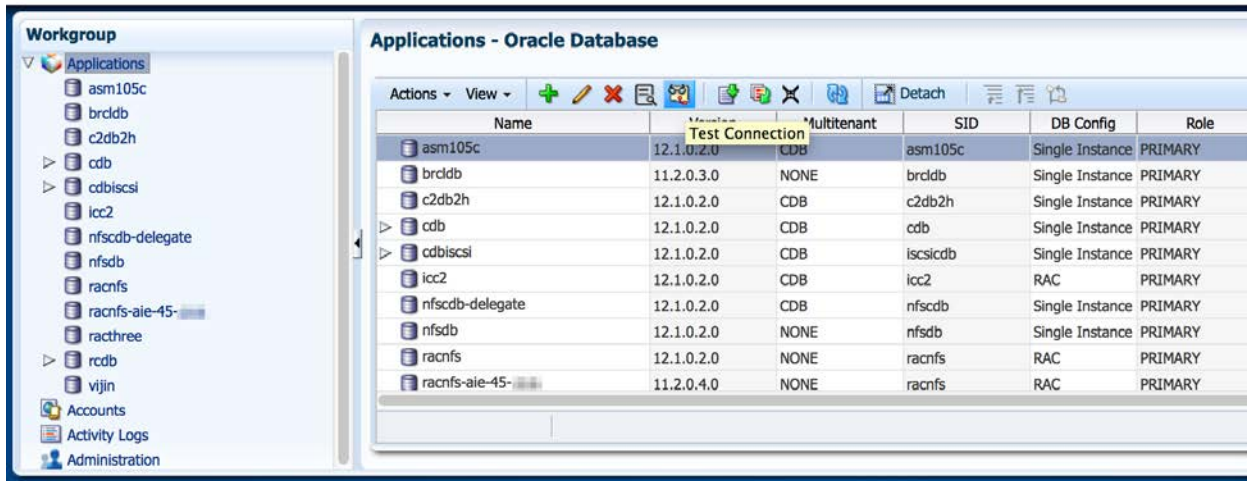


Figure 3. Selecting Test Connection for an application

Next to the Test Connection icon, the Details icon also provides information for highlighted Applications and Accounts (both storage and application host accounts), which can also be retrieved using the Details option in the corresponding Actions menu.

In figure 4 in the Applications view, the Actions pull-down menu allows you to add, modify, remove, retrieve details, or test (check on the accessibility of) an application, as well as import (clone) an RMAN image, and refresh or deprovision a cloned application. Selecting the last menu item, Refresh, refreshes the screen display. Selecting the corresponding icons to the right of the Actions and View pull-down menus also accomplishes the same actions. Mouse over each of the icons for a text descriptor of its function.

ORACLE Snap Management Utility

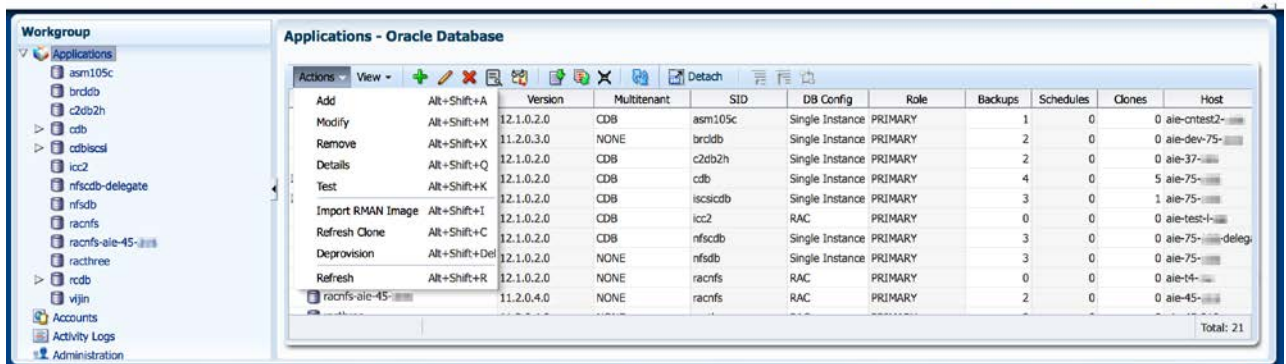


Figure 4. Applications navigation tree expanded and main display with Actions menu chosen

The main Applications panel lists a summary subtree of clones associated with a particular application, so that the relationships of source-target is easily viewed from this window. The clone's system identifier (SID) next to Type identifies the clone's characteristic, such as split clone or thin clone, and icons and prefixes in the database instance name duplicate this information. The Workgroup window's navigation tree can expand in the Applications category to show this information and subordinate relationships as well. You can expand and collapse the Workgroup navigation tree by selecting the directional arrow to the left side of the main/center display.

For a complete list of the BUI icons and their meanings, see Appendix C.

When you select a listed application, the main panel display changes to feature tabs for Snap Backups, where all the backup and clone operations are accessed, Schedules, where timeframes for scheduled operations are set, and Account Settings, which displays and accepts modifications for Application, Application Host, and Storage information.

The following sections provide basic information and procedures for common tasks within the four workgroup, or feature, categories.

Managing Accounts Using the BUI

SMU accomplishes snap backup operations by coordinating functions of the application (database), the application host where the application is running, and the storage system where the application's data resides. The accounts of applications are managed on the Applications page while the accounts of application hosts or storage systems are managed on the Accounts page.

An application account is associated with an application host account and a storage system account. An application host account and a storage system account can be associated to many application accounts. Because of this relationship, the application host account and the storage system account must be created before an application account associated with them is created. An application host account or a storage system account cannot be deleted while they are in use by an application account.

Important: Be sure to review the section "User Permission Requirements for Accessing Operations" in this document for further information on accounts. Further information on delegation is also included in that section.

The SMU BUI's Accounts panel, seen in Figure 5, contains tabs for the registered host accounts and storage accounts. These accounts must be created before taking any snap backups.

Note that the name of an account, including host, storage, and application accounts, must be unique within an SMU instance. An account name consists of alphanumeric characters or one of these permissible special characters: _ (underscore), . (period), @, and - (hyphen). An account name may not start with a number. The maximum account name length is 64 characters.

Selecting the Application Host Accounts tab lists all the accounts created for connecting to the hosts. The displayed columns under Application Host Accounts show the following items:

Account (Name) – Unique name of the account to access the application host.

Host Name – Domain name or IP address of the application host.

Type (previously called protocol) – Either SSH2 or WINRS.

OS Type – Indicates whether the application host is a UNIX-based (Oracle Linux or Oracle Solaris) or a Windows application host.

Release – Release number of the host operating system.

Delegate – Indicates whether the host requires a delegation tool for access; either NONE or SUDO.

User – User name to use to establish the connection to the application host.

Applications – Number of applications assigned under the selected account. Note that this number is a hyperlink which, when selected, produces a pop-up application list showing the associated application names.

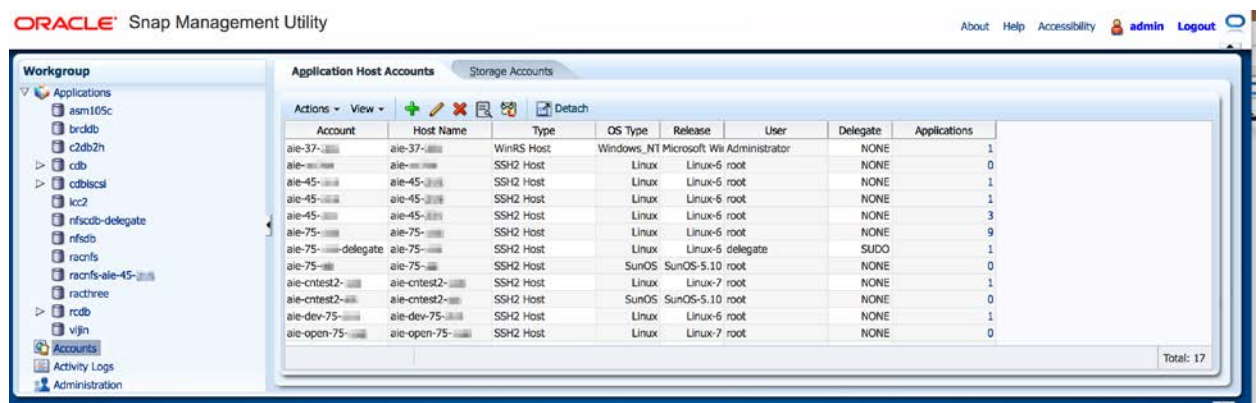


Figure 5. Application Host Accounts tab listing under Accounts in SMU BUI

Similarly to the Application Host Accounts tab, the Storage Accounts tab lists all the accounts created for connecting to the available Oracle ZFS Storage Appliance(s). The provided column information, as seen in the following figure, includes:

Account – Unique name of the account to access the storage system (Oracle ZFS Storage Appliance).

Storage Name – Domain name or IP address of the Oracle ZFS Storage Appliance.

Type – The type of storage system. Type is ZFS Storage Appliance by default.

Version – The release version of the operating system running on the Oracle ZFS Storage Appliance.

Cluster – Designates a clustered storage system configuration; either Yes or No.

User – User name to use to establish connection to the Oracle ZFS Storage Appliance.

Applications – Number of applications assigned under the selected account.

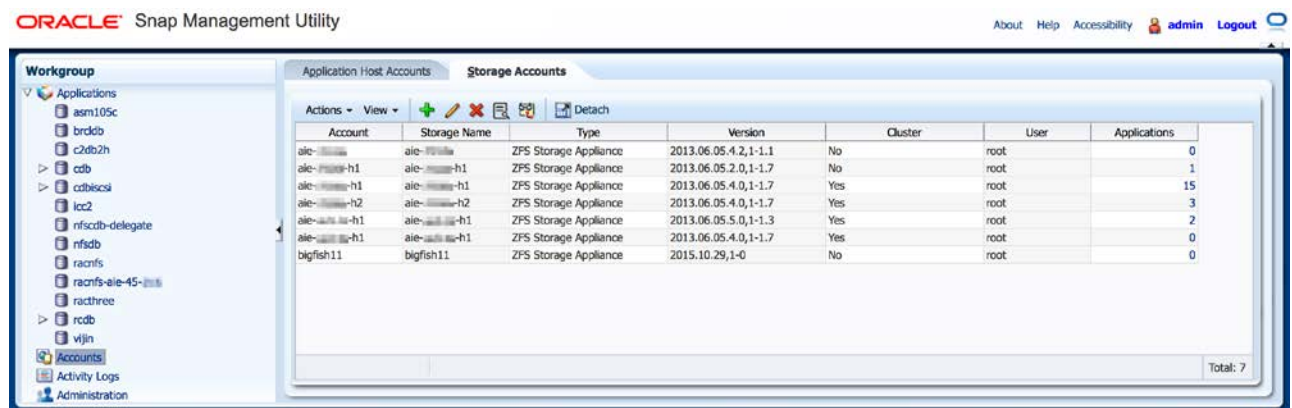


Figure 6. Listing for Storage Accounts tab

From the Accounts BUI page, you can easily perform the following tasks under the Application Host Accounts and Storage Accounts tabs by following the provided steps.

Adding a new host account

1. Selecting the Application Host Accounts tab, choose Actions, and then Add.
2. Provide entries for the following: a unique name for Account, Host Name, Protocol (either SSH2 or WINRS), Port, Delegation, User, Password, then Confirm password.
3. Click on OK.

Modifying a host account

1. From the list of application host accounts, select and highlight a row (host account) to edit.
2. Click on Actions, then Modify.
3. Make any needed changes to the fields displayed in the Modify Application Host Account window that displays and click OK.

Testing an application host account

1. From the navigation tree, select Accounts. This shows all the accounts previously enrolled in SMU in the main Accounts window.
2. Select the Application Host Accounts tab.
3. From the list of enrolled accounts in the main Application Host Accounts window list, select and highlight a row (application host account) to test.
4. Click on Actions, then Test, or select the Test Connection icon.

Once the test is complete, the resulting information is displayed.

Deleting a host account

Note that when an application host account is in use by an application, the host account cannot be deleted. Attempting to delete a host account in use results in an error.

1. From the list of application host accounts, select and highlight the row (host account) to delete.
2. Click on Actions, then Remove.
3. A warning message displays, indicating that deleting a host account will disable snap backup operations on the applications running on the host. Click on Yes to confirm that you want to delete this host account.

Adding a new storage account

For clustered storage: be sure to reference the section "Special Considerations with Clustered Systems" at the beginning of this document.

1. Selecting the Storage Accounts tab, choose Actions, and then Add.
2. Provide entries for the following: a unique Account name, Storage Name, Type (no entry required; only Oracle ZFS Storage Appliance is supported and automatically populates), Port number, User name, Password, then Confirm Password.
3. Click on OK.

Testing a storage account

1. From the navigation tree, select Accounts. This shows all the accounts previously enrolled in SMU in the main Accounts window.
2. Select the Storage Accounts tab.
3. From the list of enrolled accounts in the main Storage Accounts window list, select and highlight a row (storage account) to test.
4. Click on Actions, then Test, or select the Test Connection icon.

Once the test is complete, the resulting information is displayed.

Modifying a storage account

1. From the list of storage accounts, select and highlight a row (storage account entry) to edit.
2. Click on Actions, then Modify.
3. Make any needed changes to the fields in the Modify Storage Account window and click OK.

Deleting a storage account

Note: SMU will not delete a storage account unless the associated application account has already been deprovisioned or deleted.

1. From the list of storage accounts, select and highlight the row (storage account) to delete.
2. Click on Actions, then Remove.
3. A warning message displays, indicating that deleting a storage account will disable snap backup operations on the applications that consume the storage. Click on Yes to confirm that you want to delete this storage account.

Managing Applications Using the BUI

The following figure shows the display for the selection Applications, with the pull-down menu for choosing options for Views. Basic information such as the application's host account, storage account, and whether it has a backup schedule or existing snap backups, is readily seen for the applications/databases registered. The columns list the following information:

Name – Unique name of the account to access the application (database) running on the application host.

Version – Version of the application (database). Currently, only Oracle Database is supported. The example shows Oracle Database 11gR2 and 12c versions.

Multitenant – Indicates whether the application is a multitenant database, holding pluggable databases.

SID – The database system identifier (SID) or SID prefix without node index number in case the database is a Real Application Cluster (RAC) or RAC One Node. Identifies the characteristic of a clone.

DB Config – Indicates whether the database host is a clustered system or not (single instance).

Role – Lists either PRIMARY or STANDBY application (database).

Backups – Indicates the number of existing copies of the application (database) SMU has already taken.

Schedules – Indicates the number of existing snap backup schedules configured for the application (database).

Clones – Indicates the number of all clones created from this application (database).

Host – Account name of the application host.

Storage – Account name of the storage system.

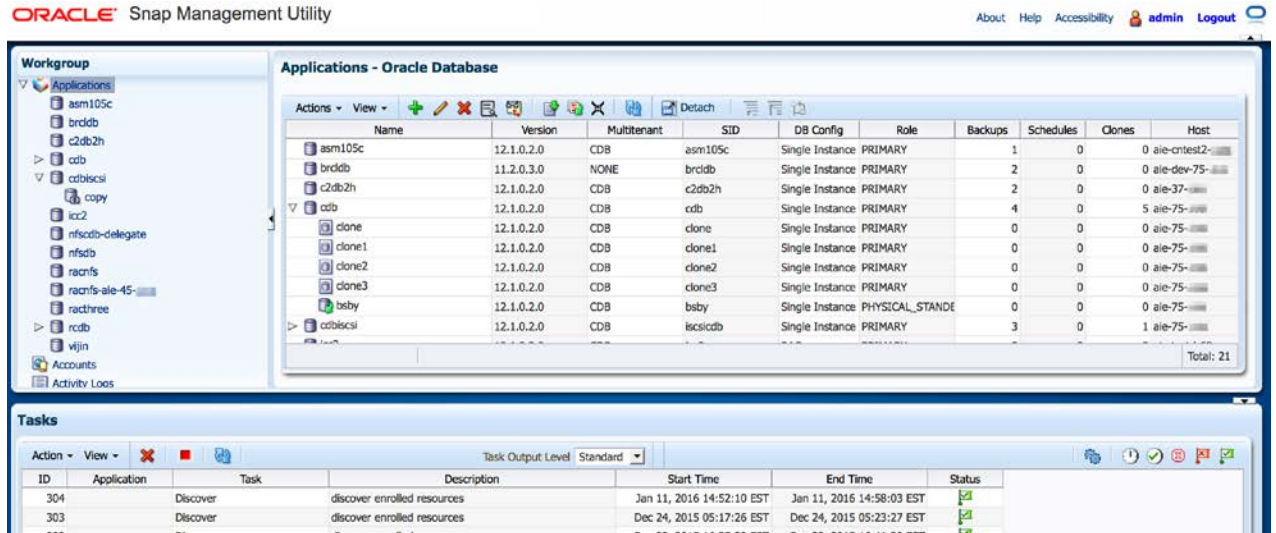


Figure 7. Viewing the list of Applications, their related clones and snap backups, and details for all

Important: Note the test icon – A clickable icon that will check validation of application credentials, appearing in the Actions and View line's list of icons. Be sure when you add or modify accounts to use the test icon to validate the changes.

From this Application BUI page, you can easily perform the following tasks by following the provided steps.

Enrolling a new application for snap backups

Note that an application host account and a storage system account must already exist so that an application account can use them when the application account is created. If you have not created the application host account and the storage system account to use with the application account, you can create them using the instructions in the "Adding a new host account" and the "Adding a new storage account" sections.

1. From the left-side Workgroup panel/navigation tree, click on Applications.
2. Click on Actions, then Add. The Add Application Account dialog window is displayed.
3. Fill in entries for application account Name, Type (no entry required; automatically populates with Oracle Database), SID/SID Prefix, Database Unique Name, Listener Port, Database Configuration (drop-down window to select among Single Instance, RAC, RAC One Node), User (populates automatically with default sys user), Password (password of the sys user), Confirm Password, and drop-down windows for Host Account and Storage Account.

Note that SID Prefix without RAC node index must be specified for the SID/SID Prefix field in case of RAC or RAC One Node database.

4. Click OK. The newly enrolled application appears in the Application table and also under Applications listed on the left side of the panel.

Testing an application

1. From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU in the main Application window.
2. From the list of enrolled applications in the main Application window list, select and highlight a row (application) to test.
3. Click on Actions, then Test, or select the Test Connection icon.

Modifying an application

1. From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU.
2. From the list of enrolled applications, select and highlight a row (application) to edit.
3. Click on Actions, then Modify.
4. Make any needed changes in the fields displayed in the Modify Application Account window and click OK.

Browsing application details

1. From the list of enrolled applications, select and highlight a row (application) to view the application's details.
2. Click on Actions, then Details.

A pop-up window will display all the related configuration information, including its host and storage information.

Removing an application

Note that you cannot delete an application account when the application has snap backups or backup schedules configured on it. You should delete all snap backups and backup schedules prior to deleting the application. Attempting to delete an application when the application has snap backups or backup schedules results in an error.

1. From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU.
2. From the list of enrolled applications, select and highlight a row (application) you want to remove.
3. Click on Actions, then Remove.
4. A warning message displays, indicating that deleting an application will disable snap backup configuration for the application. Click on OK to affirm that you want to remove this row/application.

Importing an RMAN backup image

Note that as part of the RMAN Backup Image import operations, an application account for the clone application is created. A new node for the clone application appears under the Applications node of the navigation tree. The application summary table also lists the newly created clone application. If an RMAN Backup Image import operation fails, the clone application node is removed upon the page being refreshed in the next polling cycle. You can press the Refresh button to make sure you have the latest updated application summary.

Important: If you intend to take online backups of clones created by importing an RMAN backup, the backed-up datafiles and archived logs must be in separate backup shares. This setup can be accomplished by allocating files to different RMAN channels. The following RMAN run block script employs the channel parameter to separate the datafiles and archived logs and should be run:

```
"
connect target /
.
configure controlfile autobackup on;
run {
# this undoc command will ensure that no autobackup is generated
# at the end of this script
set nocfau;
allocate channel ch01 device type disk format '/rman01/%U';
allocate channel ch02 device type disk format '/rman02/%U';
backup as copy database channel ch01 plus archivelog channel ch02;
backup as copy current controlfile channel ch01;
}
"
```

IMPORTANT: Currently, importing an RMAN backup image from a RAC One Node configuration is not supported.

1. From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU.

To clone from an RMAN image copy that is stored in multiple shares, you can specify one or more share mountpoints, separated by commas, as the value of the RMAN Image path, which is the list of filesystem shares in which the backup is contained. Shares are identified by their mountpoint property (a path like `/export/backup`).

Note that for multihead cases, SMU will check the heads automatically as it searches for the specified shares. Users can choose either head of a dual-head system.

2. Click on Actions, then Import RMAN Image. The Import RMAN Backup Image wizard's dialog windows will display.

3. Required entries are prompted in two steps: locating the RMAN image and specifying the database account information. Choose the RMAN image's storage location(s) from a drop-down list, enter a value for RMAN Image Path or, in the case of multiple-share storage locations, provide a comma-separated list of values, and click on Next. The step 2 dialog window requires the following to create the new database account for the imported RMAN backup image:
 - Host – Choose from the drop-down list displaying all registered host accounts.
 - Account Name – In which to create the RMAN clone database.
 - Database Name – Database name for the clone database to use. (This is the DB_NAME setting for the clone database).
 - SID/SID Prefix – If the clone is a single instance, use the specified SID as is. If the clone is a RAC (cluster), the specified SID prefix accompanies the index numbers used to assign the SID to members of the RAC.
 - Database Unique Name – Global name of the clone database. (Can be the alias name; maps to the DB_UNIQUE_NAME.)
 - Listener Port – The port number of SQL*Net listener. The default setting is 1521, but it can be changed.
 - Database Configuration – Single Instance, RAC, or RAC One Node.
 - Database Home – Path to the Oracle Database software home.
 - User – Automatically populated with `sys`.
 - (DB) Password.
 - Confirm Password.
4. Click Next. Step 3 of the RMAN Import wizard provides a dialog box with the following required settings:
 - System Global Area Size (in the format `integer|K|M|G`) – Specifies the shared memory for the database instances.
 - Open Mode – Options are Read Write, Read Only, and Mounted for database access settings.
 - Log Mode – Sets archiving on (ArchiveLog) or off.

A selectable Information text label that provides details on the settings is listed below the required fields.

5. When these fields are properly populated, select Next to save them and go to the next step, which provides a Summary Review of the entered information.

- To accept the summary review information and proceed with the action, click Finish. The newly enrolled application should appear in the Application table as well as under the left-side Workgroup panel, under Applications.

Note: Selecting Actions and Refresh will reload the BUI page so that the latest updated Applications information appears onscreen.

Refreshing a clone

This operation will update, or refresh, the selected clone to a more recent backup of the source database. The refresh operation combines the required two-step process for creating these updates – deprovisioning the existing clone and then provisioning a new clone – into one action (refresh clone). It is only performed on thin clones.

- From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU.
- From the list of enrolled applications, select and highlight the particular clone that falls under the application of interest for the refresh.
- Select Actions, then Refresh Clone.
- A pop-up dialog window will be displayed, with an entry field for the database backup to use to update the clone. Select the appropriate backup from the available choices populated in the pull-down menu. The dialog window requires you to affirm the selection and operation in a checkbox. Check the box and select OK to execute the operation.

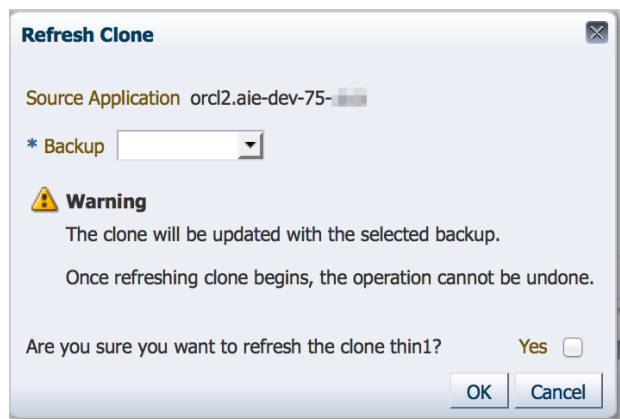


Figure 8. Selecting the source database backup for a clone refresh operation on example clone named thin1

Deprovisioning an application (clone only)

Note that only clones created from SMU can be deprovisioned. Deprovisioning removes all artifacts of the clone as well as the application account for the clone. Once a deprovision operation gets started, it cannot be undone. Canceling a deprovisioning operation may leave the clone in an unknown state. SMU should refresh the BUI page when the deprovision operation has completed and the next polling cycle is hit. You can press the Refresh button to make sure you have the latest updated application summary.

From the navigation tree, select Applications. This shows all the applications previously enrolled in SMU.

1. From the list of enrolled applications, select and highlight a row (application) to deprovision.
2. Click on Actions, then Deprovision.
3. A warning message displays, indicating that deprovisioning deletes a clone from the host and detaches the backend storage. You are prompted to fill in the clone name that you wish to deprovision. Provide the clone name and click on OK to affirm that you want to deprovision this application from the host and storage accounts.

Managing Administration Using the BUI

All user administration activities can be performed in the Administration screen. Administration is the last category selection under Workgroup on the left side of the screen. The Administration screen contains three tabs: Users, Notification, and General Settings.

Only the Snap Management Utility administrator has privileges to add, modify, or remove other users' access to the SMU. Clicking on the Users tab brings up a table that shows all the existing users of the SMU. The following columns are displayed in the table, as seen in figure 9:

User – The user name is used for login to the SMU.

Type – Indicates whether the user is a local or LDAP user.

Full Name – The actual name of the user.

Directory Server – The Directory Server name that holds this user information.

Directory – For LDAP users, the location where their user details reside on the Directory Server (the LDAP search base for the user authentication).

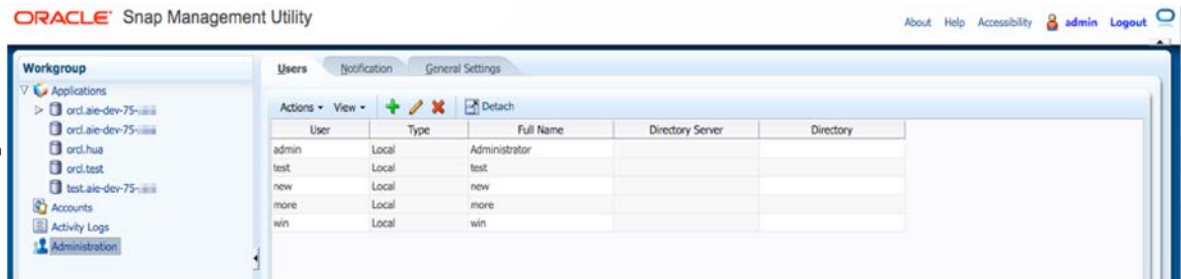


Figure 9. Administration information under Users tab

The following BUI instructions provide details on adding, modifying, or deleting user entries in the Snap Management Utility.

Adding a new user from the Users tab

1. Selecting the Users tab, choose Actions, and then Add User.
2. Choose either the Local User or LDAP User tab.
3. Provide entries for the following for Local User: User, Full Name, Password, then Confirm Password. For LDAP User, provide the User, and user's Directory Server and Directory.
4. Click on OK.

Modifying a user entry from the Users tab

1. From the list of users, select and highlight a row (user) to edit.
2. Click on Actions, then Modify.
3. Make the needed changes to the fields displayed in the Modify User dialog window and click OK to commit to the changes.

Deleting (Remove) a user from the Users tab

1. From the list of user accounts, select and highlight the row (user) to delete (remove).
2. Click on Actions, then Remove.
3. As indicated in the warning message that displays, once the user entry is deleted, the user will no longer be able to log in to the SMU. Click on OK to confirm that you want to remove this user.

SMU allows users to subscribe to notices for any important events that might occur while the SMU is running the tasks initiated by the users. For example, users can subscribe to receive email alerts when any SMU task is cancelled or failed.

The display for Administration's Notification Tab presents two columns, as seen in figure 10: a list of all the events the users subscribed to (Subscribed Events) and the email addresses to which notices are sent (Email Address).

ORACLE Snap Management Utility




Figure 10. Viewing email alerts subscriptions in the Notification window under Administration

The following instructions for adding, modifying, and removing notifications are similar to the Users add, modify, and remove instructions.

Adding a new notification subscription

1. Selecting the Notification tab, choose Actions, and then Add.
2. Provide entries for the following: Event Types (Choose among the options All, Task Cancelled, and Task Failed. Selecting All will initiate email notices for all SMU events; currently only Task Cancelled and Task Failed events are supported.), and Email (to which the notification is sent).
3. Click on OK.

Modifying a notification subscription from the Notification tab

1. From the list of notification subscriptions, select and highlight a row (subscription) to edit.
2. Click on Actions, then Modify. The dialog box 'Modify Notification Subscription' is displayed.

3. Modify the details to the following elements:
 - Event Types: Currently SMU can send the email alerts when a task is canceled or when it is failed. Selecting the “All” checkbox will let you receive all supported notification events from SMU. You can individually select subscription to either “Task Cancelled” or “Task Failed” using the checkboxes. Unselecting the checkboxes removes the subscription to those events.
 - Email: Change the email address to which notices are sent.
4. Make the needed changes to the fields and click OK.

Deleting a notification subscription from the Notification tab

1. From the list of notification subscriptions, select and highlight the row (notification subscription) to delete.
2. Click on Actions, then Remove.
3. Click on OK to confirm that you want to remove this notification subscription.

The changes to Notification will become effective immediately and the user will start (or stop, if that is the case) receiving the subscribed notifications whenever they occur on SMU.

Managing status polling, and task and activity log displays in the General Settings tab

SMU operates using asynchronous-based commands. For example, when SMU creates a backup or a clone of a database, it automatically determines the dependent mount points and the corresponding LUNS and shares on the Oracle ZFS Storage Appliance and initiates the backup process on the Oracle ZFS Storage Appliance. Depending upon the size of the database to be copied, the operation can take a considerable amount of time. In order to allow other operations on the SMU, every user-initiated activity related to the database backup is designed to work asynchronously. The moment any such operation is started, the task is listed in the Tasks panel at the bottom of the SMU BUI. Then onwards, the SMU BUI periodically polls the SMU server and/or Oracle ZFS Storage Appliances to check the status (or state) of the initiated tasks and update their status information in the Tasks panel.

In the General Settings tab’s display, as seen in the following figure, users can use the Browser Auto-Refresh panel to set the poll intervals for their own individual settings. The general settings are stored and applied across the SMU logins of that particular user. The following fields are used:

Disable Polling checkbox – Checking and unchecking this box will, respectively, disable and enable the polling process.

Polling Interval (secs) – Indicates how often the SMU BUI should poll the SMU server and the Oracle ZFS Storage Appliances to determine the running status of the task. A user can enter a new value either by typing the value or using the increment/decrement arrows.

Polling Timeout (secs) – Indicates the amount of time in seconds to pause polling when the user is not actively using the SMU BUI. The polling timeout allows the user login session to expire when the user does not use the BUI for awhile. The actual user login session expiration time is determined by adding the polling timeout to the SMU BUI session timeout (30 minutes). Users can enter a new value either by typing the value or using the increment/decrement arrows.

Once the Update button is selected, the new changes take effect immediately.

IMPORTANT: If SMU is managing a heavy load of items and the BUI auto-refresh polling interval setting is too short, the BUI may keep reloading a page before it can render it. This issue may block navigation to the General Settings panel and, thus, the ability to change the auto-refresh setting.

To address this, disable the auto-refresh BUI process by changing the `bui.autorefresh.disable` parameter to `true` in the `$SMU_HOME/etc/smu.conf` file. After changing the parameter, restart SMU, then navigate to the General Settings tab and adjust the auto-refresh polling interval to a larger span.

The next panel, Task Monitoring, provides control over the viewable number of tasks in the Tasks window display using the Maximum Task Display settings. Either enter the desired maximum line count or use the increment/decrement arrows to change this setting to customize this display for your particular screen display parameters.

Similar display controls for number of rows per page can be set in the last two panels, Snap Backups and Activity Logs.

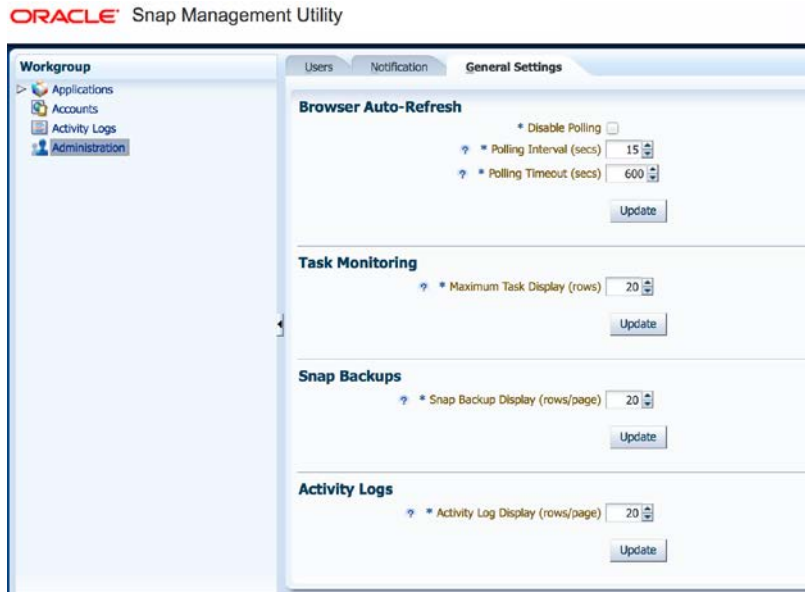


Figure 11. Polling frequency and various maximum line display settings under Administration and General Settings tab

Operating and Managing Snap Backups and Clones in the BUI

Backups can be created manually or automatically. Automatic backups can be scheduled to occur on a recurring basis. Automatic backups can also have a retention policy setting the number of backups to retain at a time.

Two types of backups are supported. Offline backups are backups taken when the database is shut down. The software will shut down the database temporarily and then restart it after taking the snap backups. Online backups are backups taken when the database is placed into backup mode while remaining online. Online backups snapshot the database shares in a particular order.

In the BUI, indicate which application to target for a snap backup by clicking on Applications on the left-side panel under Workgroup panel, then selecting the application from the list by clicking on it. This will bring up another panel in the middle of the screen with three tabs: Snap Backups, Schedules, and Account Settings, as seen in the following figure.

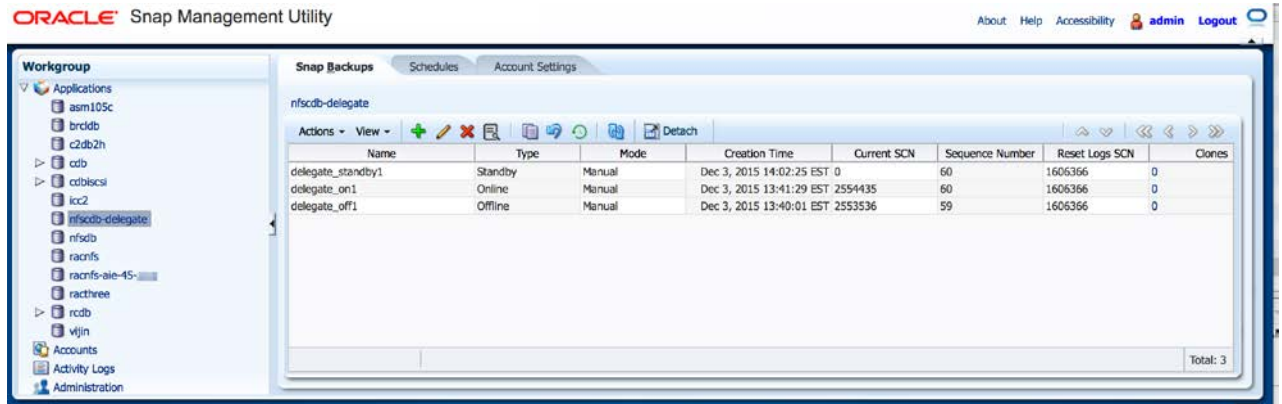


Figure 12. List of snap backups for selected application

The Snap Backups tab provides a table displaying the list of snap backups taken for the selected application. The table includes the following details:

Name – The unique name for the snap backup that will be taken. The name consists of alphanumeric characters, with a maximum length of 64 characters, and including permissible special characters (only "_" [underscore] and "." [period]). Snap backup names must not begin with a number.

Type – Indicates the type of snap backup (the method to create the snap backup): Online, Offline or Standby.

Mode – Indicates whether the snap backup was scheduled for automatic execution or manually initiated.

Creation Time – The time when the snap backup was created.

Current SCN – Indicates the database system change number at the time of the backup.

Sequence Number – Indicates the database log sequence number at the time of the backup.

Reset Logs SCN – Indicates the database system change number when the logs were last reset; this is used to identify the database incarnation. A database incarnation is created whenever you open the database with the RESETLOGS option. After a database point-in-time recovery or recovery with a backup control file, you must open the database with the RESETLOGS option, thereby creating a new incarnation of the database. The database requires a new incarnation to avoid confusion when two different redo streams have the same SCNs, but occurred at different times. If you apply the wrong redo to your database, then you will corrupt it.

Clones – Number of clones that have been created from this snap backup. The number is a hyperlink that, once selected, produces a separate pop-up window showing the list of clones and their associated clone account names. Each clone account name is also a hyperlink which, when selected, produces the backup information window for the selected clone account.

On the right side of the display, scroll arrows for up, down, first, last, previous and next task items are selectable. The Jump To tab, when selected, presents a dialog box to choose a task ID for display.

The following instructions include basic operations for taking and managing snap backups and clones in the Snap Management Utility BUI.

Creating a New Snap Backup

1. Click on Actions, then select Backup. The following figure shows the possible actions, including Backup, in the Actions pull-down window.

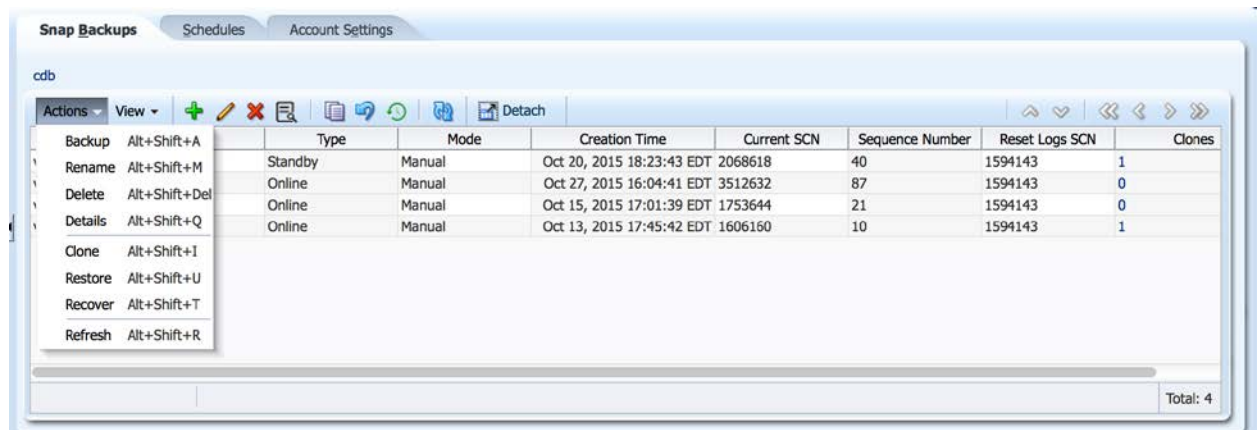


Figure 13. Selecting an operation in the Actions tab list

2. A new dialog window pops up, as seen in the following figure. Enter a unique name in the Name field for the snap backup you are creating.

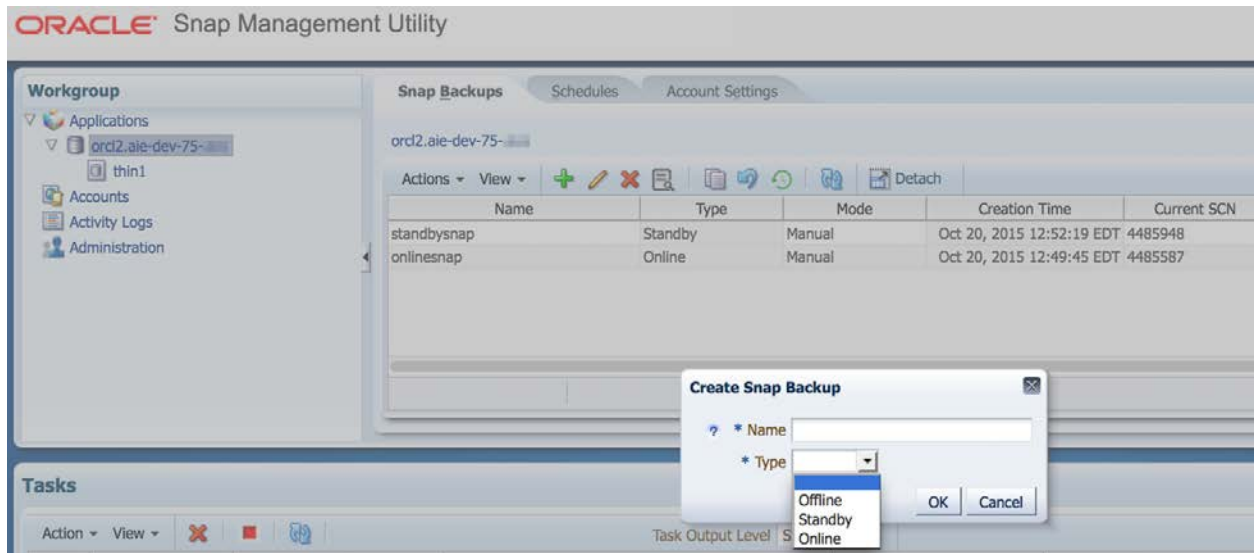


Figure 14. Specifying name and type to create a new backup

3. Select the Type from the drop-down list for the snap backup you want to create: Offline, Online, or Standby.
4. Click OK.

You will see a task in the Tasks Panel at the bottom of the screen that has been created for taking this snap backup. If the snap backup succeeds, it displays in the table under the Snap Backups tab, and the corresponding task in the Tasks panel displays a green check in the checkbox in the Status column. If the snap backup task fails, the Output column indicates a red crossmark (x). Clicking on the status indicator for the particular operation provides a task output popup display which, in the case of a failed status, includes the reason for the failure within the detailed task output.

The following figure shows a series of successful operations, including backups, displayed in the Tasks pane.

ID	Application	Task	Description	Start Time	End Time	Status
4		Discover	discover enrolled resources	Dec 15, 2015 18:17:11 EST	Dec 15, 2015 18:17:28 EST	✓
3	thin1	Clone Backup	clone orcl2.aie-dev-75- onlinesnap to thin1	Oct 20, 2015 12:54:17 EDT	Oct 20, 2015 12:57:26 EDT	✓
2	orcl2.aie-dev-75- Backup	Backup	backup orcl2.aie-dev-75- as standbysnap	Oct 20, 2015 12:51:01 EDT	Oct 20, 2015 12:52:19 EDT	✓
1	orcl2.aie-dev-75- Backup	Backup	backup orcl2.aie-dev-75- as onlinesnap	Oct 20, 2015 12:47:42 EDT	Oct 20, 2015 12:49:45 EDT	✓

Figure 15. Listing of backup operation results in the Tasks window

Renaming a snap backup

1. From the list of displayed snap backups, select and highlight the row (backup) that you want to rename.

2. Click on Actions and Rename.
3. A dialog window labeled Rename Snap Backup will be displayed. Provide the new name in the Name field.
4. Click on OK.

IMPORTANT: If a snap backup has been cloned, the rename of the snap backup cannot be performed. SMU will block this rename operation.

Cloning a Snap Backup

SMU can be used to clone a database and create a new primary database. This database is created using Oracle ZFS Storage Appliance clone technology and, depending on the type of clone operation that is chosen, can leverage ZFS cloning, remote replication service, SMU rollback features, and a combination of all of them. SMU currently supports three options for clone creation: Primary Thin Clone, Primary Clone Copy, and Data Guard Standby Clone.

Clones can be either thin-cloned datasets that share common data blocks with the original dataset – which allows the clones to be space efficient and created very quickly – or they can be a "full copy" – fully duplicating (replicating) the database, so that the clone copy will take up the same amount of storage at its target location as the original database does in its location. (For this reason, be sure to allocate the correct size for the target storage pool.) The primary thin clone operation, true to its name, produces thin-cloned datasets, while clone copy and Data Guard standby clone operations produce full replicated datasets that are independent of the source database.

The Supported Operations section near the beginning of this document provides further details on processes for each of the different clone operations.

SMU supports cloning single instance, RAC and RAC One Node databases, which are differentiated in the `cluster_database` property setting in their respective application accounts. When the property is set to true, the application account represents a clustered (RAC) database. When the property is set to false, the application account represents a single instance database. When selecting the target account in the cloning dialog window, you do not need to explicitly specify (define again) whether the selected target is a single instance, RAC or RAC One Node.

Important: If you are cloning an instance of Oracle Database that uses its transparent data encryption (TDE) feature, you must manually copy the encrypted database's wallet file to the target host, modify the `sqlnet.ora` file to point to the wallet file, then restart the clone. See Appendix D: Cloning Wallet Files for an Encrypted Database of this document for further information.

Ensure that the targeted host has adequate resources, including shared memory, to handle the addition of another database.

In order to create a database clone from a snap backup on another host, the target host must meet certain requirements:

- The Oracle user must have the same uid/gid as the source database host, since cloning a snap backup produces a new share with the same files and ownership as the original share.
- If the Oracle home is different on the target host, you must specify the Oracle home as a clone option. By default, SMU uses the same Oracle home setting as the source database.
- The Oracle Database software must be the same version as on the source database host. SMU does not perform downgrades or upgrades during the clone operation.

If dNFS is desired for the clone database, it must be configured in the target Oracle home before the clone operation is performed.

Use the following steps to clone a database.

1. From the navigation tree, click “Applications” and click on the application for which a snap backup(s) has already been taken.
2. From the list of the snap backup entries that are displayed, select and highlight the row (snap backup) that you want to clone.
3. Click on Actions and Clone. The pop-up dialog window, called Create Clone, will provide three major options, as seen in the following figure, for the type of clone to create:
 - Primary Thin Clone
 - Primary Clone Copy
 - Data Guard Standby Clone

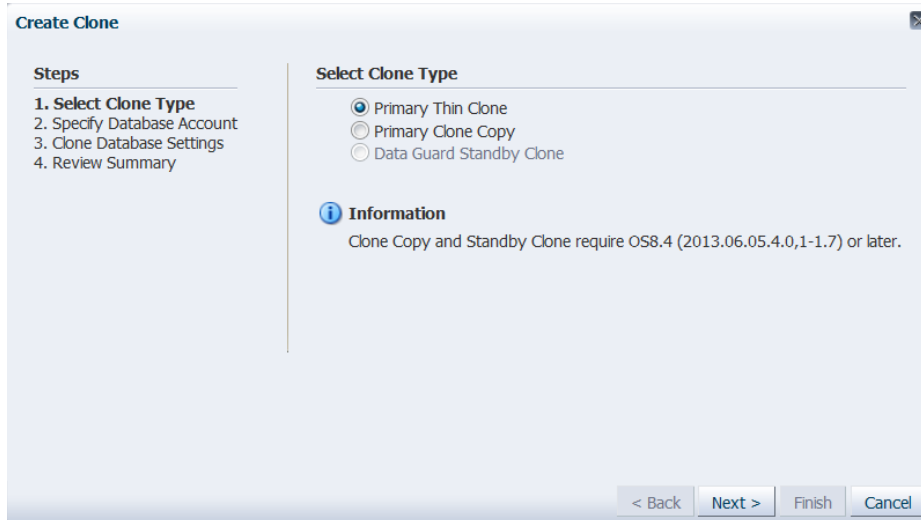


Figure 16. The Create Clone dialog window listing type of clone to create

Note that for each of these choices, the required chain of steps is listed on the left side of the window. All of these clone creation operations share four entry steps, prompted by a dialog window, in common. Two additional steps, as noted in the following list, exist for Primary Clone Copy and Data Guard Standby Clone. The entry windows/steps are:

- Select Clone Type
- Specify Database Account
- Clone Database Settings
- Select Storage (for Primary Clone Copy and Data Guard Standby Clone type options only)
- Select Storage Pool (for Primary Clone Copy and Data Guard Standby Clone type options only)
- Review Summary

Choose the button next to the desired type of clone to create, and select Next.

4. . The Specify Database Account window displays the required field entries:

- Host Account – Choose from the drop-down list displaying all registered host accounts.
- Application Name – In which to create the clone database.
- Database Name – Automatically populates (Oracle Database is the default).
- SID – If the clone is a single instance, use the specified SID as is. If the clone is a RAC (cluster), the specified SID prefix accompanies the index numbers used to assign the SID to members of the RAC.

- Database Unique Name – Global name of the clone database.
- Listener Port – The port number of SQL*Net listener. The default setting is 1521, but it can be changed.
- Database Configuration – Select from the pull-down menu: Single Instance, RAC, or RAC One Node.
- Database Home – Path to the Oracle Database software home.
- User – Automatically populated with `sys`.
- Password
- Confirm Password

The following figure shows the Specify Database Account - Primary Thin Clone dialog window.

Figure 17. Specifying Database Account information for a Thin Clone operation

The following figure shows the Specify Database Account - Clone Copy dialog window. Note that this window displays the additional two clone operation steps for a clone copy or standby clone operation compared to a thin clone operation. But the entries required under Specify Database Account are the same for all three types of clone operation.

Create Clone

Steps

1. Select Clone Type
- 2. Specify Database Account**
3. Clone Database Settings
4. Select Storage
5. Select Storage Pool
6. Review Summary

Specify Database Account - Clone Copy

* Host Account

? * Application Name

? * Database Name

? * SID

? * Database Unique Name

* Listener Port

* Database Configuration

? * Database Home

User

* Password

* Confirm Password

< Back Next > Finish Cancel

Figure 18. Specifying Database Account information for a Clone Copy operation

5. Upon completing the required fields, select Next to continue to Step 3, Clone Database Settings.

Create Clone

Steps

1. Select Clone Type
2. Specify Database Account
- 3. Clone Database Settings**
4. Review Summary

Clone Database Settings - Thin Clone

* System Global Area Size(int|K|M|G)

Open Mode

Log Mode

Information

#Clone database settings

< Back Next > Finish Cancel

Figure 19. Providing clone database settings in step 3 of a clone operation

6. The Clone Database Settings window has the following required fields:

- System Global Area Size (in the format integer|K|M|G) – Specifies the shared memory for the database instances.

- Open Mode – Options are Read Write, Read Only, and Mounted for database access settings.
- Log Mode – Sets archiving on (Archivelog) or off.

A selectable Information text label that provides details on the settings is listed below the required fields.

When these fields are properly populated, select Next to save them and go to the next step. In the case of thin clone operations, it is the final step, Create Clone Review Summary. But for either clone copy or Data Guard standby clone operations, two additional specification steps are required.

7. For clone copy or standby clone operations, the Select Storage window appears, with a prepopulated field displaying the source database and an entry field with a populated pull-down menu of choices for Target Storage. As previously noted, target storage must already be configured before a clone copy operation can be started. After choosing the target storage, select the Next button to save and proceed to the Select Storage Pool step.



Figure 20. Selecting target storage for a clone copy operation

The Select Target Pool dialog window, as seen in the following figure, is displayed and prepopulated with the source storage, target storage, and storage project name under which the source database is located. The selectable entry field with a pull-down menu of choices is labeled Target Pool. A field labeled Target Project is prepopulated but editable, if changing the target project name is desired. The default name is in the format `smu-<db_name>`. The source project and target pool should be mapped in a 1-to-1 relationship.

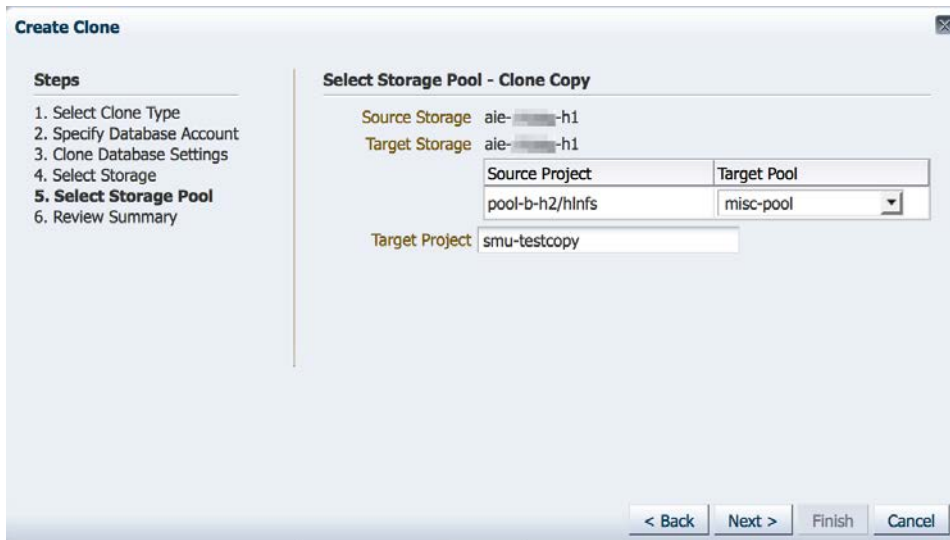


Figure 21. Selecting the target storage pool for a clone copy operation

Once entries are set, click Next to save them and proceed to the final configuration step, the Review Summary.

8. The Create Clone Review Summary window appears with details on the clone for verification. Selecting Finish starts the clone process.

In the Task Panel at the bottom of the screen, you will see the task that has been created for the cloning operation. If the cloning succeeds, the BUI displays the account name of the cloned application (database) in the navigation tree under the Applications node, and in the table under the Snap Backup tab. The corresponding task in the Task panel displays a green check in the Status column. If the cloning task fails, the Status column displays a red cross-check (x). Clicking on the status indicator icon produces a detailed Task Output display that, in the case of a failure, indicates the reason for the failure.

Rolling Back (Restoring) a Snap Backup

SMU can be used to restore a database from an on-disk backup. The restore is accomplished using ZFS rollback technology. Rollbacks allow you to revert a dataset back to a point in time without having to copy or delete any data.

You can only restore a database to the specified backup if there are no database clones made from any newer backups of that database. Since backups are based on ZFS snapshots, restoring to a snap backup automatically destroys any newer snap backups.

The following figure shows a selected Restore operation. Choosing the reverse arrow icon, listed horizontally next to the Actions and View menus, will also initiate the operation.

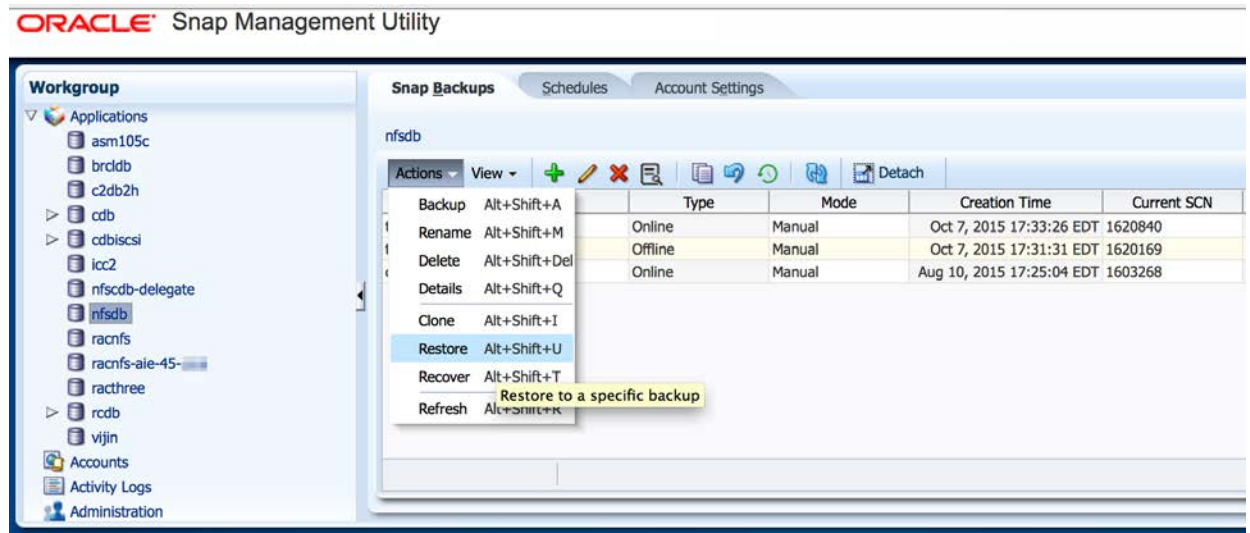


Figure 22. Choosing Restore to roll back to the previous selected database snap backup

Important: This step will revert all the active data, and any more recent snap backups or clones will be deleted. This operation cannot be undone.

1. From the list of applications, select and highlight the application in which you want to perform a rollback.
2. A list of existing snap backups for the selected application will be displayed. Select the snap backup you wish to restore (roll back to) by clicking on it.
3. Click on Actions and Restore.
4. A warning reminder that all existing snap backups and clones that occurred after the selected restore target will be deleted pops up. To execute the restore, select OK.

Recovering a Snap Backup (Recovering a Database) from a Point in Time

SMU can recover a database to a specified point in time, between snap backups. The point in time can be designated either by a time, a change number, or a sequence number. Similarly to the restore operation that is accomplished using ZFS rollback technology, recovery allows you to revert a dataset back to a point in time without having to copy or delete data. SMU first rolls back to a snap backup previous to the designated point in time for recovery, then uses archive logs to fill in the blanks between that previous snap backup and the user's selected point in time. The point in time for recovery can be designated by time, by change, or by sequence.

You can only recover a database to the specified point in time if there are no database clones made from any newer backups of that database. Since backups are based on ZFS snapshots, restoring to a snap backup automatically destroys any newer snap backups.

The following figure shows the selection of a Recover operation. Choosing the reverse arrow icon, listed horizontally next to the Actions and View menus, will also initiate the operation.

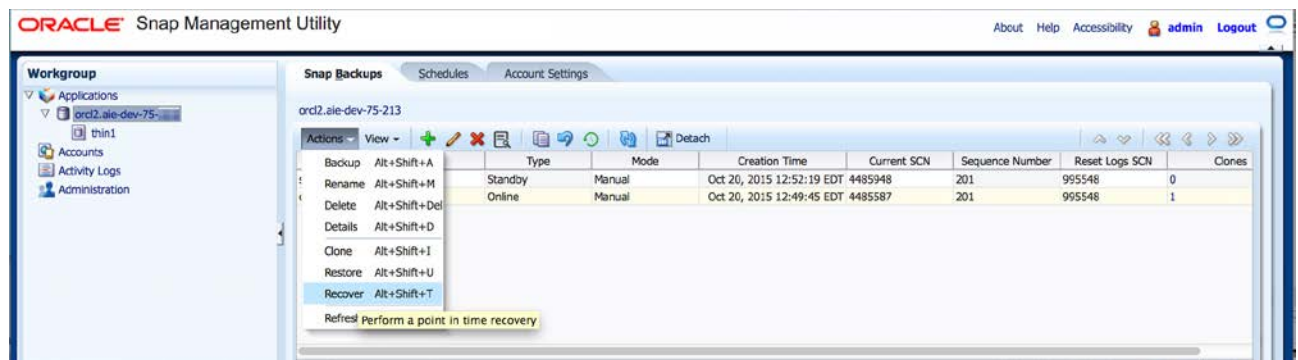


Figure 23. Choosing Recover to roll back a database snap backup to a specified point in time

Important: This step will revert all the active data, and any more recent snap backups or clones will be deleted. This operation cannot be undone.

1. From the list of applications, select and highlight the application in which you want to perform a recovery.
2. A list of existing snap backups for the selected application will be displayed. Click on Actions and Recover.
3. The recover wizard dialog window appears, with the following Recovery Point options:
 - Date – Selecting this date- and time-based option generates a calendar popup for time designation.
 - SCN – Recovers up to the specified database system change number (SCN).
 - Sequence – Recovers up to the specified log sequence number.
 - Recover to an Ancestor Incarnation – Checking this option activates the Reset Logs Change Number input box for specifying the ancestor incarnation number to recover to.
 - Reset Logs Change Number – Recovers to ancestor incarnation number.

Choose one of these Recovery Point options and select Next at the bottom right of the dialog window. SMU will validate the chosen selection.

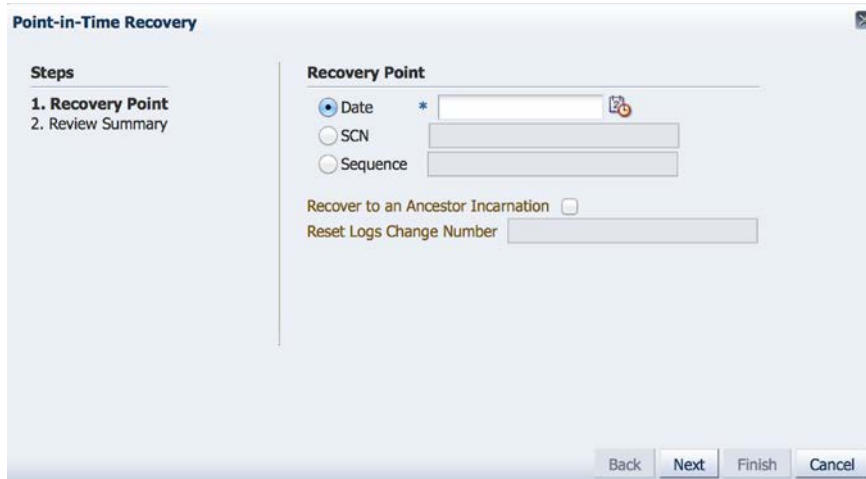


Figure 24. Point-in-Time Recovery option screen

Choose one of these search parameters and select Next at the bottom right of the dialog window. SMU will validate the chosen search parameter.

- If the chosen option cannot provide the desired recovery, the Review Summary window issues a warning notice to go back and select another option. If a suitable backup is determined by SMU based on the recovery point search parameter, the Review Summary window will appear with details for the recovery, similar to the following figure.

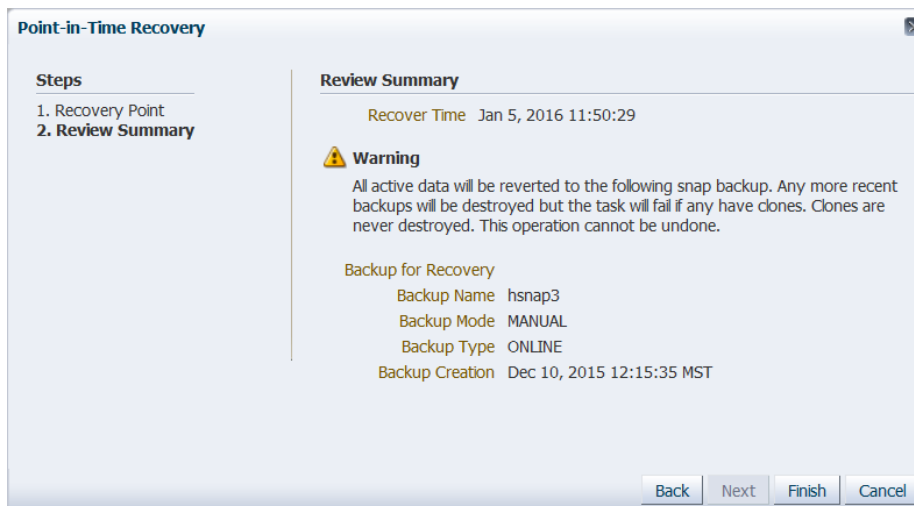


Figure 25. Recover Review Summary of chosen parameter for valid snap backup recovery

- Select Finish to execute the recovery operation, or Back to choose a different option.

Refreshing a Clone (Updating to Current Source Database)

The Refresh Clone operation can update the data in an existing clone to match the current source database contents. This operation removes the need for the more complicated process of deprovisioning an existing clone and then creating a new one in order to ensure an up-to-date clone for dev/test operations. Refresh Clone operations are only permitted on thin clones.

Use the following steps to perform a Refresh Clone operation:

1. In the Applications window, select the clone you wish to update to match the source database.
2. Select Actions and Refresh Clone from the Actions pull-down menu, as seen in the following figure.

IMPORTANT: If you select a database application in the navigation tree under Workgroup>Applications, the Snap Backups tab and screen will be displayed. From there, if you select the Actions pull-down menu, there is a Refresh option for updating/refreshing the screen display that is NOT related to the Refresh Clone option. Be sure to stay in the Applications window as seen in the following figure to correctly select the Refresh Clone operation.

ORACLE Snap Management Utility

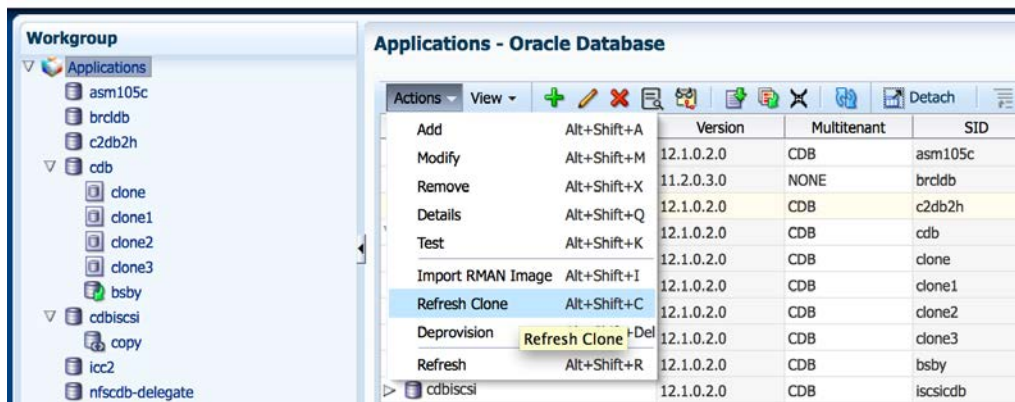


Figure 26. Choosing the Refresh Clone action in the Applications screen

Deleting a Snap Backup

Important: You can only delete a backup if it does not have any dependent clones. When performing a delete task, SMU checks the backup to see if any clones were made from it and the clones are still active. If there are active clones, SMU will fail the task and not allow the backup to be deleted.

1. From the list of the snap backups that are displayed, select and highlight the row that you want to delete.

2. Click on Actions and Delete.
3. A warning message indicating all data within the snap backup will be lost, and the operation cannot be undone, appears. Click on OK to confirm the delete operation on this row.

Running Simultaneous Snap Operations

SMU allows snaps operations on different applications to run in parallel, so that you can take snap backups on multiple applications or create more than one clone at the same time. SMU performs snap operations on the same application sequentially in order to preserve the dependencies among snap operations and resources. Monitor snap operation tasks currently running or waiting to execute in the Task Monitoring panel. You can configure the length of the visible task queue by selecting the General Settings tab of the Administration window.

This task queue setting is also configurable in the `$SMU_HOME/etc/smu.conf` file by setting the parameter `task.queue.length`. The parameter's default is 100. Restart SMU in order for the parameter setting change to take effect.

Note that SMU will block a submitted task if it has any dependency on a currently running task. A dependency exists if applications share the same project or if a snap operation that modifies the application shares is currently running. As described in the section "Managing Activity Logs Using the BUI" that follows, The Task Output Level pull-down menu provides level-of-detail display options – Summary, Standard, and Detailed – that can be very useful for task monitoring.

You can change the task output level on either the Task Monitoring panel or the Activity Logs table, and these two BUI displays' settings are independent of each other. So, for example, changing the task output level to Detailed on the Task Monitoring panel leaves the setting for the Activity Log table unchanged.

Managing Snap Backup Initiation in the Schedules Tab

The following figure shows the BUI displayed by selecting the Schedules tab.




Figure 27. Viewing a selected application's scheduled backups

The displayed table shows a list of snap backup schedules for the selected application and includes the following items:

Frequency – Hour, Day, Week, Month.

Day – Day of the week.

Hour – The hour the snap backup will be taken.

Minute – The minute the snap backup will be taken.

Retention – Number of snap backups that will be retained.

Backup Type – Offline or Online.

Description – Brief description of this entry entered by the user.

Adding a schedule

1. Click on Actions and Add.
2. Enter the values for Frequency, Day of Month, Day of Week, Hour, Minute, Retention, Backup Type and Description.
3. Click OK.

This entry will show up in the table under the Schedules tab for this application. The snap backups will be taken for this application according to the schedule just created.

Editing a schedule

1. Select an application of your choice and an entry from the table under the Schedules tab that you want to edit.
2. Click on Actions and Modify.
3. Edit the values for Frequency, Day of Month, Day of Week, Hour, Minute, Retention, Backup Type and Description.
4. Click OK.

The edited schedule will be reflected in the table under the Schedules tab for the selected Application.

IMPORTANT: If a backup schedule is modified while its auto backup is currently in process, this backup will not be managed by the new retention policy for the edited schedule. The backup should be manually deleted when the backup is no longer used.

Deleting a schedule

1. Select an application of your choice and an entry you want to delete from the table under the Schedules tab.
2. Click on Actions and Remove.
3. Click OK to confirm the delete operation. This will delete the selected schedule you created for taking snap backups for the selected application.

Monitoring Snap Operation Tasks

The Task Monitoring Table lists tasks currently running, completed, or waiting in the queue to be processed. The Task Monitoring Table displays ID, Application, Task (snap operation), Description, Start Time, End Time, and Status. The Application column displays the name of the snap operation target application in a hypertext link. Clicking on the application name, you can navigate to the Snap Backups Summary page of the application.

The Status column displays the status of the task with a status icon. Clicking on the status icon launches the task output window where you can view the detailed task output. By selecting a particular status icon among the set of status icons on the right upper corner of the Task Monitoring Table panel, you can list tasks only in that status.

Canceling currently running tasks

To cancel a currently running task, select the task and click on the cancel button. It is very important to be aware that canceling a task may leave the underlying application or resources in an inconsistent state. Therefore, cancel a task with caution.

Deleting task history

Delete any completed tasks of no further interest for review from the Task Monitoring Table by selecting the task and clicking on the delete button. If you need to track back to the deleted task, you can still view its task output on the Activity Logs page.

Managing Account Settings Using the Account Settings Tab

The following figure shows the basic account information displayed using the Account Settings tab for a selected application. In this tab, you can make modifications to the Application, Application Host, and Storage settings by selecting their corresponding right-sided change buttons (seen as Modify to the right of the Application details screen in the figure) and filling out the appropriate editable fields. The display for each of the three settings windows operate in windowshade fashion so that the selected setting fills as much of the display screen as needed and available.

If a selected application account has backups under it, neither the Application Host nor Storage settings can be modified; the Change Host/Change Storage buttons will not display.

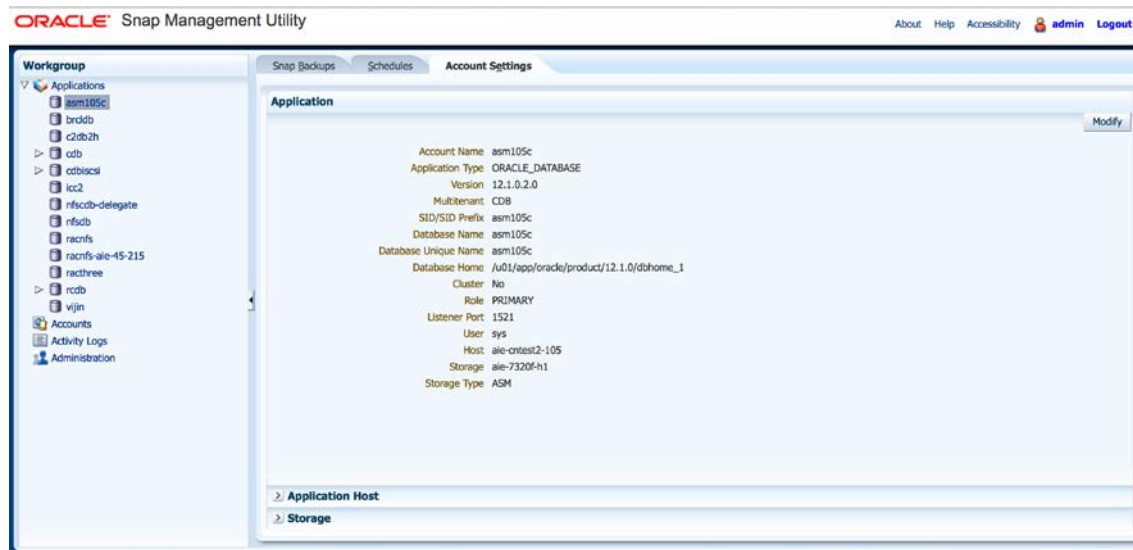


Figure 28. Displaying account configuration in the Account Settings tab

Managing Activity Logs Using the BUI

Located below the Accounts category in the left-side Workgroup area of the BUI, the Activity Logs category contains information on the past user-initiated activity that has taken place with the Snap Management Utility. The Activity Logs display parameters and navigation are highly customizable for the user. The Task Output Level pull-down menu provides these level-of-detail display options: Summary, Standard, and Detailed.

The Summary level presents only higher-level task actions. The Standard level presents task actions and action-level information or messages resulting from a task action. The Detailed level provides task actions and operational details for each action within that task.

Standard level view is set by default. You can change the task output level on either the Task Monitoring panel or the Activity Logs table, and these two BUI displays' settings are independent of each other. So, for example, changing the task output level to Detailed on the Task Monitoring panel leaves the setting for the Activity Log table unchanged.

The Actions and View bar also include the following separate action icons, whose functions can also be selected in either the Actions or View pull-down menus. Mousing over them displays their name/function:

- Filter Activity Logs – Filter by Log ID, User, Application, Action, Task or Status.

- Export Activity Logs – Specify beginning and ending log ID numbers for the range to export. The exported range of logs will appear in the Tasks window, and can be detached as a separate table display window for easier inspection as well as maintaining a record of them before purging logs.
- Purge Activity Logs – Delete a selected range, specified by beginning and ending log IDs, from the Activity Logs. As noted in Export Activity Logs, use that function to save a record of logs before purging them.

Clicking on “Activity Logs” generates a display with the following column items:

- ID – The system-generated unique identification number for each of the activities.
- User – The SMU user name for who initiated the particular activity.
- Application – The name of the application for which the activity was initiated.
- Action – The action that was initiated by the user.
- Task – The numeric ID assigned for a task generated from a user-initiated action.
- Description – A brief description of the activity.
- Time – The time the activity was initiated.
- Status – Indicates whether the activity was successfully completed or not.

These column displays can be individually hidden or selected for display using the View pull-down menu, Columns setting.

The following figure shows the BUI screen for Activity Logs.

The screenshot shows the Oracle Snap Management Utility (SMU) BUI interface. The top navigation bar includes the Oracle logo, the text "Snap Management Utility", and links for "About", "Help", "Accessibility", "admin", and "Logout".

The main interface is divided into two primary sections: "Activity Logs" and "Tasks".

Activity Logs Section:

- Left sidebar: "Workgroup" tree showing "Applications" (orc2.aie-dev-75-213), "Accounts", "Activity Logs", and "Administration".
- Table: "Activity Logs" with columns: ID, User, Application, Action, Task, Description, Time, Status.

ID	User	Application	Action	Task	Description	Time	Status
73	admin		Login	user	user admin logged in	Dec 17, 2015 21:52:36 EST	SUCCEEDED
72	admin		Login	user	user admin logged in	Dec 16, 2015 18:15:50 EST	SUCCEEDED
71	admin		Login	user	user admin logged in	Dec 16, 2015 16:56:33 EST	SUCCEEDED
70	admin		Login	user	user admin logged in	Dec 16, 2015 10:13:09 EST	SUCCEEDED
69	admin		Login	user	user admin logged in	Dec 15, 2015 18:28:35 EST	SUCCEEDED
68	admin		Login	user	user admin logged in	Dec 15, 2015 18:18:27 EST	SUCCEEDED
67	admin		Discover Resources	discover	discover enrolled resources	Dec 15, 2015 18:17:28 EST	SUCCEEDED
66	admin		Modify Application Account	modify	modify application account thin1	Dec 15, 2015 18:17:28 EST	SUCCEEDED
65	admin		Modify Application Account	modify	modify application account orc2.aie-dev-75-213	Dec 15, 2015 18:17:27 EST	SUCCEEDED
64	admin		Modify Storage Account	modify	modify storage account aie-7330a-h1	Dec 15, 2015 18:17:25 EST	SUCCEEDED
63	admin		Modify Host Account	modify	modify host account aie-dev-75-213	Dec 15, 2015 18:17:12 EST	SUCCEEDED
62	admin		Discover Resources	discover	discover enrolled resources	Dec 15, 2015 18:17:11 EST	SUBMITTED

Tasks Section:

- Table: "Tasks" with columns: ID, Application, Task, Description, Start Time, End Time, Status.

ID	Application	Task	Description	Start Time	End Time	Status
4		Discover	discover enrolled resources	Dec 15, 2015 18:17:11 EST	Dec 15, 2015 18:17:28 EST	Success
3	thin1	Clone Backup	clone orc2.aie-dev-75-213:onlinesnap to thin1	Oct 20, 2015 12:54:17 EDT	Oct 20, 2015 12:57:26 EDT	Success
2	orc2.aie-dev-75-213	Backup	backup orc2.aie-dev-75-213 as standbysnap	Oct 20, 2015 12:51:01 EDT	Oct 20, 2015 12:52:19 EDT	Success
1	orc2.aie-dev-75-213	Backup	backup orc2.aie-dev-75-213 as onlinesnap	Oct 20, 2015 12:47:42 EDT	Oct 20, 2015 12:49:45 EDT	Success

Figure 29. Displaying the activity logs in the SMU BUI

When a user action is performed as a task, the activity log(s) for the user action will display a task ID in the Task field. The task ID is a hyperlink that links to the corresponding line in the Tasks panel below the Activity Logs window.

You can view completed task output by clicking the hyperlink in the Task field in the Activity Logs window or, by clicking the associated Output field icon for the task in the Tasks panel, you can view not only completed task output but also running task output in real-time.

On the right side of the display, scroll arrows for up, down, first, last, previous and next task items are selectable. The Jump To tab, when selected, presents a dialog box to choose a task ID for display.

A user action performed as a task will have two activity log entries: one for task submission and another for task completion (status could be succeeded, failed, or canceled, for example). Neither intermediary (sub)actions or task deletions will record activity logs.

Filtering Activity Logs

The Activity Logs window contains configurable filters that, once selected, are persistent per that user's setup. The optional fields for use as filters include: Log ID, User, Application, Action, Task, and Status.

The filters for Log ID and Task use the form of an ID range that can be expressed with a comparison operator followed by a number (ID). The supported comparison operators are listed in the following table.

If you are viewing the Activity Logs window and do not see the expected rows in the activity table, check to see if any filter criteria are specified. Change or reset the criteria to view the activity logs you are interested in.

TABLE 9. ACTIVITY LOG FILTER SPECIFICATIONS		
OPERATOR	SEARCH PARAMETER	EXAMPLE
> <i>n</i>	Search for logs whose ID is greater than <i>n</i>	>10 returns logs whose ID is greater than 10
>= <i>n</i>	Search for logs whose ID is greater than or equal to <i>n</i>	>=20 returns logs whose ID is greater than or equal to 20
= <i>n</i>	Search for logs whose ID equals <i>n</i>	=30 returns logs for ID number 30
<= <i>n</i>	Search for logs whose ID is less than or equal to <i>n</i>	<=40 returns logs whose ID is less than or equal to 40
< <i>n</i>	Search for logs whose ID is less than <i>n</i>	<50 returns logs whose ID is less than 50
<i>n</i> - <i>nn</i>	Search for logs whose ID is between <i>n</i> and <i>nn</i>	60-70 returns logs whose ID is between 60 and 70

The filter for the rest of the fields is the text search filter. This filter allows the wildcard character %. Example uses are:

`%xyz` – Search for the text fields whose value ends with `xyz`.

`xyz%` – Search for the text fields whose value starts with `xyz`.

`%xyz%` – Search for the text fields whose value contains `xyz`.

When no wildcard characters are used, the filter finds the exact match. The text comparison of the text search filter is case sensitive.

By default, activity logs are fetched up to 10,000 entries. The max rows to fetch can also be configured using the `actlog.max_fetch` parameter of the `smu.conf` file. Note that this parameter setting will apply across the particular SMU application, so it will apply to all SMU users. If the `smu.conf` file has been changed, the Snap Management Utility must be stopped and restarted in order to activate the new `smu.conf` file parameter settings.

Exporting Activity Logs

To export a range of activity log entries:

1. Either select Actions and Export in the Actions pull-down menu, or select the Export Activity Log icon next to the Filter (funnel-shaped) icon.
2. Provide values for the task ID range you would like to export, in Begin ID and End ID.
3. Select OK.
4. The Export Activity Logs operation will appear in the Tasks window, as shown in the following figure. Click on the green Success checkmark icon in the Status column of the newly executed Export Activity Logs task to

see the details of this operation, including the name of the zip file that the operation creates to hold the activity logs for the specified task IDs.

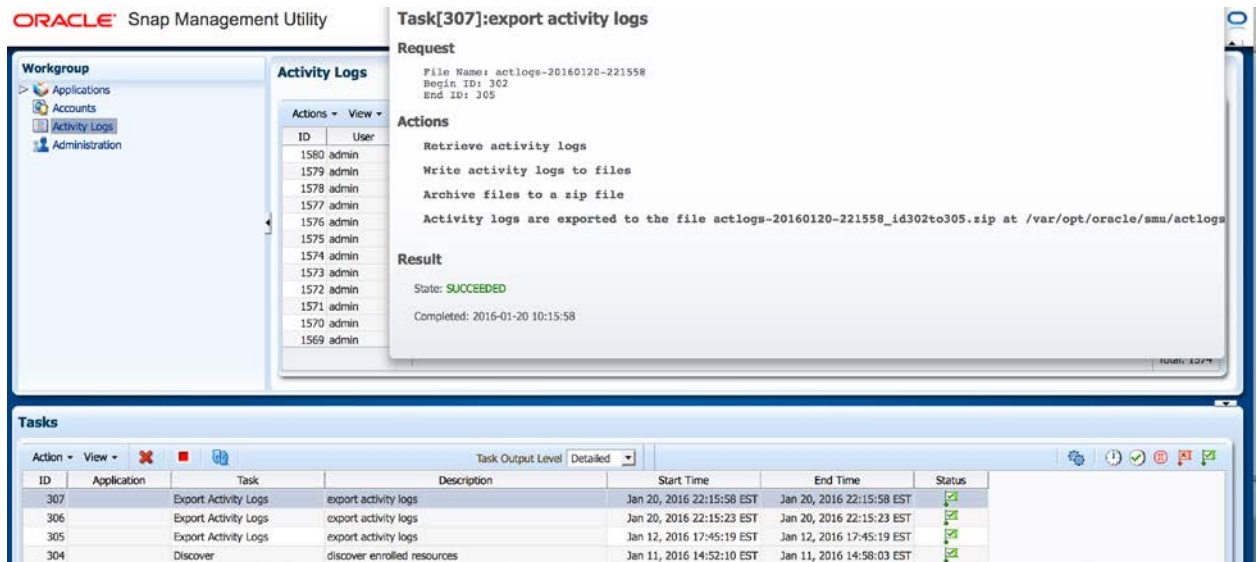


Figure 30. Task Output window showing zip file of exported activity logs

Purging Activity Logs

To purge a selected range of tasks in the Activity Logs window:

1. Either select the Purge Activity Logs icon (crossbar) or select Actions, Purge. A popup window provides a warning reminder that, upon the purge operation's completion, the selected task IDs will be permanently deleted and should first be backed up to a file using the Export operation.
2. Provide entries for the Begin ID and End ID task numbers which indicate the range of task IDs to purge.
3. Verify that you are committing to deletion of these entries by typing "Yes" in the Confirm box.
4. Select OK.

Using the Command-Line Interface

Once you have accessed the CLI you are presented with the SMU shell. The prompt `smu>` is displayed. You can type SMU commands at the SMU prompt. The following command categories are available:

TABLE 10. AVAILABLE COMMAND SETS IN SNAP MANAGEMENT UTILITY

COMMAND CATEGORY	COMMAND	DESCRIPTION
Accounts	<code>accounts</code>	Accounts used to access remote resources such as databases, database hosts, or RAC cluster nodes and Oracle ZFS Storage Appliances.
Activities	<code>activities</code>	A record of all actions performed by SMU users that can be used for auditing purposes.
Alerts	<code>alerts</code>	Actions SMU will take when events such as task failures or task cancellations occur.
Backups	<code>backups</code>	Backups created by the user manually or automatically by a schedule.
Certificates	<code>certs</code>	Certificates used by the embedded SMU WinRS and web servers. For the v1.0 release, only the self-signed SSL certificate is managed.
Keys	<code>keys</code>	SSH public keys of SMU users. This allows you to log in to SMU using SSH with key-based authentication instead of password-based authentication.
Schedules	<code>schedules</code>	Schedule automatic backups with optional retention policies.
Tasks	<code>tasks</code>	Background tasks. SMU performs all snapshot-based operations as background tasks since they involve logging into and coordinating remote resources and executing a series of commands that can take a long time to complete.
Users	<code>users</code>	SMU users. Both local and directory users are supported.

Each command category supports a set of subcommands to manage the category objects. Most categories follow the `add/modify/remove` idiom, which allows you to, respectively, add new objects, modify existing objects and remove objects. All categories support the `get` and `list` methods. These methods are used to display the objects in property list or tabular format.

The following example shows the syntax of each subcommand supported in each command category.

To create a new object:

```
add [category options] [-o property=value] ... name
```

where:

category options – Specifies one or more category options (depending on the command category). Additionally, you can specify general properties of the object.

-o *property=value* – Sets the specified property. Multiple -o options can be specified.

name – Creates the new object with the specified name.

To display properties for a given object:

```
get [-H] [-o all | field[,...]] all | property[,...] [name] ...
```

where:

-H – Displays output in a form more easily parsed. Headers are omitted and fields/columns are separated by a single tab instead of an arbitrary amount of space.

-o *field* – Sets the fields to display, which includes one or more of name, property, or value. Present multiple fields as a comma-separated list. The default value is: *name,property,value*. The special value *all* will display all properties.

If no object names are specified, then the command displays properties for all objects in that command category. For each property, these columns are displayed:

<i>name</i>	Object name or identifier
<i>property</i>	Property name
<i>value</i>	Property value

To list property information for the given datasets in table form:

```
list [-H] [-o property[,...]] [ -s property ] ... [ -S property ] ... [name] ...
```

where:

-H – Suppresses printing of headers. Fields are separated by a single tab instead of arbitrary white space.

-o *property* – A comma-separated list of properties to display.

- s *property* – A property for sorting the output by column in ascending order based on the value of the property. Multiple properties can be specified at one time using multiple -s property options. Multiple -s options are evaluated from left to right and establish the sorting precedence.
- S *property* – Same as the -s option but sorts by property in descending order.

To modify one or more properties of an object:

```
modify [-o property=value] ... name
```

where:

- o *property=value* – Sets the specified property. Multiple -o options can be specified.
- name* – Name of the object whose property will be modified.

To remove an object:

```
remove [-F] name
```

where:

- name* – Lists the object to be deleted from the command category.
- F – Forcibly removes the object. Without this option, SMU will prompt the user to confirm before removing the object.

Managing Accounts

SMU accesses and coordinates various application and system resources while performing operations. In order to access these resources, the user must supply the accounts for SMU to use. Each account has a type and protocol. The type identifies what type of resource the account is for: application, host, or storage. The protocol identifies what method is used to access the resource. Each protocol has a set of properties that must be specified. Some properties have defaults; others require setting their values. The following table lists the accounts supported by SMU along with their type and properties:

TABLE 11. SUPPORTED ACCOUNTS ACCESSED BY SNAP MANAGEMENT UTILITY

ACCOUNT TYPE	ACCOUNT PROTOCOL	PROPERTIES (DEFAULT VALUE)
APPLICATION	ORACLE_DATABASE (Oracle Database)	db_configuration (SI) host db_unique_name (orcl) oracle_sid (orcl) password port (1521) storage
HOST	SSH2 (Secure Shell Version 2)	hostname password port (22) user (root) delegate (none)
HOST	WINRS (Window Remote Shell)	hostname password port (5986) user (Administrator)
STORAGE	SUN_ZFS_STORAGE (Oracle ZFS Storage Appliance)	hostname password port (22) user (root)

At a minimum, you must add one account of each type in order to perform any SMU tasks. Additionally, the application account must reference a host and storage account through the properties, which is how you associate or link the application to its hardware resources; namely, the host the application is running on and the storage the application is using.

The following table lists actions that can be performed on accounts. Command examples follow.

TABLE 12. PERMISSIBLE ACTIONS ON ACCOUNTS

ACCOUNTS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add an account	<code>accounts add [-t <i>type</i>] [-p <i>protocol</i>] [-o <i>option</i>] ... <i>name</i></code>
get	Get account properties	<code>accounts get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>name</i>] ...</code>
list	List accounts	<code>accounts list [-t <i>type</i>] [-a -c] [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [<i>name</i>] ...</code>
modify	Modify an account	<code>accounts modify [-o <i>option</i>] ... <i>name</i></code>
remove	Remove an account	<code>accounts remove [-F] <i>name</i></code>
test	Test an account	<code>accounts test <i>name</i></code>

To set or modify the account password property

Each account has a password property. This property value can be supplied on the SMU accounts command line or can be entered interactively with no character echoing. SMU will prompt for the password if it is not specified on the accounts command line in the following fashion when adding a new account:

```
Type password:
Re-type password:
```

To modify an account password interactively, clear the password on the accounts modify command line:

```
smu> accounts modify -o password= <account name>
Type password:
Re-type password:
smu>
```

To add an Oracle ZFS Storage Appliance account

```
smu> accounts add -t STORAGE -p SUN_ZFS_STORAGE -o hostname=<hostname> -o user=<user>
-o port=<port> -o password=<password> <account name>
```


To add a UNIX host account

```
smu> accounts add -t HOST -p SSH2 -o hostname=<hostname> -o user=<user> -o port=<port>
-o password=<password> -o delegate=<NONE/SUDO> <account name>
```

To add a Windows host account

```
smu> accounts add -t HOST -p WINRS -o hostname=<hostname> -o user=<user> -o
port=<port> -o password=<password> <account name>
```

To add an Oracle Database account

```
smu> accounts add -t APPLICATION -p ORACLE_DATABASE -o
db_configuration=<SI/RAC/RACOneNode> -o host=<host account> -o oracle_sid=<SID> -o
db_unique_name=<DB Unique Name> -o password=<password> -o port=<port> -o
storage=<storage account> <account name>
```

To get all of the properties of every account

```
smu> accounts get all
```

To list accounts

```
smu> accounts list
```

To list accounts of a specified type*

```
smu> accounts list -t type
```

* Values for *type* can be application, host, or storage (Both uppercase and lowercase letters are accepted for these entries.)

To list application accounts with their associated details*

```
smu> accounts list -a
```

*including their associated properties, host and storage accounts, origin of the clone and clone method

To list clones of the specified applications

```
smu> accounts list -c [app]
```

To modify an account

```
smu> accounts modify -o property=value ... <account name>
```

To remove an account

```
smu> accounts remove <account name>
```

You cannot remove application accounts that have backups. All backups of the application must be deleted before you can remove the application account.

You cannot remove host accounts that are referenced from an application account. All application accounts that reference the host account must be removed first.

You cannot remove storage accounts that are referenced from an application account. All application accounts that reference the storage account must be removed first.

Managing Activities

SMU keeps a record of all commands run by the user through the CLI or BUI using an activity log. Each activity record has the following fields:

TABLE 13. ACTIVITY RECORD FIELDS

ACTIVITIES FIELD	DESCRIPTION
ID	Activity integer identifier
time	The time of the activity
status	The status of the action (succeeded, failed, submitted, cancelled)
user	The user performing the activity
action	The action performed by the user

The activity log will grow over time as the software is used. Trim the activity log by removing older records. The following table lists permitted actions that can be performed on activities.

TABLE 14. PERMITTED ACTIONS ON ACTIVITIES

ACTIVITIES SUBCOMMANDS	DESCRIPTION	SYNOPSIS
get	Get activity properties.	activities get [-H] [-o "all" field[,...]] <"all" property[,...]> [id] ...
list	List activity.	activities list [-H] [-o property[,...]] [-s property] ... [-S property] ... [id] ...
purge	Purge activities.	activities purge <id1-id2 id>

To get all of the properties of activities

```
smu> activities get all
```

To list activities

```
smu> activities list
```

To purge a range of activities

```
smu> activities purge <id1-id2>
```

Managing Alerts

The software can be configured to alert the user when certain events occur. The following event alerts are supported:

TABLE 15. CONFIGURABLE EVENT ALERTS

EVENT	DESCRIPTION
TASK CANCELLED	A task has been canceled.
TASK FAILED	A task has failed.

The following actions can be performed when an event occurs:

TABLE 16. PERMISSIBLE ACTIONS FOR EVENT OCCURRENCES

ACTION	DESCRIPTION	PROPERTIES
EMAIL	Send an email message.	address subject

The EMAIL action requires that a list of email addresses and a subject line be specified for the mail message. You can specify multiple recipients for the address property by listing multiple email addresses separated by a comma.

The following actions can be performed on alerts.

TABLE 17. PERMISSIBLE ALERT ACTIONS

ALERTS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add an alert.	<code>alerts add [-o <i>option</i>] ... <i>event</i> <i>action</i></code>
get	Get alert properties.	<code>alerts get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>event:action</i>] ...</code>
list	List alerts.	<code>alerts list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [<i>event:action</i>] ...</code>
modify	Modify an alert.	<code>alerts modify [-o <i>option</i>] ... <i>event</i> <i>action</i></code>
remove	Remove an alert.	<code>alerts remove [-F] <i>event action</i></code>

To send an email alert when a task fails

```
smu> alerts add -o address=<address-list> -o "subject=<subject>" TASK_FAILED EMAIL
```

To get all of the properties of every alert

```
smu> alerts get all
```

To list alerts

```
smu> alerts list
```

To modify an alert

```
smu> alerts modify -o property=value ... <event> <action>
```

To remove an alert

```
smu> alerts remove <event> <action>
```

Managing Backups

The software allows you to create on-disk backups of databases that use the Oracle ZFS Storage Appliance for storage. These backups are based on creating ZFS snapshots of the database shares, so they are only suitable for a specific class of use cases in which the user wants to make a quick backup for development or testing purposes. The backups can be used to restore, recover, or clone the database. Backups are added and removed by running the appropriate task (see Managing Tasks for more information).

The following actions can be performed on backups:

TABLE 18. PERMISSIBLE BACKUP MANAGEMENT ACTIONS

BACKUPS MANAGEMENT SUBCOMMANDS	DESCRIPTION	SYNOPSIS
get	Get backup properties.	<code>backups get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>app:name</i>] ...</code>
list	List backups.	<code>backups list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [-t <i>app</i>[,...]] [<i>app:name</i>] ...</code>

To get all of the properties of every backup

```
smu> backups get all
```

To list backups

```
smu> backups list
```

Managing Certificates

The software uses a single self-signed certificate for encrypting Windows Remote Shell (WinRS) client and web browser sessions. The certificate is generated the first time SMU is started and is saved to a keystore file for subsequent use.

The following actions can be performed on certificates.

TABLE 19. PERMISSIBLE CERTIFICATE ACTIONS

CERTS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
get	Get certificate properties.	<code>certs get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>alias</i>] ...</code>
list	List certificates.	<code>certs list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [<i>alias</i>] ...</code>

To get all of the properties of every certificate

```
smu> certs get all
```

To list certificates

```
smu> certs list
```

Managing Keys

The software allows you to administer SSH2 public keys that can be used to perform key-based authentication in place of password-based authentication when connecting to SMU using an SSH client. A key has the following properties:

TABLE 20. KEY PROPERTIES

PROPERTY	DESCRIPTION
alias	An alias for the key.
encoded	The SSH public key encoded in ssh-dss or ssh-rsa format. See http://www.ietf.org/rfc/rfc4253.txt for more information.

The following actions can be performed on keys.

TABLE 21. PERMISSIBLE KEYS ACTIONS

KEYS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add a key.	<code>keys add alias encoded</code>
get	Get key properties.	<code>keys get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [alias] ...</code>
list	List keys.	<code>keys list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [alias] ...</code>
remove	Remove a key.	<code>keys remove [-F] alias</code>

To add an ssh-dss key with the alias mykey

```
smu> keys add mykey
AAAAB3NzaC1kc3MAAACBAJM3cknqShlHI8E9EXWXYgM/XeLl+0jccFG3C/W7C7dD6dLxlAOW5Tv67le/ils1N9
be8KEIZDdX85/wnRkyRomhjHMs7TEYDzHRoWS5gzBMr93pkkNiWdo02B9fUo2RpAlimBHQ+G09PQCoLPlSKfYC
lup0UiKg8H3XDCSu+tGBAAAAFQDwhT/KGyLrZCpaYNbieuzfgYoykwAAAIBEC/pEe2k8Gt8IgiYsZj7iw/aHoc
u4/Kri2cqoCDTKHWK08IE+ZHq0tUR6vvfOfdthsCX3Qiqh3ufZYltK0BR5qTS5AWW5cESnIN/orWCxvbrkuLOx
tAaq1GjZaV+cGLmJvWartIYhB68j7NxbsUoyW9yRb0TQCvv+J/rLGWi/GAAAAIEAiu0+IfKWNapbsf09TBbWao
AZavj8c1z7KUFyXmu99fJjKOAVL0K4uVNmsnwy4glMGVmxEQVxcmZh/WPfltvRkr2n5TmlA1DtJkIr3FOJ7XXs
zCt7M1q6JHEy/oHqjWnkZTYhp8LOGigZBh9mfU4B5id+TZQRtdm2ggRU82H6JEg=
```

To get all of the properties of every key

```
smu> keys get all
```

To list keys

```
smu> keys list
```

To remove a key

```
smu> keys remove <alias>
```

Managing Schedules

The software allows you to schedule automatic backups at regular intervals such as hourly, daily, weekly or monthly. Additionally, each schedule can retain a set number of backups. For example, you can schedule automatic backups to occur daily and retain the seven (7) most recent backups.

The following schedule frequencies are available along with the day, hour and minute ranges.

TABLE 22. AUTOMATIC BACKUP SCHEDULE PARAMETERS

SCHEDULE FREQUENCIES	DAY RANGE	HOUR RANGE	MINUTE RANGE
hour	-1	-1	0-59
day	-1	0-23 (0 = midnight)	0-59
week	1-7 (day of the week, 1=Sun, 7=Sat)	0-23 (0 = midnight)	0-59
month	1-31 (day of the month)	0-23 (0 = midnight)	0-59

The following actions can be performed on schedules.

TABLE 23. PERMISSIBLE ACTIONS AND RELATED SUBCOMMANDS FOR SCHEDULES

SCHEDULES SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add a schedule.	<code>schedules add [-k keep] [-d desc] [-o option] ... app freq day hour minute</code>

TABLE 23. PERMISSIBLE ACTIONS AND RELATED SUBCOMMANDS FOR SCHEDULES

SCHEDULES SUBCOMMANDS	DESCRIPTION	SYNOPSIS
get	Get schedule properties.	<code>schedules get [-H] [-o "all" field[,...]] <"all" property[,...]> [app:freq:day:hour:minute] ...</code>
list	List schedules.	<code>schedules list [-H] [-o property[,...]] [-s property] ... [-S property] ... [app:freq:day:hour:minute] ...</code>
modify	Modify a schedule.	<code>schedules modify [-k keep] [-d desc] [-o option] ... app freq day hour minute</code>
remove	Remove a schedule.	<code>schedules remove [-F] app freq day hour minute</code>
rename	Modify frequency, day, hour, minute of a schedule.	<code>schedules rename app:freq:day:hour:minute freq day hour minute</code>

To schedule an online backup at the top of every hour

```
smu> schedules add -o type=online <app> hour -1 -1 0
```

To schedule an offline backup at midnight every day

```
smu> schedules add -o type=offline <app> day -1 0 0
```

To schedule a backup every Friday at 5:00 p.m.

```
smu> schedules add <app> week 6 17 0
```

To schedule a backup on the 15th day of every month at 4:30 a.m.

```
smu> schedules add <app> month 15 4 30
```

To get all of the properties of every schedule

```
smu> schedules get all
```

To list schedules

```
smu> schedules list
```

To modify a schedule

```
smu> schedules modify [-k keep] [-d desc] -o property=value ... <app> <freq> <day>
<hour> <minute>
```

To modify frequency, day, hour, minute of a schedule to the 15th day of every month at 4:30 a.m

```
smu> schedules rename <app>:<freq>:<day>:<hour>:<minute> month 15 4 30
```

IMPORTANT: If a backup schedule is modified while its auto backup is currently in process, this backup will not be managed by the new retention policy for the edited schedule. The backup should be manually deleted when the backup is no longer used.

To remove a schedule

```
smu> schedules remove [-F] <app> <freq> <day> <hour> <minute>
```

Managing Tasks

The software uses tasks to carry out database backup and recovery operations in addition to database cloning. Tasks are designed to run in the background and can take a long time to finish. The task manager is used to submit and run tasks. The following commands can be run as tasks.

TABLE 24. PERMISSIBLE TASK COMMANDS

COMMAND	DESCRIPTION	SYNOPSIS
backup	Back up a database (create a snap backup).	backup [-o type=ONLINE OFFLINE STANDBY] <app> <backup name>
clone	Clone a database from a snap backup.	clone [-o option]<source app> <source backup> <target app>

<pre>clone -o method -o sourceproject -o target pool -o project (for clone copy)</pre>	<p>Create a clone copy specifying through the listed options the method (type) for creating the new copy, the source pool/project shares (for the source database) to be copied, the target pool for this new copy, and the identifying project for the copy.</p>	<pre>clone -o method={copy thin standby} -o sourceproject=[pool/project, [pool/project]...] -o targetpool=[pool[,pool]...] -o project=[project prefix] source_app backup target_app</pre> <p>Note that a blank space exists between -o option subcommand statements.</p>
<pre>deprovision</pre>	<p>Deprovision (delete) a clone database.</p>	<pre>deprovision <app></pre>
<pre>delete</pre>	<p>Delete a snap backup.</p>	<pre>delete <app> <backup name></pre>
<pre>import</pre>	<p>Clone a database from an RMAN image copy backup.</p>	<pre>import [-o option] <mountpoint> <app></pre>
<pre>recover</pre>	<p>Recover a database instance between backups, either by designated SCN change, time, or sequence, and optionally by database incarnation.</p>	<pre>recover -o change=<scn> "time=<yyyy-MM-dd HH:mm:ss>" sequence=<seqno> [-o resetlogs_change=<scn>] <app></pre>
<pre>refresh</pre>	<p>Updates an existing clone's data to the current data of the source database clone. Note that no options need to be specified; all options used to configure the existing clone are reused to configure the refreshed clone.</p>	<pre>refresh <source backup> <target app></pre>
<pre>rename</pre>	<p>Rename an existing backup.</p>	<pre>rename <app> <old backup name> <new backup name></pre>
<pre>restore</pre>	<p>Restore from a snap backup.</p>	<pre>restore [-F] <app> <backup name></pre>

The following table provides further details for command options.

TABLE 25. AVAILABLE TASK COMMAND OPTIONS

COMMAND	OPTION	DESCRIPTION
backup	type	The type of snap backup to create. Options are OFFLINE, ONLINE or STANDBY. OFFLINE is the default when this option is not specified.

TABLE 25. AVAILABLE TASK COMMAND OPTIONS

COMMAND	OPTION	DESCRIPTION
clone	oracle_home	The Oracle Home, or Oracle home directory, the clone database should use. The default is to use the same Oracle Home of the database whose backup is being cloned. The Oracle Home (database version) of a clone operation's target can only vary from the original Oracle Home by the version's fifth digit.
clone	db_name	The database name the clone database should use. The default is to set the database name to the oracle_sid of the clone database account.
clone	type	The type of clone to create. Currently, the only supported type is PRIMARY.
clone	method	The method to use to create a clone copy. The method options are THIN, COPY, and STANDBY. THIN is the default when this option is not specified.
clone	sourceproject	Specifies the source storage projects holding the storage shares of the source database. SMU will use a default value if this is not specified.
clone	targetpool	Specifies the target storage pool to copy the clone shares to. SMU will use a default value if this is not specified. The sourceproject and targetpool should be mapped in a 1-to-1 relationship.
clone	project	The name of the project where the clone shares are to be located. The project prefix for replicated projects defaults to <i>smu-<code><db_name></code></i> . SMU will use a default value if this is not specified.
clone	open_mode	Specifies the desired status of the cloned database after a clone task is completed. Valid values are MOUNTED, READ_ONLY, or READ_WRITE. All clone creation methods can specify this option. For standby clone, only MOUNTED is a valid value.

TABLE 25. AVAILABLE TASK COMMAND OPTIONS

COMMAND	OPTION	DESCRIPTION
clone	sga_target	Specifies the total size of all SGA components of the clone database. The unit specification can be G, M, or K. If left empty, the default is byte. All clone creation methods can specify this option.
clone	log_mode	Specifies the log_mode of the clone database. Values can be archivelog or noarchivelog. All clone creation methods can specify this option.
clone	protection_mode	For standby clone operation, specifies the protection mode of the data guard configuration. Values can be MAXIMUM_AVAILABILITY, MAXIMUM_PERFORMANCE, or MAXIMUM_PROTECTION.
import	oracle_home	The Oracle Home the clone database should use. There is no default. This option must be specified. The Oracle Home (database version) of a clone operation's target can only vary from the original Oracle Home by the version's fifth digit.
import	db_name	The database name the clone database should use. There is no default. This option must be specified.
recover	change time sequence	The option for choosing the recovery point. This option must be specified. Choices are change, time, or sequence, in the form: change=NNNNN "time=YYYY-MM-DD HH:MM:SS" sequence=NNNNNN Note that quotation marks encapsulating the time option must be included as listed because of empty spaces included in the time option syntax.
recover	resetlogs_change	-o resetlogs_change=<resetlogs_change_#>

The following actions can be performed on tasks.

TABLE 26. PERMISSIBLE TASK ACTIONS

TASKS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add a task.	<code>tasks add [-f] [-F] <i>command</i> [<i>options</i>] <i>arguments</i></code>
cancel	Cancel a task.	<code>tasks cancel [-F] <i>id</i></code>
cat	Display task output. Value options for <code>-L</code> indicate level of detail, with 1=basic, 3=most detailed.	<code>tasks cat [-L {1 2 3}] <i>id</i></code>
get	Get task properties.	<code>tasks get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>id</i>] ...</code>
list	List tasks.	<code>tasks list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [-t <i>type</i>[,...]]... [-T <i>state</i>[,...]] [<i>id</i>] ...</code>
remove	Remove a task (must have finished running).	<code>tasks remove [-F] <i>id</i></code>
tail	Tail task output (watch output as it is written). Value options for <code>-L</code> indicate level of detail, with 1=basic, 3=most detailed.	<code>tasks tail [-c <i>chars</i>][-f] [-n <i>lines</i>] [-L {1 2 3}] <i>id</i></code>
wc	Display task output size in characters and lines.	<code>tasks wc [-c] [-l] <i>id</i></code>

To add a default (offline) backup task

```
smu> tasks add backup <app> <backup name>
```

To add an online backup task

```
smu> tasks add backup -o type=ONLINE <app> <backup name>
```

To add a recover task where the recovery point is based on a selected change number (SCN)

```
smu> tasks add recover -o change=NNNNN <app>
```

To add a recover task where the recovery point is based on a selected time *(Note that quotation marks encapsulating the time option must be included as listed because of empty spaces included in the time option syntax.)

```
smu> tasks add recover -o "time=<yyyy-MM-dd HH:mm:ss>" <app>
```

To add a recover task where the recovery point is based on a selected sequence number

```
smu> tasks add recover -o sequence=<sequence number> <app>
```

To add a recover task for recovery to the current incarnation using the time 10:18 on 7-31-2015

```
smu> tasks add -f -F recover -o "time=2015-07-31 10:18:00" app1
```

To add a recover task for recovery to the current incarnation using SCN

```
smu> tasks add -f -F recover -o change=1040320 app1
```

To add a recover task for recovery to the current incarnation using sequence number

```
smu> tasks add -f -F recover -o sequence=100 app1
```

To add a recover task for recovery to an ancestor incarnation using time

```
smu> tasks add -f -F recover -o "time=2015-10-22 11:36:00" -o resetlogs_change=1029840  
app1
```

To add a recover task for recovery to an ancestor incarnation using SCN

```
smu> tasks add -f -F recover -o change=904000 -o resetlogs_change=1029840 app1
```

To add a recover task for recovery to an ancestor incarnation using sequence number

```
smu> tasks add -f -F recover -o sequence=19 -o resetlogs_change=1029840 app1
```

To add a restore tasks

```
smu> tasks add restore <app> <backup name>
```

To add a refresh clone task where the existing clone data is updated/refreshed to the latest data contents of the source database clone

```
smu> tasks add refresh <backup> <target_app>
```

To add a thin clone task with an SGA size of 2g and open mode read only

```
smu> tasks add clone -o method=thin -o sga_target=2g -o open_mode=READ_ONLY orclspl
backup1 orclspl1
```

To create a clone copy task where the source database and its snapshots reside in pool "pool-a-h1" and project "datapool" and pool "pool-b-h2" and project "fraproject" and the clone is to be created in the target pools "pool-c" and "pool-a-h1" and project "targetproject"

```
smu> tasks add clone -o method=copy -o sourceproject=pool-a- h1/dataproject,pool-b-
h2/fraproject -o targetpool=pool-c,pool-a-h1 -o project=targetproject orclspl backup1
orclspl1
```

To create a standby clone where the source database and its snapshots reside in pool "pool-a-h1" and project "datapool" and pool "pool-b-h2" and project "fraproject" and the clone is to be created in target pool "pool-c" and "pool-a-h1" and project "targetproject" with the data guard configuration protection mode set to maximum performance.

```
smu> tasks add clone -o method=standby -o sourceproject=pool-a- h1/dataproject,pool-b-
h2/fraproject -o targetpool=pool-c,pool-a-h1 -o project=targetproject -o
protection_mode=maximum_performance orclspl backup1 orclspl1
```

To add an import task where the RMAN backup in share <mountpoint> is cloned to app <app>

```
smu> tasks add import <mountpoint> <app> -o oracle_home=<oracle-home> <backup-shares-
mountpoint-list> <app>
```

Where <oracle-home> is the Oracle home to use for the clone database (and must be specified), <backup-shares-mountpoint-list> is a comma-separated list of the backup share mountpoints to create the clone from, and <app> is the database account to use for the clone database.

The mountpoint list can specify share mountpoints from either head/controller of a clustered Oracle ZFS Storage Appliance. The backup can span both heads of a clustered Oracle ZFS Storage Appliance. For more descriptions of the mountpoint property, refer to the List of Shares table in the Oracle ZFS Storage Appliance on-line help, section Shares:Shares.

To get all of the properties of tasks

```
smu> tasks get all
```

To list tasks

```
smu> tasks list
```

To list tasks of a specific type

```
smu> tasks list -t <type[,...]>
```

where the value for type can be backup, restore, clone, import, recover, refresh, discover, delete, deprovision, rename, export-log (case insensitive)

To list tasks of specific state(s)

```
smu> tasks list -T <state[,...]>
```

where the value for state can be SUBMITTED, PENDING, RUNNING, CANCELLED, FAILED, SUCCEEDED (case insensitive)

To display task <id> output

```
smu> tasks cat <id>
```

To watch task <id> output

```
smu> tasks tail -f <id>
```

To remove a completed task

```
smu> tasks remove <id>
```

You cannot remove tasks that have not finished running. You must either wait until the task completes or cancel the task before execution.

Managing Users

In order to access the user interfaces, you must first log in to your user account. These user accounts represent the list of users who can use the software. The software supports two types of users: local and LDAP. Local users are only defined within the context of the software. LDAP users are defined within the enterprise.

The following actions can be performed on users.

TABLE 27. PERMISSIBLE ACCOUNT ACTIONS

USERS SUBCOMMANDS	DESCRIPTION	SYNOPSIS
add	Add a user.	<code>users add [-t <i>type</i>] [-o <i>option</i>] ... <i>name</i></code>
get	Get user properties.	<code>users get [-H] [-o "all" <i>field</i>[,...]] <"all" <i>property</i>[,...]> [<i>name</i>] ...</code>
list	List users.	<code>users list [-H] [-o <i>property</i>[,...]] [-s <i>property</i>] ... [-S <i>property</i>] ... [<i>name</i>] ...</code>
modify	Modify a user.	<code>users modify [-o <i>option</i>] ... <i>name</i></code>
remove	Remove a user.	<code>users remove [-F] <i>name</i></code>

To set or modify the user password property

Local users have a password property. Supply this property value either on the SMU user's command line or enter it interactively with no character echoing. When adding a new account, SMU will prompt for the password if it is not specified on the user's command line:

```
Type password:
Re-type password:
```

To modify a user password interactively, clear the password on the `users modify` command line:

```
smu> users modify -o password= <account name>
Type password:
Re-type password:
smu>
```

To add a local user

```
smu> users add -t LOCAL -o gecos=<user fullname> -o password=<password> <name>
```

To add an LDAP user

```
smu> users add -t LDAP -o server=<directory server> -o directory=<directory>  
<username>
```

To get all of the properties of users

```
smu> users get all
```

To list users

```
smu> users list
```

To modify a user

```
smu> users modify -o property=value ... <user name>
```

To remove a user

```
smu> users remove <user name>
```

You cannot remove the admin user.

Troubleshooting – General Information

If encountering difficulties with operations in the Snap Management Utility, first verify that none of the restrictions listed in the quick reference table that follows have been overlooked.

For a list of known issues and required actions, go to My Oracle Support (MOS) at <http://support.oracle.com>. Do a search for Doc ID 1522925.1. Some of the more common issues have been contained in this document in a Troubleshooting Common Issues table that follows in Appendix E.

To report an issue that is still not resolvable, submit a service request (SR) to <http://support.oracle.com>. Please provide the following information along with your SR.

If the issue occurred using the CLI:

- Provide the command line output and all error messages output by the command.
- If the issue is due to a task failure, provide the task information and output.

If the issue occurred using the BUI:

- Indicate which user interface pane was being used, which table was displayed, and what table action was attempted.
- Describe any error dialog that appeared and what information was contained in it.
- If the issue is due to a task failure, provide the task information and output.

Additionally provide the following files from the software data directory:

- For Oracle Linux – `/var/opt/oracle/smu`
- For Oracle Solaris – `/var/opt/ORCLsmu`
- For Windows – `C:\ProgramData\Oracle\Oracle Snap Management Utility`
- `smu.log.N`
- `SmuService.log.N` (only present on Windows hosts)

SMU Restrictions Quick Reference Table

The following table provides a quick reference reminder of the restrictions that are applicable to each SMU operation. Failing to understand and follow any of these restrictions can cause the related operation to fail.

OPERATION	RESTRICTIONS
Clone copy	<ul style="list-style-type: none"> • Clone copy of ASM/iSCSI LUN database to a different storage appliance is not supported. • For ASM/iSCSI LUN databases, the iSCSI initiators of the clone target database host must be configured in the same iSCSI Initiator Group to which the clone source database host is mapped. • Remote Replication feature must be enabled on Oracle ZFS Storage Appliance. • Clone copy cannot be used for a source database whose shares span across the storage cluster heads. • However, clone copying from a single head to a target storage cluster is permitted, and the shares of a target database can be placed across the heads. • Replication target(s) pointing to the target storage (cluster head) must be configured on the source storage (cluster head). • The specified source project and target pool specified for the clone copy must be mapped in a 1-to-1 relationship.
Deprovision snap clone	<p>The target database host must have the same version of Oracle software as the source database host; no upgrades are performed. Upgrading Oracle software on the clone will disable deprovisioning of the clone. The database compatibility of the clone sets to the version of Oracle software installed on the target database host. The database must be in the mounted state so that SMU can obtain the required information on it. No listener, tnsname or enterprise manager are configured for the clone.</p>
Refresh clone	Can only be performed on thin clone.
Snap backup	<p>Offline, online or standby backups only. Offline backups allow one or more shares with database files spread across them in any fashion. Online backups require datafiles and archived logs be in separate shares; during an online backup the datafile shares are snapped first, the online logs are then archived, and finally the archived log shares are snapped. Online backups are only supported with filesystem (NFS or dNFS) storage type.</p>

OPERATION	RESTRICTIONS
	<p>ASM database only supports offline backup; it is not possible to separately snap datafiles and archived log files that are in a diskgroup.</p> <p>When creating an offline backup of a clustered database, any database nodes that have been stopped but not disabled will be restarted at the end of the backup task when the software restarts the database.</p>
Snap clone from RMAN backup	<ul style="list-style-type: none"> • The backup must be in image copy format. • Only filesystem (NFS, dNFS) storage type clones are supported. • The backup must contain appropriate files (controlfile, datafiles and archivelogs). Files must be in %U format. • Only cloning from hot (online) backup is supported. • The controlfile must have appropriate rows in the v\$datafile_copy about each datafile (change_check# must be correct). • Clone shares are placed in the same project as the backup shares. • The target database host must have the same version of Oracle software as the source database host; no upgrades are performed. • The target database's host oracle user must have the same uid as the source database's host oracle user. • No listener, tnsname or enterprise manager are configured for the clone. • You must specify Oracle home for clone database to use (a path like /u01/app/oracle/product/11.2.0/dbhome_1). The Oracle home must not end with a "/" (a slash character). • SMU must be able to mount a copy of the backup controlfile in order to query key system views for information about the backup. This requires starting up a temporary instance that uses the original database name, so there cannot be another database mounted on the target database host that uses the same name. • The clone will use the default settings for the memory-related parameters (ASMM memory model and 2GB SGA size). • The selected share(s) must contain a single RMAN backup set to be used in the clone operation. No other RMAN files or Oracle files may reside on the share(s). • The target host(s) must have data path connectivity to the Oracle ZFS Storage Appliance.

OPERATION	RESTRICTIONS
Snap clone from snap backup	<ul style="list-style-type: none"> • Clones are exact copies of the origin/source; use the same memory model and memory sizes. • Clone shares are placed in the same project as original shares. • Clone LUNs are placed in the same initiator and target groups as the origin shares. • The target database's host oracle user must have the same uid and gid as the source database's host Oracle user. • The clone will use the same memory model as the original database. • If dNFS is desired for the clone database, it must be configured in the target Oracle home before the clone operation is performed.
Snap recover	<ul style="list-style-type: none"> • Not supported on Windows dNFS or CDB databases. • Cannot be performed on Data Guard configurations.
Snap restore	<ul style="list-style-type: none"> • Can only restore to backup if newer backups do not have any clones.
Standby clone	<ul style="list-style-type: none"> • Primary database must be managed by SMU; SMU should be able to access primary database for snap backup and modification of configuration. • Avoid manually interaction. • No pre-existing standby database(s) should exist. SMU may fail to create/ mistakenly remove standby redo log for primary database. • No preconfigured Data Guard configuration; SMU will fail to create DG configuration for user. • If the listener to be used for standby has been created on host, its information must be added correctly to <code>\$TNS_ADMIN/network/admin/listener.ora</code> or <code>endpoint_listener.ora</code> (for RAC database). • Cross-platform standby is not supported. • Standby on Windows host is not supported. • Clone from Single Instance to RAC or RAC One Node is not supported. • Primary database should be read write open mode and archive log mode. • If primary database has protection mode as maximum availability or maximum protection, SMU will fail. Remove the last standby with <code>LogXptMode=SYNC</code> from the Data Guard configuration. SMU will clean up standby from host; the user must change protection mode of the Data Guard configuration and remove that standby from configuration manually.

Appendix A: References

- *Oracle ZFS Storage Appliance Administration Guide*
http://download.oracle.com/docs/cd/E22471_01/index.html
- Oracle Snap Management Utility for Oracle Database full documentation set
http://docs.oracle.com/cd/E39520_01/index.html
- My Oracle Support (MOS)
<http://support.oracle.com>
- Notes on RMAN Cloning, search for Doc ID 1210656.1
My Oracle Support (MOS)
<http://support.oracle.com>
- Notes on Time Settings with Oracle Database, found in My Oracle Support (MOS):
 - Doc ID 340512.1 (Timestamps & time zones - Frequently Asked Questions)
 - Doc ID 1627439.1 (How to Diagnose Wrong Time (SYSDATE and SYSTIMESTAMP) After DST Change, Server Reboot, Database Restart or Installation When Connecting to a Database on a Unix Server)
 - Doc ID 1209444.1 (How to Change Timezone for 11gR2 Grid Infrastructure)
- *Oracle® Database 12c Release 1 (12.1)*
http://www.oracle.com/pls/db121/portal.portal_db?selected=4&frame=
- *Oracle® Database Backup and Recovery User's Guide, 12c Release 1 (12.1)*
http://docs.oracle.com/cd/E16655_01/backup.121/e17630/toc.htm
- *Oracle® Database Backup and Recovery User's Guide, 11g Release 2 (11.2)*
http://docs.oracle.com/cd/E11882_01/backup.112/e10642/toc.htm
- *Oracle® Database Platform Guide*
11g Release 2 (11.2) for Microsoft Windows
http://docs.oracle.com/cd/E11882_01/win.112/e10845/architec.htm
- *Oracle® Database Installation Guide*
11g Release 2 (11.2) for Linux
http://docs.oracle.com/cd/E11882_01/install.112/e24321/toc.htm
- *Creating Files on a NAS Device for Use with Oracle Automatic Storage Management*

- Sudo authorization delegation tool information and download
<http://www.sudo.ws>
- Oracle Solaris 10-compatible Sudo tool available on Oracle Solaris Companion CD
<http://www.sunfreeware.com>

Appendix B: Glossary

Ancestor incarnation	The parent of a parent incarnation of a database is an ancestor incarnation. The parent incarnation is the database incarnation from which the current incarnation branched following an OPEN RESETLOGS operation. Any parent of an ancestor incarnation is also an ancestor incarnation.
ASM	Automatic Storage Management. A type of filesystem that organizes database files into disk groups for ease of management, as ASM automates and manages the underlying database files.
CDB	Container Database, used with multitenant container database, which is a database that can hold numerous pluggable user databases.
Clone	A clone is an instantaneously created, read-writable copy of a snap backup. One or more clones can be created from a single snap backup. Clones are presented to users as a normal filesystem. The usual operations can be performed on clones. Clones are typically used in a test, development, QA, or backup environment.
Clone copy	A clone copy is a "full" database clone that holds its own data shares independent of the original database. Clone copy can create a clone on pools of different storage from the source database.
Cluster	Multiple interconnected computers or servers that appear as if they are one server to end users and applications.
Cold backup	See proper term, offline backup.
Dependency	Applications that share the same project are considered dependent on each other.
Deprovision	The process of changing the state of a storage asset (share) as a usable resource for an application to unavailable.
dNFS	direct Network File System. An NFS client that optimizes I/O on Network Attached Storage (NA) devices.
FRA	Flash Recovery Area. A configured area of disk storage where backup components such as datafile image copies, archive logs, and controlfiles are held.
Hot backup	See proper term, online backup.
iSCSI	Internet Small Computer System Interface. A protocol that allows data packets to be transmitted using TCP/IP.
Multitenant	A new option for Oracle Database 12c, Oracle Multitenant delivers a new architecture that allows a multitenant container database to hold many pluggable databases. An existing database can simply be adopted with no application changes required. Oracle Multitenant fully complements other options, including Oracle Real Application Clusters and Oracle Active Data Guard.
NFS	Network File System. A filesystem protocol used in Network Attached Storage (NAS) systems that allows the sharing of files, based on access privileges, among remote clients and the primary server.
Offline backup	Offline (also called cold) backups are backups taken when the database is shut down. The software will shut down the database temporarily and then restart it after taking the snap backups.
Online backup	Also referred to as hot backup. Online backups are taken when the database is placed into backup mode while remaining online. Online backups take snap backups of the database shares in a particular order and in between changing the database mode and archiving the current logs.

Project	An Oracle ZFS Storage Appliance project defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. Projects are also used to group logically related shares together so their common attributes can be accessed from a single point. All filesystems and LUNs are grouped into projects. Typically, every application has its own project.
RAC	Oracle Real Application Clusters. A tool enabling the clustering of Oracle databases.
Recover	Referred to as database recover or snap recover, this operation can restore a database to a point in time in between backups. Recovery can be time-based, change based (using a system change number), or by log sequence number.
Refresh clone	A refresh clone updates or refreshes the data of an existing clone to match the latest data contents of the source database backup from which it was originally created.
Remote replication	Remote replication is a feature of the Oracle ZFS Storage Appliance that facilitates replication of projects and shares to and from one Oracle ZFS Storage Appliance to another.
Replication action	The replication action is a configuration object on a source storage specifying a project or share, a target storage, synchronization policy, and so on. A replication action is a part of the remote replication process used in SMU's clone copy operation.
Replication package	The replication package is the configuration object, counterpart to the replication action described in the previous row, but used on target storage.
RMAN	A feature of the Oracle Database, Oracle Recovery Manager is a comprehensive tool for easily managing backup and recovery of Oracle Database and provides a common interface, either through a command line or Enterprise Manager, for backup tasks across different host operating systems.
Sever	A replication package can be converted into a local, writable project acting like a local project. Severing the replication connection of a replication package converts the package to a local project. The Clone Copy severs a replication package with a new project name.
Shares	Shares are filesystems and LUNs exported over supported data protocols to clients of the Oracle ZFS Storage Appliance. A share is created under a project. Filesystems export a file-based hierarchy and can be accessed over CIFS, NFS, HTTP/WebDAV and FTP. LUNs export block-based volumes and can be accessed through iSCSI.
SID	The Oracle System ID that uniquely identifies a particular database on a system. The variable for the identifier is ORACLE_SID.
Snap backup	A snap backup is a read-only, point-in-time copy of a filesystem, instantaneously created with no space allocated initially. Blocks are allocated as changes are made to the base filesystem (copy on write). Snap backup data can be directly accessed for backup purposes, Snap backups are initiated either manually or through automated scheduling at specified intervals, Any reads to the snap backup blocks are served by the base filesystem's block. As the changes occur to the base filesystem, the older block referenced by the snap backup and the new, changed block are referenced by the filesystem. A project-level snap backup is the same as taking snap backups on all the shares within the project.
Standby backup	A database backup that is designated as the type standby when it is created. This is the first step or process in creating a standby clone that is used as a data guard database.
Standby clone	A clone that is created from a database backup of type standby. This clone is then designated in the clone operation as a type data guard standby clone.
Storage pool	A storage pool is created among a set of physical disks. Filesystems are then created over the storage pool. One or more storage pools are created over the available physical disks.

Sudo	A privilege authorization program that allows a substitute user ("su") to execute programs ("do") using another user's security privileges, often the root user. The substitute user establishes a password connected to the desired privileged access that sudo verifies from its configuration file, then grants the requested access. Sudo was developed for Unix-like operating systems.
Thin clone	A clone of a snap backup which is created using Oracle ZFS Storage Appliance clone technology. A thin clone is created on the same storage pool as the source snap backup (called a clone point) and shares base data blocks with the base share and the source snap backup. As such, a thin clone has data dependencies to the source snap backup and the base share. Neither the base share or the source snap backup can be deleted unless the thin clone is removed.
Wallet file	An Oracle wallet file stores the master encryption key to an Oracle Database protected using transparent data encryption, which encrypts sensitive table data in the datafiles. Without access to the Oracle wallet, which is stored outside of the database, the database table data cannot be read and/or copied.

Appendix C: Icon Set for the SMU Browser User Interface

The following table lists all the icons used in the Browser User Interface and their associated meanings. Mousing over any of these icons produces a display of that icon's function.




































Icon	Action or entity it represents	Further information
	Application and Accounts details image icon	Image icon for Application and Accounts (Application Host and Storage) details.
	Add action	Operation for Add Application, Add Backup, Add Schedule, Add Storage, Add Host, Add Notification, Add User.
	Administrator image icon	Administration node in the navigation tree.
	Clone creation icon	Operation for Create Primary clone.
	Clone copy	Identifying clone copy node in the navigation tree.
	Refresh clone	Operation for Refresh clone.
	Database Application image icon	Application node in the navigation tree.
	Delete action	Operation for Remove Application, Delete task from Tasks queue, Delete Backup, Remove Schedule, Remove Host, Remove Storage, Remove user, Remove Notification.
	Edit/update action	Operation for Modify Application, Rename Backup, Modify Schedule, Modify Host, Modify Storage, Modify user, Modify Notification.
	Application root image icon	Image icon of the applications root node in the navigation tree.
	Export action	Operation for Export Activity Logs in the activity logs panel.
	File import action	Operation for Import RMAN Backup Image in Applications panel.
	Filter action	Operation for Filter Activity Logs in Activity Logs panel.

TABLE 28. ICON SET FOR SMU BROWSER USER INTERFACE		
Icon	Action or entity it represents	Further information
	User image icon	Login User image icon in upper right header area of the main page.
	Help for a topic	Topic helper next to each field in a entry form panel.
	Running Tasks display option	One of the options in upper right of Task panel for all the Running Tasks display. Also used as a task status indicator in the task table.
	Activity Logs image icon	Activity Logs node image icon in the navigation tree.
	Canceled Tasks display option	One of the options in upper right of Task panel for the Canceled Tasks display. Also used as a task status indicator in the task table.
	Failed Tasks display option	One of the options in upper right of Task panel for all the Failed Tasks display. Also used as a task status indicator in the task table.
	Pending Tasks display option	One of the options in upper right of Task panel for all the Pending Tasks display. Also used as a task status indicator in the task table.
	Succeeded Tasks display option	One of the options in upper right of Task panel for all the successful Tasks display. Also used as a task status indicator in the task table.
	Recover Backup action	Operation for Recover Backup in the Application panel.
	Restore Backup action	Operation for Restore Backup in the Application panel.
	Refresh panel/table action	Operation for Refresh Applications Table, Refresh Backups Table, Refresh Tasks Table.
	Deprovision Application action; Purge Activity Logs action	Operation for Deprovision Application in Applications panel and Purge Activity Logs in Activity Logs panel.
	Shuttle down (disabled is gray); shuttle down (enabled is blue)	Shuttle down operation for page navigation in Activity Logs panel and Backups panel.

Icon	Action or entity it represents	Further information
	Shuttle to the leftmost position (disabled: gray); (enabled: blue)	Shuttle to the leftmost position operation for page navigation in Activity Logs panel and Backups panel.
	Shuttle right (disabled: gray); (enabled: blue)	Shuttle right operation for page navigation in Activity Logs panel and Backups panel.
	Shuttle to the rightmost position (disabled: gray); (enabled: blue)	Shuttle to the rightmost position operation for page navigation in Activity Logs panel and Backups panel.
	Shuttle up (disabled is gray); shuttle up (enabled is blue)	Shuttle up operation for page navigation in Activity Logs panel and Backups panel.
	Shuttle left (enabled: blue); (disabled: gray)	Shuttle left operation for page navigation in Activity Logs panel and Backups panel.
	Cancel Task action	Operation for Cancel Task in Tasks table.
	All Tasks display option	One of the options in the upper right of Task panel for All Tasks display.
	Test Application action	Operation for Test Application to test for valid, accessible database.
	Thin clone image icon	Thin clone node image icon in the navigation tree.

Appendix D: Cloning Wallet Files for an Encrypted Database

When an Oracle Database is encrypted using the transparent data encryption feature, a random key that serves as the master encryption key is generated and stored in an Oracle wallet. This Oracle wallet file must be copied to a designated target host before its associated database instance can be cloned to that target. Use the following procedure to clone the wallet file to the target.

Perform the following configuration and verification steps on the target host.

1. Add wallet information into the `sqlnet.ora` file on the remote host.

```
[oracle@aie-4200f admin]$ vi sqlnet.ora
[oracle@aie-4200f admin]$ cat sqlnet.ora
ENCRYPTION_WALLET_LOCATION=
(SOURCE=(METHOD=FILE)(METHOD_DATA=
(DIRECTORY=/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/)))

[oracle@aie-4200f admin]$ mkdir -p
/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/
```

2. Clone the wallet files to the target host. If you are using the encryption wallet, you only have a .p12 file.

```
[oracle@aie-4200f admin]$ scp oracle@aie-
4200x.us.oracle.com:/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/* .

oracle@aie-4200x.us.oracle.com's password:

cwallet.sso
 100% 3499    3.4KB/s   00:00

ewallet.p12
 100% 3421    3.3KB/s   00:00

[oracle@aie-4200f admin]$ ls

cwallet.sso ewallet.p12 samples shrept.lst sqlnet.ora tnsnames.ora tnsnames.ora.bak

[oracle@aie-4200f admin]$ mv cwallet.sso
/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/

[oracle@aie-4200f admin]$ mv ewallet.p12
/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/
```

3. Log in to the clone instance and check the wallet information.


```
[oracle@aie-4200f admin]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 11.2.0.3.0 Production on Tue Jun 4 23:03:14 2013
Copyright (c) 1982, 2011, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.3.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options
```

```
SQL> startup force
ORACLE instance started.
```

```
Total System Global Area 1269366784 bytes
Fixed Size 2227984 bytes
Variable Size 754974960 bytes
Database Buffers 503316480 bytes
Redo Buffers 8847360 bytes
```

```
Database mounted.
Database opened.
```

```
SQL> select * from v$encryption_wallet;
```

```
WRL_TYPE
```

```
WRL_PARAMETER
```

```
STATUS
```

```
file
```

```
/u01/app/oracle/product/11.2.0/dbhome_1/encryption_wallet/
OPEN
```

4. Verify that the encrypted column can be seen.

```
SQL> conn scott/tiger
```

```
Connected.
```

```
SQL> desc emp1;
```

Name	Null?	Type
EMPNO		NUMBER(4)
ENAME		VARCHAR2(10)
JOB		VARCHAR2(9)
MGR		NUMBER(4)
HIREDATE		DATE
SAL		NUMBER(7,2) ENCRYPT
COMM		NUMBER(7,2)
DEPTNO		NUMBER(2)

```
SQL> select sal from emp1;
```

```

SAL
-----
800
1600
1250
2975
```

```
1250
2850
2450
3000
5000
1500
1100
  SAL
-----
   950
  3000
  1300
14 rows selected.

SQL>
```

Appendix E: Troubleshooting Common Issues

The following issues have also been reported and listed as MOS notes under Doc ID 1522925.1, Snap Management Utility for Oracle Database – Information and Troubleshooting, on the My Oracle Support web site listed in the References section. If you do not see a listing for your issue, check the MOS notes for the latest updates to the list.

ISSUE DESCRIPTION	INFORMATION/RESOLUTION
ORA-27102: out of memory while cloning a database.	If you receive ORA-27102 when cloning a snap backup or RMAN backup, there is not enough shared memory available for the database clone. Either add more shared memory or delete other databases that are running on the host or cluster. Alternatively, consider cloning the database to another host or cluster that has available space.
ORA-01034: ORACLE not available during snap backup	Before creating a snap backup of the database, SMU must connect to and query the database for vital information including the list of files the database is using. This error indicates that the database instance that SMU tried to connect to is shut down and not running. Restart the database instance so SMU can operate correctly. If this is a RAC database you can also modify the host account for the database to use one of the other cluster nodes that is up and running.
Cluster nodes that were down before offline snap backup are up and running afterward.	SMU uses the <code>srvctl stop</code> and <code>srvctl start</code> commands to temporarily shut down the database when performing an offline database. For a cluster database, if some nodes were down before the backup, they will be brought up after the backup. If you want a particular cluster node to remain down after the backup, you must use the <code>srvctl disable</code> command to disable the node so that the <code>srvctl start</code> command will not restart the node.
"Auth fail" or "HTTP 401" when performing a task.	When executing a task, SMU logs in to one or more host and storage systems. If the user and/or password for the account is incorrect, the task will fail, displaying either "Auth fail" if the account is for a Linux or Oracle Solaris storage system, or an "HTTP 401" error if the account is for a Windows system. Test account settings prior to executing tasks to help ensure that the account settings are correct. Use the <code>accounts test</code> command from the CLI or click the test button in the column of the account for which you want to check settings.
SMU failed on first mount of the clone of an RMAN backup – permission denied.	FFAS: Error creating clone using Oracle Snap Management Utility Mount failed, reason given by server: Permission denied . If root exceptions for the host are enabled using an invalid CIDR, this can be the result. Fixing the CIDR should enable the mount to succeed. Example invalid exception: <code>root=@192.168.10.100/20..</code>

ISSUE DESCRIPTION	INFORMATION/RESOLUTION
Invalid application file layout. Remote share X has already been backed up.	This error indicates that the database file layout does not allow an online backup operation on this database. In order to take an online backup of the database, the datafiles and archived logs must reside in separate shares. During an online backup, snapshots of the datafile shares are taken first while the database is in backup mode. Next, the current redo logs are archived. And then snapshots of the archived log shares are taken. If, during the backup sequence, SMU detects that shares have already had snaps taken of them, it will fail the online backup task and display this error.
host X login: timeout: socket is not established.	This error indicates that the software could not connect to the database host (Linux or Oracle Solaris) or storage appliance. This error occurs when the database host or storage appliance are not reachable or do not respond to the connection request within a timeout period. Verify that the database host or storage appliance are up and reachable over the network and try the operation again.
Could not find all shares or shares were unavailable due to pool status.	This error message can occur during a task when SMU searches for the shares to operate on. A key feature SMU provides is the ability to map shares from their external attributes (mountpoint or lunguide) to their internal appliance identifier (pool/collection/project/share). This error message indicates that the shares SMU was looking for either do not exist on the Oracle ZFS Storage Appliance or are not available because the storage pool they are in is in a degraded state or other than online state. You can encounter this error if you specify the wrong storage account with a database account. In particular this error will occur when using ASM databases and the wrong storage account is specified for the database. SMU cannot determine which external storage system an iSCSI LUN is using and so will only search the storage that was linked to by the database account.
ORA-19809 occurs when creating a snap clone database.	This error message indicates that the size specified for the flash recovery area (FRA) is too small to support the clone database. SMU sets the size of the FRA for the clone database based on the <code>db_recovery_file_dest_size</code> initialization parameter of the database that was backed up. Since a snap clone is an identical copy of the source database, including the size of each redo log, the size parameter should be adequate for the clone database. To address this issue, create the clone database from a backup of the source database that has a suitable FRA size.

ISSUE DESCRIPTION	INFORMATION/RESOLUTION
<p>The WS-Management service cannot process the request because the request contained invalid selectors for the resource.</p>	<p>This error occurs with Windows hosts when the shell session that SMU established has been idle too long. SMU uses Windows Remote Shell to connect to the host and establish a session. Windows Remote Shell will automatically log the session out if the idle timeout period expires. The software alternates issuing commands to the host and storage. It is possible for the host session to be idle while SMU sends commands to the storage.</p> <p>To resolve this, increase the WinRS idle timeout period to a larger value (the software requires the timeout period to be 2 hours or more).</p> <pre>C:\>winrm set winrm/config/winrs @{IdleTimeout="7200000"}</pre>
<p>BUI always displays fetching data or displays it frequently.</p>	<p>The BUI is designed to refresh itself regularly so that it can display the current status of tasks and other items in the various UI panes. The amount of data to display can grow over time. Completed tasks are retained until removed or deleted by the user. It is possible to accumulate a large number of tasks that prevent the BUI from refreshing its display properly. Either delete completed tasks that are no longer needed or disable the UI refresh by modifying the global refresh settings.</p>
<p>The WS-Management service cannot process the request. The maximum number of concurrent operations for this user has been exceeded. Close existing operations for this user, or raise the quota for this user.</p>	<p>This error indicates that the WinRM setting <code>MaxConcurrentOperationsPerUser</code> is set too low. The software's recommended value for this setting is 1500. The software executes many SQL Plus, RMAN and system commands on the host while performing operations, greater than the number of commands allowed by default.</p> <p>To modify this setting, run the following command:</p> <pre>C:\>winrm set winrm/config/service @{MaxConcurrentOperationsPerUser="1500"}</pre>
<p>Clone database task hangs when target database host is Linux running UEK kernel and dNFS is enabled in the target Oracle home.</p>	<p>This error occurs when the UEK kernel 2.6.32-300.11.1.el5uek is running on the database host. More information on this issue is available in Doc ID 1460787.1. To resolve, you must upgrade your kernel to 2.6.32-300.26.1 or disable dNFS. The software does not support the use of an <code>oranfstab</code> file in this release.</p>
<p>Clone database task hangs during control file creation on Windows database host,</p>	<p>This error occurs because the WinRM setting <code>MaxTimeoutms</code> is set too low. Some of the commands SMU runs can take a while to complete. The software requires that this parameter be set to a value high enough to allow these long-running commands to complete. This error can also be verified by examining the software log for an exception like the following: "Exception in thread: "Thread-416" javax.xml.ws.soap.SOAPFaultException: The WS-Management service cannot complete the operation within the time specified in</p>

ISSUE DESCRIPTION	INFORMATION/RESOLUTION
	<p>OperationTimeout."</p> <p>To address the issue, modify the setting of the MaxTimeoutms setting and run the clone task again.</p> <pre>C:\>winrm set winrm/config @{MaxTimeoutms="7200000" }</pre>
<p>No rows for datafile X in v\$datafile_copy system view.</p>	<p>This error occurs during an RMAN clone operation when SMU cannot find a row for one of the backed-up data files in the backed-up control file v\$datafile_copy system view. This indicates that the data file is not a part of the backup set or the control file was backed up before the data file was backed up. SMU requires that each data file in the backup has a row in this system view so that it can calculate the maximum SCN (system change number) to recover the database to. To resolve this issue, create a valid backup that contains the data files, archived logs and control file.</p>
<p>Database already in backup mode.</p>	<p>This error indicates that a database that SMU attempted to back up was already in backup mode. Another process or program could be backing up the database. SMU will not back up a database that is already in backup mode. If the database is not being backed up by other utilities, then the database must be taken out of backup mode manually before SMU can successfully back up the database.</p>
<p>Cannot map disk <hostname ip address>:<lun guid></p>	<p>This error indicates that SMU could not find the clone disk on the database host or node. When cloning an ASM database, SMU clones the appropriate snapshot on the appliance to create new LUNs. SMU then uses operating system-specific commands to discover the clone LUNs from the database host or node. If the clone LUNs cannot be discovered, then SMU reports this error. The usual cause is that the appropriate iSCSI targets have not been logged into the database host or node. The software performs no SAN configuration and requires that all iSCSI targets be logged into before any ASM database clone operations are performed.</p>
<p>Does SMU need to communicate with Recovery catalog at main site for deploying Dev. DB with image copy (Clone) at DR site?</p>	<p>SMU does not currently use the RMAN catalog. SMU requires that the backup shares contain a single full image copy backup. SMU will scan the backup shares and identify the control files, data files and archive logs in the backup shares. It will then snap and clone the backup shares and mount the clone shares on the target database host and proceed to configure and start a clone database that uses the files as is.</p>
<p>If the source database is RAC, should the target DB be RAC as well?</p>	<p>No. SMU can detect if the backup is of a single instance or RAC database and will perform the appropriate processing on the clone based on whether it is targeted for a single instance or RAC environment. In other words you can create single instance or RAC clones no matter if the source is single instance or RAC.</p>

ISSUE DESCRIPTION	INFORMATION/RESOLUTION
<p>Multiple operating systems (Redhat/IA Linux, SPARC/Oracle Solaris) exist at the source side. Should the same platforms be prepared for target databases?</p>	<p>In general, we recommend having the same platform. However, SMU can migrate a clone database from one operating system to another as long as both are the same Endian architecture. This means if the source database is Linux/IA you can create a clone on an Oracle Solaris/IA host. You cannot clone from little endian to big endian. So if your production environment is using SPARC/Oracle Solaris, you can only create clones on a SPARC/Oracle Solaris host.</p>
<p>Should the backup format be image copy? Are the files data and control needed?</p>	<p>Yes, the backup must be in image copy format. See "RMAN Clone" section in table 28, Quick Reference for SMU Operations Restrictions. Backup files must use the RMAN %U format specification and must contain only one control file, one or more data files and one or more archived logs.</p> <p>Note the sample RMAN export runblock:</p> <pre>run { set nocfau; # back up the control file explicitly allocate channel ch1 device type disk format '/backup/smu/%U'; allocate channel ch2 device type disk format '/backup/smu/%U'; allocate channel ch3 device type disk format '/backup/smu/%U'; allocate channel ch4 device type disk format '/backup/smu/%U'; backup as copy database plus archivelog; backup as copy current controlfile; }</pre>
<p>WARNING: A Bean Validation provider is not present, therefore bean validation @ is disabled</p>	<p>This warning displays in the <code>/var/opt/ORCLsmu/smu.log</code> but has nothing to do with SMU functionality; rather, it is an ADF informational notice. It can be ignored when searching the log in troubleshooting.</p>







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2013, 2014, 2015, 2016 Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0615

Oracle Snap Management Utility for Oracle Database, v1.3.0 User Guide
January 2016

Oracle ZFS Storage Appliance Software Solutions



Oracle is committed to developing practices and products that help protect the environment: